The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation, or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR: So last lecture, we laid out the foundations of quantum stat mech. And to remind you, in quantum mechanics we said we could very formally describe a state through a vector in Hilbert space that has complex components in terms of some basis that we haven't specified. And quantum mechanics describes, essentially what happens to one of these states, how to interpret it, et cetera.

But this was, in our language, a pure state. We were interested in cases where we have an ensemble and there are many members of the ensemble that correspond to the same macroscopic state. And we can distinguish them macroscopically. If we think about each one of those states being one particular member, psi alpha, some vector in this Hilbert space occurring with some probability P alpha, then that's kind of a more probabilistic representation of what's going on here for the mech state.

And the question was, well, how do we manipulate things when we have a collection of states in this mixed form? And we saw that in quantum mechanics, most of the things that are observable are expressed in terms of matrices. And one way to convert a state vector into a matrix is to conjugate it and construct the matrix whose elements are obtained-- a particular rho vector by taking element and complex conjugate elements of the vector.

And if we were to then sum over all members of this ensemble, this would give us an object that we call the density matrix, rho. And the property of this density matrix was that if I had some observable for a pure state I would be able to calculate an expectation value for this observable by sandwiching it between the state in the way that quantum mechanic has taught us.

In our mech state, we would get an ensemble average of this quantity by multiplying this density matrix with the matrix that represents the operator whose average we are trying to calculate, and then taking the trace of the product of these two matrices.

Now, the next statement was that in the same way that the micro-state classically changes as a function of time, the vector that we have in quantum mechanics is also changing as a function of time. And so at a different instant of time, our state vectors have changed. And in principle, our density has changed. And that would imply a change potentially in the average that we have over here.

We said, OK, let's look at this time dependence. And we know that these psis obey Schrodinger equation. I h bar d psi by dt is h psi. So we did the same operation, I h bar d by dt of this density matrix. And we found that essentially, what we would get on the other side is the commutator of the density matrix and the matrix corresponding to the Hamiltonian, which was reminiscent of the Liouville statement that we had classically.

Classically, we had that d rho by dt was the Poisson bracket of a rho with H versus this quantum formulation.

Now, in both cases we are looking for some kind of a density that represents an equilibrium ensemble. And presumably, the characteristic of the equilibrium ensemble is that various measurements that you make at one time and another time are the same.

And hence, we want this rho equilibrium to be something that does not change as a function of time, which means that if we put it in either one of these equations, the right-hand side for d rho by dt should be 0. And clearly, a simple way to do that is to make rho equilibrium be a function of H.

In the classical context, H was a function in phase space. Rho equilibrium was, therefore, made a function of various points in phase space implicitly through its dependence on H. In the current case, H is a matrix in Hilbert space. A function of a matrix is another function, and that's how we construct rho equilibrium. And that function will also commute with H because clearly H with H is 0. Any function of H with H will be 0.

So the prescription that we have in order now to construct the quantum statistical mechanics is to follow what we had done already for the classical case. And classically following this equation, we made postulates relating rho-- what its functional dependence was on H. We can now do the same thing in the quantum context. So let's do that following the procedure that we followed for the classical case.

So classically, we said, well, let's look at an ensemble that was micro-canonical. So in the micro-canonical ensemble, we specified what the energy of the system was, but we didn't allow either work to be performed on it mechanically or chemically so that the number of elements was fixed. The coordinates, such as volume, length of the system, et cetera, were fixed.

And in this ensemble, classically the statement was that rho indexed by E is a function of H. And clearly, what we want is to only allow H's that correspond to the right energy. So I will use a kind of shorthand notation as kind of a delta H E.

Although, when we were doing it last time around, we allowed this function to allow a range of possible E-values. I can do the same thing, it's just that writing that out is slightly less convenient. So let's stick with the convenient form. Essentially, it says within some range delta A allow the Hamiltonians. Allow states, micro-states, whose Hamiltonian would give the right energy that is comparable to the energy that we have said exists for the macro-state.

Now clearly also, this definition tells me that the trace of rho has to be 1. Or classically, the integral of rho E over the entire phase space has to be 1. So there is a normalization condition. And this normalization condition gave us the quantity omega of E, the number of states of energy E, out of which we then constructed the entropy and we were running away calculating various thermodynamic quantities. So now let's see how we evaluate this in the quantum case. We will use the same expression, but now me realize that rho is a matrix.

So I have to evaluate, maybe elements of that matrix to clarify what this matrix looks like in some basis. What's the most convenient basis?

Since rho is expressed as a function of H, the most convenient basis is the energy basis, which is the basis that diagonalizes your Hamiltonian matrix. Basically, there's some vectors in this Hilbert space such that the action of H on this will give us some energy that I will call epsilon n, some eigenvalue. An n. So that's the definition of the energy basis.

Again, as all basis, we can make these basis vectors n to be unit length and orthogonal to each other. There is an orthonormal basis.

If I evaluate rho E in this basis, what do I find?

I find that n rho m. Well, 1 over omega E is just the constant. It comes out front. And the meaning of this delta function becomes obvious. It is 1 or 0. It is 1 if, let's say, Em equals to the right energy. Em is the right energy for the ensemble E.

And of course, there is a delta function. So there is also an m equals to n. And it is 0, clearly, for states that have the wrong energy. But there is an additional thing here that I will explain shortly for m not equal to n. So let's parse the two statements that I have made over here.

The first one, it says that if I know the energy of my macro-state, I clearly have to find wave functions to construct possible states of what is in the box that have the right energy. States that don't have the right energy are not admitted. States that are the right energy, I have nothing against each one of them. So I give them all the same probability. So this is our whole assumption of equal equilibrium a priori probabilities.

But now I have a quantum system and I'm looking at the matrix. And this matrix potentially has off-diagonal elements. You see, if I am looking at m equals to n, it means that I am looking at the diagonal elements of the matrix. What this says is that if the energies are degenerate-- let's say I have 100 states, all of them have the right energy, but they are orthonormal.

These 100 states, when I look at the density matrix in the corresponding basis, the off-diagonal elements would be 0. So this is the "or." So even if this condition is satisfied, even if Em equals to E, but I am looking at off-diagonal elements, I have to put 0. And this is sometimes called the assumption of random phases.

Before telling you why that's called an assumption of random phases, let's also characterize what this omega of E is. Because trace of rho has to be 1, clearly omega of E is the trace of this delta h E. Essentially as I scan all of my possible energy levels in this energy basis, I will get 1 for those that have the right energy and 0 otherwise. So basically, this is simply number of states of energy E, potentially with minus plus some delta E if I want to include that.

Now, what these assumptions mean are kind of like this. So I have my box and I have been told that I have energy E. What's a potential wave function that I can have in this box? What's the psi given that I know what the energy E is?

Well, any superposition of these omega E states that have the right energy will work fine. So I have a sum over, let's say, mu that belongs to the state such that H-- such that this energy En is equal to the energy that I have specified. And there are omega sub E such states. I have to find some kind of an amplitude for these states, and then the corresponding state mu. I guess I should write here E mu.

Now, the first statement over here is that as far as I'm concerned, I can put all of these a mu's in whatever proportion that I want. Any linear combination will work out fine. Because ultimately, psi has to be normalized. Presumably, the typical magnitude a m squared, if I average over all members of the ensemble, should be 1 over omega. So this is a superposition. Typically, all of them would contribute.

But since we are thinking about quantum mechanics, these amplitudes can, in fact, be complex. And this statement of random phases is more or less equivalent to saying that the phases of the different elements would be typically uncorrelated when you average over all possible members of this ensemble.

Just to emphasize this a little bit more, let's think about the very simplest case that we can have for thinking about probability. And you would think of, say, a coin that can have two possibilities, head or tail. So classically, you would say, head or tail, not knowing anything else, are equally likely.

The quantum analog of that is a quantum bit. And the qubit can have, let's say, up or down states. It's a Hilbert space that is composed of two elements. So the corresponding matrix that I would have would be a 2 by 2. And according to the construction that I have, it will be something like this. What are the possible wave functions for this system?

I can have any linear combination of, say, up and down, with any amplitude here. And so essentially, the amplitudes, presumably, are quantities that I will call alpha up and alpha down. That, on average, alpha squared up or down would be 1/2. That's what would appear here.

The elements that appear here according to the construction that I have up there, I have to really take this element, average it against the complex conjugate of that element. So what will appear here would be something like e to the i phi of up minus phi of down, where, let's say, I put here phi of up. And what I'm saying is that there are a huge number of possibilities. Whereas, the classical coin has really two possibilities, head or tail, the quantum analog of this is a huge set of possibilities. These phis can be anything, 0 to 2 pi. Independent of each other, the amplitudes can be anything as long as the eventual normalization is satisfied. And as you sort of sum over all possibilities, you would get something like this also.

Now, the more convenient ensemble for calculating things is the canonical one, where rather than specifying what the energy of the system is, I tell you what the temperature is. I still don't allow work to take place, so these other elements we kept fixed. And our classical description for rho sub T was e to the minus beta H divided by some partition function, where beta was 1 over kT, of course.

So again, in the energy basis, this would be diagonal. And the diagonal elements would be e to the minus beta epsilon n. I could calculate the normalization Z, which would be trace of e to the minus beta H. This trace is calculated most easily the basis in which H is diagonal. And then I just pick out all of the diagonal elements, sum over n e to the minus beta epsilon.

Now if you recall, we already did something like this without justification where we said that the states of a harmonic oscillator, we postulated to be quantized h bar omega n plus 1/2. And then we calculated the partition function by summing over all states. You can see that this would provide the justification for that.

Now, various things that we have in classical formulations also work, such that classically if we had Z, we could take the log of Z. We could take a derivative of log Z with respect to beta. It would bring down a factor of H. And then we show that this was equal to the average of the energy. It was the average of Hamiltonian, which we then would identify with the thermodynamic energy.

It is easy to show that if you take the same set of operations over here, what you would get is trace of H rho, which is the definition of the average that you find.

Now, let's do one example in the canonical ensemble, which is particle in box. Basically, this is the kind of thing that we did all the time. We assume that there is some box of volume v. And for the time being, I just put one particle in it. Assume that there is no potential. So the Hamiltonian for this one particle is just its kinetic energy p1 squared over 2m, maybe plus some kind of a box potential.

Now, if I want to think about this quantum mechanically, I have to think about some kind of a basis and calculating things in some kind of basis. So for the time being, let's imagine that I do things in coordinate basis. So I want to have where the particle is. And let's call the coordinates x. Let's say, a vector x that has x, y, and z-components.

Then, in the coordinate basis, p is h bar over i, the gradient vector. So this becomes minus h bar squared Laplacian over 2m. And if I want to look at what are the eigenstates of his Hamiltonian represented in coordinates basis, as usual I want to have something like H1 acting on some function of position giving me the energy.

And so I will indicate that function of position as x k. And I want to know what that energy is. And the reason I do that is because you all know that the correct form of these eigenfunctions are things like e to the i k dot x. And since I want these to be normalized when I integrate over all coordinates, I have to divide by square root of V. So that's all nice and fine.

What I want to do is to calculate the density matrix in this coordinate basis. So I pick some point x prime rho x. And it's again, a one particle. Yes?

AUDIENCE: I was thinking the box to be infinite or-- I mean, how to treat a boundary because [INAUDIBLE].

PROFESSOR: So I wasn't careful about that. The allowed values of k that I have to pick here will reflect the boundary conditions. And indeed, I will be able to use these kinds of things. Let's say, by using periodic boundary condition. And if I choose periodic boundary conditions, the allowed values of k will be quantized. So that, say, kx would be multiples of 2 pi over Lx times some integer.

If I use really box boundary conditions such as the wave function having to vanish at the two ends, in reality I should use the sine and cosine type of boundary wave functions, which are superpositions of these e to the i kx and e to the minus i kx. And they slightly change the normalization-- the discretization.

But up to, again, these kinds of superpositions and things like that, this is a good enough statement. If you are really doing quantum mechanics, you have to clearly be more careful. Just make sure that I'm sufficiently careful, but not overly careful in what I'm doing next. So what do I need to do?

Well, rho 1 is nicely diagonalized in the basis of these plane waves. In these eigenstates of the Hamiltonian. So it makes sense for me to switch from this representation that is characterized by x and x prime to a representation where I have k. And so I can do that by inserting a sum over k here. And then I have rho 1.

Rho 1 is simply e to the minus beta. It is diagonal in this basis of k's. There is a normalization that we call Z1, partition function. And because it is diagonal, in fact I will use the same k on both sides of this. So the density matrix is, in fact, diagonal in momentum. This is none other than momentum, of course.

But I don't want to calculate the density matrix in momentum where it is diagonal because that's the energy basis we already saw, but in coordinate basis. So now let's write all of these things.

First of all, the sum over k when l becomes large, I'm going to replace with an integral over k. And then I have to worry about the spacing that I have between k-values. So that will multiply with a density of states, which is 2 pi cubed product of lx, ly, lz, which will give me the volume of the box.

And then I have these factors of-- how did I define it?

kx is e to the i k dot x. And this thing where k and x prime are interchanged is its complex conjugate. So it becomes x minus x prime.

I have two factors of 1 over square root of V, so that will give me a factor of 1 over V from the product of these things. And then I still have minus beta h bar squared k squared over 2m. And then I have the Z1.

Before proceeding, maybe it's worthwhile for me to calculate the Z1 in this case. So Z1 is the trace of rho 1. It's essentially the same thing. It's a sum over k e to the minus beta h bar squared k squared over 2m. I do the same thing that I did over there, replace this with V integral over k divided by 2 pi cubed e to the minus beta h bar squared k squared over 2m.

I certainly recognize h bar k as the same thing as momentum. This is really p squared. So just let me write this after rescaling all of the k's by h bar. So then what I would have is an integral over momentum. And then for each k, I essentially bring down a factor of h bar. What is 2 pi h bar?

2 pi h bar is h. So I would have a factor of h cubed. I have an integration over space that gave me volume. e to the minus beta h bar-- p squared over 2m. Why did I

write it this way?

Because it should remind you of how we made dimensionless the integrations that we had to do for classical calculations. Classical partition function, if I wanted to calculate within a single particle, for a single particle Hamiltonian, I would have this Boltzmann weight. I have to integrate over all coordinates, over all momenta. And then we divided by h to make it dimensionless. So you can see that that naturally appears here. And so the answer to this is V over lambda cubed with the lambda that we had defined before as h over root 2 pi m kT.

So you can see that this traces that we write down in this quantum formulations are clearly dimensionless quantities. And they go over to the classical limit where we integrate over phase space appropriately dimensional-- made dimensionless by dividing by h per combination of p and q.

So this V1 we already know is V-- Z1 we already know is V over lambda cubed. The reason I got a little bit confused was because I saw that the V's were disappearing. And then I remembered, oh, Z1 is actually proportional to mu over lambda cubed. So that's good. So what we have here from the inverse of Z1 is a 1 over V lambda cubed here.

Now, then I have to complete this kind of integration, which are Gaussian integrations. If x was equal to x prime, this would be precisely the integration that I'm doing over here. So if x was equal to x prime, that integration, the Gaussian integration, would have given me this 1 over lambda cubed. But because I have this shift-- in some sense I have to shift k to remove that shift in x minus x prime. And when I square it, I will get here a factor, which is exponential, of minus x minus x prime squared-- the Square of that factor-- times the variance of k, which is m kT over h bar squared. So what do I get here?

I will get 2. And then I would get essentially m kT. And then, h bar squared here.

Now, this combination m kT divided by h, same thing as here, m kT and h. Clearly, what that is, is giving me some kind of a lambda. So let me write the answer.

Eventually, what I find is that evaluating the one-particle density matrix between two points, x and x prime, will give me 1 over V. And then this exponential factor, which is x minus x prime squared. The coefficient here has to be proportional to lambda squared. And if you do the algebra right, you will find that the coefficient is pi over 2. Or is it pi?

Let me double check. It's a numerical factor that is not that important. It is pi, pi over lambda squared. OK, so what does this mean?

So I have a particle in the box. It's a quantum mechanical particle at some temperature T. And I'm trying to ask, where is it? So if I want to ask where is it, what's the probability it is at some point x?

I refer to this density matrix that I have calculated in coordinate space. Put x prime and x to be the same. If x prime and x are the same, I get 1 over V.

Essentially, it says that the probability that I find a particle is the same anywhere in the box. So that makes sense. It is the analog of throwing the coin. We don't know where the particle is. It's equally likely to be anywhere.

Actually, remember that I used periodic boundary conditions. If I had used sine and cosines, there would have been some dependence on approaching the boundaries. But the matrix tells me more than that.

Basically, the matrix says it is true that the diagonal elements are 1 over V. But if I go off-diagonal, there is this [INAUDIBLE] scale lambda that is telling me something. What is it telling you?

Essentially, it is telling you that through the procedures that we have outlined here, if you are at the temperature T, then this factor tells you about the classical probability of finding a momentum p. There is a range of momenta because p squared over 2m, the excitation is controlled by kT. There is a range of momenta that is possible. And through a procedure such as this, you are making a superposition of plane waves that have this range of momentum. The superposition of plane waves, what does it give you?

It essentially gives you a Guassian blob, which we can position anywhere in the space. But the width of that Gaussian blob is given by this factor of lambda, which is related to the uncertainty that is now quantum in character. That is, quantum mechanically if you know the momentum, then you don't know the position. There is the Heisenberg uncertainty principle.

Now we have a classical uncertainty about the momentum because of these Boltzmann weights. That classical uncertainty that we have about the momentum translates to some kind of a wave packet size for these-- to how much you can localize a particle.

So really, this particle you should think of quantum mechanically as being some kind of a wave packet that has some characteristic size lambda. Yes?

AUDIENCE: So could you interpret the difference between x and x prime appearing in the [INAUDIBLE] as sort of being the probability of finding a particle at x having turned to x prime within the box?

PROFESSOR: At this stage, you are sort of introducing something that is not there. But if I had two particles-- so very soon we are going to do multiple particles-- and this kind of structure will be maintained if I have multiple particles, then your statement is correct. That this factor will tell me something about the probability of these things being exchanged, one tunneling to the other place and one tunneling back. And indeed, if you went to the colloquium yesterday, Professor [? Block ?] showed pictures of these kinds of exchanges taking place. OK?

But essentially, really at this stage for the single particle with no potential, the statement is that the particle is, in reality, some kind of a wave packet. OK?

Fine. So any other questions?

The next thing is, of course, what I said. Let's put two particles in the potential. So let's go to two particles. Actually, not necessarily even in the box. But let's say what kind of a Hamiltonian would I write for two particles?

I would write something that has a kinetic energy for the first particle, kinetic energy for the second particle, and then some kind of a potential of interaction. Let's say, as a function of x1 minus x2.

Now, I could have done many things for two particles. I could, for example, have had particles that are different masses, but I clearly wanted to write things that were similar. So this is Hamiltonian that describes particles labeled by 1 and 2. And if they have the same mass and the potential is a function of separation, it's certainly h that is for 2 and 1 exchange. So this Hamiltonian has a certain symmetry under the exchange of the labels.

Now typically, what you should remember is that any type of symmetry that your Hamiltonian has or any matrix has will be reflected ultimately in the form of the eigenvectors that you would construct for that. So indeed, I already know that for this, I should be able to construct wave functions that are either symmetrical or anti-symmetric under exchange.

But that statement aside, let's think about the meaning of the wave function that we have for the two particles. So presumably, there is a psi as a function of x1 and x2 in the coordinate space. And for a particular quantum state, you would say that the square of this quantity is related to quantum probabilities of finding particles at x1 and x2.

So this could, for example, be the case of two particles. Let's say oxygen and nitrogen. They have almost the same mass. This statement would be true. But the statement becomes more interesting if we say that the particles are identical. So for identical particles, the statement of identity in quantum mechanics is damped at this stage.

You would say that I can't tell apart the probability that one particle, number 1, is at x1, particle 2 is at x2 or vice versa. The labels 1 and 2 are just thinks that I assign for convenience, but they really don't have any meaning. There are two particles. To all intents and purposes, they are identical. And there is some probability for seeing the events where one particle-- I don't know what its label is-- at this location. The other particle-- I don't know what its label is-- is at that location. Labels don't have any meaning.

So this is different. This does not have a classic analog. Classically, if I put something on the computer, I would say that particle 1 is here at x1, particle 2 is here at x2. I have to write some index on the computer.

But if I want to construct a wave function, I wouldn't know what to do. I would just essentially have a function of x1 and x2 that is peaked here and there.

And we also know that somehow quantum mechanics tells us that there is actually a stronger version of this statement, which is that psi of x1, x2 x1 is either plus psi of x1, x2 or minus psi of x1, x2 for identical particles depending on whether they are boson or fermions. So let's kind of build upon that a little bit more and go to many particles.

So forth N particles that are identical, I would have some kind of a state psi that depends on the coordinates. For example, in the coordinate representation, but I could choose any other representation. Coordinate kind of makes more sense. We can visualize it.

And if I can't tell them apart, previously I was just exchanging two of the labels but now I can permute them in any possible way. So I can add the permutation P. And of course, there are N factorial in number. And the generalization of the statement that I had before was that for the case of bosons, I will always get the same thing back. And for the case of fermions, I will get a number back that is either minus or plus depending on the type of permutation that I apply of what I started with. So I have introduced here something that I call minus 1 to the power of P, which is called a parity of the permutation.

A permutation is either even-- in which case, this minus 1 to the power of P is plus-- or is odd-- in which case, this minus 1 to P is minus 1. How do I determine the parity of permutation?

Parity of the permutation is the number of exchanges that lead to this permutation. Basically, take any permutation. Let's say we stick with four particles. And I go from 1, 2, 3, 4, which was, let's say, some regular ordering, to some other ordering. Let's say 4, 2, 1, 2. And my claim is that I can achieve this transformation through a series of exchanges. So I can get here as follows.

I want 4 to come all the way back to here, so I do an exchange of 1 and 4. I call the exchange in this fashion. I do exchange of 1 and 4. The exchange of 1 and 4 will make for me 4, 2, 3, 1. OK.

I compare this with this. I see that all I need to do is to switch 3 and 1. So I do an exchange of 1 and 3. And what I will get here is 4, 2, 1, 3. So I could get to that permutation through two exchanges. Therefore, this is even.

Now, this is not the only way that I can get from one to the other. I can, for example, sit and do multiple exchanges of this 4, 2 a hundred times, but not ninety nine times. As long as I do an even number, I will get back to the same thing. The parity will be conserved. And there's another way of calculating parity. You just start with this original configuration and you want to get to that final configuration. You just draw lines.

So 1 goes to 1, 2 goes to 2, 3 goes to 3, 4 goes to 4. And you count how many crossings you have, and the parity of those crossings will give you the parity of the permutation.

So somehow within quantum mechanics, the idea of what is identical particle is stamped into the nature of the wave vectors, in the structure of the Hilbert space that you can construct. So let's see how that leads to this simple example of particle in the box, if we continue to add particles into the box.

So we want to now put N particles in box. Otherwise, no interaction completely free. So the N particle Hamiltonian is some alpha running from 1 to N p alpha squared over 2m kinetic energies of all of these things. fine.

Now, note that this Hamiltonian, since it is in some sense built up of lots of noninteracting pieces. And we saw already classically, that things are not interacting-- calculating probabilities, partition functions, et cetera, is very easy. This has that same structure. It's the ideal gas. Now, quantum mechanically. So it should be sufficiently easy.

And indeed, we can immediately construct the eigenstates for this. So we can construct the basis, and then do the calculations in that basis. So let's look at something that I will call product state. And let's say that I had this structure that I described on the board above where I have a plane wave that is characterized by some wave number k for one particle. For N particles, I pick the first one to be k1, the second one to be k2, the last one to be kN.

And I will call this a product state in the sense that if I look at the positional representation of this product state, it is simply the product of one-particle states. So this is a product over alpha of x alpha k alpha, which is this e to the i kx k dot x over square root of it. So this is perfectly well normalized if I were to integrate over x1, x2, x3. Each one of the integrals is individually normalized. The overall thing is normalized. It's an eigenstate because if I act Hn on this product state, what I will get is a sum over alpha h bar squared k alpha squared over 2m, and then the product state back. So it's an eigenstate. And in fact, it's a perfectly good eigenstate as long as I want to describe a set of particles that are not identical.

I can't use this state to describe particles that are identical because it does not satisfy the symmetries that I set quantum mechanics stamps into identical particles. And it's clearly the case that, for example-- so this is not symmetrized since clearly, if I look at k1, k2, k1 goes with x1, k2 goes with x2. And it is not the same thing as k2, k1, where k2 goes with x1 and k1 goes with x2. Essentially, the two particles can be distinguished. One of them has momentum h bar k1. The other has momentum h bar k 2. I can tell them apart because of this unsymmetrized nature of this wave function.

But you know how to make things that are symmetric. You basically add k1 k2 plus k2 k1 or k1 k2 minus k2 k1 to make it anti-symmetrized. Divide by square root of 2 and you are done. So now, let's generalize that to the case of the N-particle system.

So I will call a-- let's start with the case of the fermionic state. In fermionic state, I will indicate by k1, k2, kN with a minus index because of the asymmetry or the minus signs that we have for fermions.

The way I do that is I sum over all N factorial permutations that I have. I let p act on the product state. And again, for two particles, you have the k1 k2, then you do minus k2 k1.

For general particles, I do this minus 1 to the power of p. So all the even permutations are added with plus. All the odd permutations are added with minus. Except that this is a whole bunch of different terms that are being added. Again, for two particles, you know that you have to divide by a square root of 2 because you add 2 vectors. And the length of the overall vector is increased by a square root of 2.

Here, you have to divide in general by the number of terms that you have, square root of N factorial. The only thing that you have to be careful with is that you cannot have any two of these k's to be the same. Because let's say these two are the same, then along the list here I have the exchange of these two. And when I exchange them, I go from even to odd. I put a minus sign and I have a subtraction. So here, I have to make sure that all k alpha must be distinct.

Now, say the bosonic one is simpler. I basically construct it, k1, k2, kN with a plus. By simply adding things together, so I will have a sum over p. No sign here. Permutation k1 through kN. And then I have to divide by some normalization.

Now, the only tricky thing about this is that the normalization is not N factorial. So to give you an example, let's imagine that I choose to start with a product state where two of the k's are alpha and one of them is beta. So let's sort of put here 1, 1, 1, 2, 3 for my k1, k2, k3. I have chosen that k1 and k2 are the same. And what I have to do is to sum over all possible permutations.

So there is a permutation here that is 1, 3, 2. So I get here alpha, beta, alpha. Then

I will have 2, 1, 3, 2, 3, 1, 3, 1, 2, 3, 2, 1, which basically would be alpha, alpha, beta, alpha, beta, alpha, beta, alpha, alpha, beta, alpha, alpha. So there are for three things, 3 factorial or 6 permutations. But this entity is clearly twice alpha, alpha, beta, plus alpha, beta, alpha, plus beta, alpha, alpha. And the norm of this is going to be 4 times 1 plus 1 plus 1. 4 times 3, which is 12.

So in order to normalize this, I have to divide not by 1/6, but 1/12. So the appropriate normalization here then becomes 1 over root 3. Now, in general what would be this N plus?

To calculate N plus, I have to make sure that the norm of this entity is 1. Or, N plus is the square of this quantity. And if I were to square this quantity, I will get two sets of permutations. I will call them p and p prime. And on one side, I would have the permutation p of k1 through kN. On the other side, I would have a permutation k prime of k1 through kN.

Now, this is clearly N factorial square terms. But this is not N factorial squared distinct terms. Because essentially, over here I could get 1 of N factorial possibilities. And then here, I would permute over all the other N factorial possibilities. Then I would try the next one and the next one. And essentially, each one of those would give me the same one. So that is, once I have fixed what this one is, permuting over them will give me one member of N factorial identical terms. So I can write this as N factorial sum over Q. Let's say I start with the original k1 through kN, and then I go and do a permutation of k1 through kN. And the question is, how many times do I get something that is non-zero?

If these two lists are completely distinct, except for the identity any transformation that I will make here will make me a vector that is orthogonal to this and I will get 0. But if I have two of them that are identical, then I do a permutation like this and I'll get the same thing back. And then I have two things that are giving the same result.

And in general, this answer is a product over all multiple occurrences factorial. So let's say here there was something that was repeated twice. If it had been repeated three times, then all six possibilities would've been the same. So I would have had 3 factorial.

So if I had 3 and then 2 other ones that were the same, then the answer would have been multiplied by 3 factorial 2 factorial. So the statement here is that this N plus that I have to put here is N factorial product over k nk factorial, which is essentially the multiplicity that is associated with repeats. Oh

So we've figured out what the appropriate bosonic and fermionic eigenstates are for this situation of particles in the box where I put multiple particles. So now I have N particles in the box. I know what the appropriate basis functions are. And what I can now try to do is to, for example, calculate the analog of the partition function, the analog of what I did here for N particles, or the analog of the density matrix. So let's calculate the analog of the density matrix and see what happens.

So I want to calculate an N particle density matrix, completely free particles in a box, no interactions. For one particle, I went from x to x prime. So here, I go from x1, x2, xN. To here, x1 prime, x2 prime, xN prime. Our answer ultimately will depend on whether I am dealing with fermions or bosons. So I introduce an index here eta. Let me put it here, eta is plus 1 for bosons and minus 1 for fermions. So kind of goes with this thing over here whether or not I'm using bosonic or fermionic symmetrization in constructing the wave functions here.

You say, well, what is rho? Rho is e to the minus this beta hn divided by some partition function ZN that I don't know. But what I certainly know about it is because I constructed, in fact, this basis so that I have the right energy, which is sum over alpha h bar squared alpha squared k alpha squared over 2m. That is, this rho is diagonal in the basis of the k's. So maybe what I should do is I should switch from this representation to the representation of k's.

So the same way that for one particle I sandwiched this factor of k, I am going to do the same thing over here. Except that I will have a whole bunch of k's. Because I'm doing a multiple system. What do I have?

I have x1 prime to xN prime k1 through kN, with whatever appropriate symmetry it is. I have e to the minus sum over alpha beta h bar squared k alpha squared over 2m. Essentially, generalizing this factor that I have here, except that I have to normalize it with some ZN that I have yet to calculate.

And then I go back from k1, kN to x1, xN. Again, respecting whatever symmetry I want to have at the end of the day.

Now, let's think a little bit about this summation that I have to do over here. And I'll put a prime up there to indicate that it is a restricted sum. I didn't have any restriction here. I said I integrate over-- or sum over all possible k's that were consistent with the choice of 2 pi over l, or whatever.

Now, here I have to make sure that I don't do over-counting. What do I mean?

I mean that once I do this symmetrization, let's say I start with two particles and I have k1, k2. Then, I do the symmetrization. I get something that is a symmetric version of k1, k2. I would have gotten exactly the same thing if I started with k2, k1. That is, a particular one of these states corresponds to one list that I have here.

And so over here, I should not sum over k1, k2, k3, et cetera, independently. Because then I will be over-counting by N factorial. So I say, OK, let me sum over things independently and then divide by the over-counting. Because presumably, these will give me N factorial similar states. So let's just sum over all of them, forget about this. You say, well, almost. Not quite because you have to worry about these factors.

Because now when you did the N factorial, you sort of did not take into account the kinds of exchanges that I mentioned that you should not include because essentially if you have k1, k1, and then k3, then you don't have 6 different permutations. You only have 3 different permutations. So actually, the correction that I have to make is to multiply here by this quantity. So now that sum be the restriction has gone into this. And this ensures that as I sum over all k's independently, the over-counting of states that I have made is taken into account by the appropriate factor.

Now, this is actually very nice because when I look at these states, I have the normalization factors. The normalizations depend on the list of k's that I have over here. So the normalization of these objects will give me a 1 over N factorial product over k nk factorial. And these nk factorials will cancel.

Now you say, hang on. Going too fast. This eta says you can do this both for fermions and bosons. But this calculation that you did over here applies only for the case of bosons. You say, never mind because for fermions, the allowed values of nk are either 0 or 1. I cannot put two of them. So these nk factorials are either-- again, 1 or 1. So even for fermions, this will work out fine. There will be appropriate cancellation of wave functions that should not be included when I do the summation over permutations with the corresponding factors of plus and minus.

So again, the minus signs-- everything will work. And this kind of ugly factor will disappear at the end of the day. OK? Fine.

Now, each one of these is a sum over permutations. Actually, before I write those, let me just write the factor of e to the minus sum over alpha beta h bar squared k alpha squared over 2m. And then I have ZN. I don't know what ZN is yet.

Each one of these states is a sum over permutations. I have taken care of the overall normalization. I have to do the sum over the permutations. And let's say I'm looking at this one. Essentially, I have to sandwich that with some combination of x's. What do I get?

I will get a factor of eta to the p. So actually, maybe I should have really within the following statement at some point.

So we introduced these different states. I can combine them together and write k1, k2, kN with a symmetry that is representative of bosons or fermions. And that's where I introduced a factor of plus or minus, which is obtained by summing over all permutations. This factor eta, which is plus or minus, raised to the p. So for bosons, I would be adding everything in phase. For fermions, I would have minus 1 to the p. I have the permutation of the set of indices k1 through kN in the product state.

Whoops. This should be product. This should be a product state. And then I have to divide by square root of N factorial product over k nk factorial. And clearly, and we will see a useful way of looking at this. I can also look at the set of states that are occupied among the original list of k's. Some k's do not occur in this list. Some k's occurs once. Some k's occur more than once. And that's how I construct these factors.

And then the condition that I have is that, of course, sum over k nK should be the total number N. That nk is 0 or 1 for fermions. That nk is any value up to that constraint for bosons.

Now, I take these functions and I substitute it here and here. What do I get?

I get the product state. So I have eta to the p. The product state is e to the i sum over, let's say, beta k of p beta x beta prime. And then from the next one, I will get e to the minus i sum over-- again, some other beta. Or the same beta, it doesn't matter. p prime beta x beta.

I think I included everything. Yes?

AUDIENCE: Question. Why do you only have one term of beta to the p?

PROFESSOR: OK. So here I have a list. It could be k1, k2, k2-- well, let's say k3, k4. Let's say I have four things. And the product means I essentially have a product of these things.

When I multiply them with wave functions, I will have e to the i, k1, x1 et cetera. Now, I do some permutation. Let's say I go from 1, 2, 3 to 3, 2, 1. So there is a permutation that I do like that. I leave the 4 to be the same. This permutation has some parity. So I have to put a plus or minus that depends on the parity.

AUDIENCE: Yeah. But my question is why in that equation over there--

PROFESSOR: Oh, OK.

AUDIENCE: --we have a term A to the p for the k.

PROFESSOR: Yes.

AUDIENCE: k factor and not for the bra.

PROFESSOR: Good. You are telling me that I forgot to put p prime and eta of p prime.

AUDIENCE: OK. I wasn't sure if there was a reason.

PROFESSOR: No. I was going step by step. I had written the bra part and I was about to get to the k part. I put that part of the ket, and I was about to do the rest of it.

AUDIENCE: Oh, sorry.

PROFESSOR: It's OK. OK. So you can see that what happens at the end of the day here is that a lot of these nk factorials disappear. So those were not things that we would have liked.

There is a factor of N factorial from here and a factor of N factorial from there that remains. So the answer will be 1 over ZN N factorial squared. I have a sum over two permutations, p and p prime, of something. I will do this more closing next time around, but I wanted to give you the flavor.

This double sum at the end of the day, just like what we did before, becomes one sum up to repetition of N factorial. So the N factorial will disappear. But what we will find is that the answer here is going to be something that is a bunch of Gaussians that are very similar to the integration that I did for one particle. I have e to the minus k squared over 2m. I have e to the i k x minus x prime. Except that all of these k's and x's and x primes have been permuted in all kinds of strange ways.

Once we untangle that, we find that the answer is going to end up to be eta of Q x alpha minus x prime Q alpha squared divided by pi 2 lambda squared. And we have a sum over alpha, which is really kind of like a sum of things that we had for the case of one particle.

So for the case of one particle, we found that the off-diagonal density matrix had elements that's reflected this wave packet nature of this. If we have multiple particles that are identical, then the thing is if I have 1, 2 and then I go 1, 1 prime, 2, 2 prime for the two different locations that I have x1 prime, x2 prime, et cetera. And 1 and 2 are identical, then I could have really also put here 2 prime, 1 prime. And I wouldn't have known the difference.

And this kind of sum will take care of that, includes those kinds of exchanges that I mentioned earlier and is something that we need to derive more carefully and explain in more detail next time.