

ANNOUNCER: The following content is provided under a creative commons license. Your support will help MIT Open Courseware 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: Last few lectures, and we started to see some of the consequences and looked at simple changes, transformations of expansion and mixtures of liquids and gases. And saw how in the framework of statistical mechanics, we could derive the thermodynamic results that you saw before, based on an empirical framework. On the thermodynamic framework we've been working with all term. But you saw them derived from a microscopic perspective. Now what I want to do is move to examples that are more commonly encountered in chemistry. One that we did last time actually was very common, or is at least a prototype for something common. That is we looked at what would be a simple model for a polymer with different configurations available to it. And just tried to understand what the basic thermodynamics is that it would exhibit. What the heat capacities would be in various limiting cases. High and low temperature. I want to continue that today. But for an even more common case.

So the model that I want to lay out is indicated on your notes for today. And essentially, you could look at it as double stranded DNA. Any kind of polymer with a whole variety of configurations. Ultimately what's going to matter is what is the set of energy levels available to the system. And so here's what it looks like. The way I've tried to depict it. I'm imagining something like, covalently bonded chains that have links between them. So of course, it's reminiscent of DNA with hydrogen bonding between the strands. And the idea in this simple model is that you can have various numbers of these hydrogen bonds. So this is the case that would be the most energetically stable. That is, all the available hydrogen bonds are formed between neighboring pairs. So this would be our lowest energy. Now let's starts breaking some hydrogen bonds. And what I'm imagining is they break starting at one end and working their way down. So the next highest energy state would be like this. So now these are still hydrogen bonded. This is not. And there's a cost in energy for breaking that hydrogen bond. We'll call it ϵ_0 . It's a positive number. So that this is higher energy than this. And now we'll go to the next one, which could have two broken bonds. Broken hydrogen bonds. And so on. And if we imagine that this is very long, there may be many many, members of the chain. Then of course this sequence of structures with corresponding energies might go on for a very long way.

So this is a simple construction that gives us a set of states with distinct energies. And so we have a series of configurational energy levels. They're evenly spaced in this model. So there's zero, E_0 , two E_0 . And so on. So those are available energy states.

And now, just based on that simple model, we should be able to figure out all the thermodynamics. We should

should be able to figure out what the equilibrium energy is at a particular temperature. What the average energy is. And so forth. And we can do it in a straightforward way. So we'll start with our molecular partition function, as always. So it's q configurational. And it's just a sum over our available states with their associated Boltzmann factors. So sum over n of $e^{-\epsilon_n / kT}$. Now the way I've modeled this, there are no degeneracies. Each energy level has just one microscopic state of the molecule corresponding to it. So there are no molecular degeneracies. Turns out we'll be able to put this into a simple form.

So the first thing is, let's approximate that we could take the sum -- not to some finite level, which would be the case if there's a finite number of elements in the chain. But all the way through infinity. And the idea here is the following. We'll assume that the chain at least has a great many members. Even if it's finite. And then at some point, this becomes a really small number. Even if ϵ is naught, the energy of one of the bonds, is less than kT . Once we multiply it by a very large number n , it becomes very much greater than kT . So the underlying assumption here is, well there will not be a significant number of molecules in the highest energy state anyway. In other words, the terms higher than a finite but large value of n between there and infinity are going to be so small that those terms will contribute negligibly to the sum anyway. And on that basis, we could make this approximation. And the incentive to make this approximation and carry the sum to infinity is that then we can put this in what turns out to be a very simple closed form.

So now if we just write out a few of the individual terms in the sum, it's one plus $e^{-\epsilon_0 / kT}$. Plus $e^{-2\epsilon_0 / kT}$. I'll write this as x . $e^{-\epsilon_0 / kT}$. I'll write this as x . I'll write this as x^2 . $e^{-2\epsilon_0 / kT}$. I'll write this as x^2 . And so on. So this has the form of a sum that looks like one plus x , plus x^2 , and so on. And because the sum goes to infinity, we can simply write it as one over one minus x . Sort of a very simple result. And now putting this back into x , it's one over one minus $e^{-\epsilon_0 / kT}$. That's our q_0 . Our molecular partition function.

So this model yields a particularly simple result for q_0 . And then everything else follows from there. So then we can write the canonical partition function Q_N . It's just q_0^N . It's just q_0 configurational to the N th power, where N is the total number of these molecules. So it's just one over one minus $e^{-\epsilon_0 / kT}$ to the N th power. And then we can start writing out the results for the various thermodynamic properties. So A configurational is $-NkT \ln q_0$. That is, remember, A is $-kT \ln Q$. And this is just going to come straight out. So it's $-NkT \ln q_0$. Or $-NkT \ln (1 / (1 - e^{-\epsilon_0 / kT}))$. Or positive $NkT \ln (1 - e^{-\epsilon_0 / kT})$. So a pretty simple form for the Helmholtz free energy. And I'm not going to write out all of the individual thermodynamic terms, but I'll write a few of them. So μ , the chemical potential for these configurations, is just dA/dN . With T and V constant. And so the only N dependence, the only dependence on the number of molecules is multiplicative. Right here. So this

just gives us $kT \log$ of one minus $e^{-\epsilon_0/kT}$. And the important point to realize here that we mentioned last time too, in the examples we considered then, especially the last one, is that because the only place that N figures in is as a multiplicative factor. What this is telling us is that we just have a chemical potential, of Helmholtz free energy per molecule. The molecules are all independent. The total energy doesn't depend on any interaction between this molecule and its neighbor somewhere else. So terms like this are simply additive. So if we work out the chemical potential, it's just one over N times A .

And I'll just write the result for u configurational. Because I do want to look at the heat capacity. So it's $NkT^2 \frac{d \log q_{\text{config}}}{dT}$. With N and V held constant. And it turns out to just be $N \frac{\epsilon_0}{kT} e^{-\epsilon_0/kT}$. One over $e^{\epsilon_0/kT} - 1$.

So now we can look at the heat capacity. C_v configurational. So it's $\frac{du_{\text{config}}}{dT}$. With constant N and V . And taking this derivative with respect to temperature. Gives us an expression of the following form. It's $Nk \frac{\epsilon_0}{kT} \frac{e^{-\epsilon_0/kT}}{e^{-\epsilon_0/kT} - 1}$. Just taking the derivative in the usual way. Now, this looks kind of complicated for the heat capacity. But let's take a look, just like we did last time for the simpler case that we treated then, let's take a look at the high and low temperature limits of what happens to the heat capacity.

So at high temperature, well we can start by just looking at the energy. That has a form that's fairly simple. So when temperature is large, this is small, and we can Taylor expand it. So then we're just going to have one plus $\frac{\epsilon_0}{kT} - 1$. Which means the ones will cancel. And we end up with a fairly simple result. So u_{config} , in that case, is just NkT . The ϵ_0 zeroes are going to cancel also. And what that says is that the heat capacity -- and of course we could take the limit of this, but we can just take the derivative of this with respect to temperature more easily. So C_v configurational. It's just N times k . In other words, the heat capacity in the high temperature limit is a constant. Last time we treated a simpler case, where there were only four configurations altogether available, to this sort of simple polymer model that we drew. Unlike this present model where we're saying there are essentially infinite number of configurations and different energies available. So many that the highest ones we're never even going to access, because they'll be much higher than kT . Here it's different, so now we have -- You know before, when we had a limited number of total states, remember what happened in the heat capacity at high temperature. What was the limiting case? Finite number of states. What's the heat capacity at high temperature? This is going to be on the exam. What's the heat capacity at high temperature, if there's a finite number of states available?

STUDENT: Zero.

PROFESSOR: It's zero. What was the low temperature limit of the heat capacity?

STUDENT: Zero.

PROFESSOR: It was also zero. Right. And the idea was for a system with a finite number of states. The idea was - - so let's say, the way we had it before, there was one state with energy zero. And we had three states with some energy ϵ . And that was it. Those were all the molecular states available. So we looked at the two cases. In one case, when the high temperature limit where kT is much bigger than any of this stuff. In that limit, the molecules are just equally likely to be in any of these states. Because there's much, much more thermal energy than this energy difference. And the molecules are constantly getting kicked around by the available thermal energy. So if you raise the temperature a little bit more, it doesn't make any difference. The molecules already are evenly distributed among the states. There's no additional configurational energy. So du/dT is zero. It can't increase any more. So in that case, the high temperature limiting heat capacity is zero. And in the other case -- In the low temperature limit, now let's say kT is really low. It's much less than ϵ . So now we'll redraw this as zero. And put the ϵ states up here. And kT is here. Well when it's like this, the temperature is so low that there's not nearly enough thermal energy to populate any of these states. And if you change the temperature by a little bit, it's still not enough thermal energy to populate any of these states. So once again, du/dT is zero. You change the temperature and the configurational energy doesn't change. So the heat capacity is zero again.

This is why it's so informative to measure heat capacities. Because you can learn an awful lot about the intrinsic structure of the material. What are the energy levels available to it? What do they do? You can learn a tremendous amount about that by making measurements of the heat capacity.

Well, now we're in a different case. We have a whole set of evenly spaced energy levels. Never ends. So now let's look at the high temperature limit. But we're never in a limit that's higher than the highest level, because we're assuming it goes on forever. So it's up here somewhere. There's kT . So what happens? Well if you raise the temperature, there still are higher lying levels that can be populated, and that will get populated. So du/dT isn't going to be zero in the high temperature limit, in this case. But it stops changing at some point. Because nothing's very different about this than having kT be, let's say, up here or up here. So in the high temperature limit, yes, the energy does change with temperature. But it changes in the same way at any temperature. In other words, the energy is just linearly varying with temperature. And the heat capacity is a constant. du/dT doesn't change anymore, once you're in the high temperature limit.

Now without me writing any expression on the board, what's the low temperature limit of the heat capacity going to be in this case? Going to be on the exam.

STUDENT: Zero?

PROFESSOR: Zero, that's still going to be the same, right? Put kT way down here. That's just like this, right? Not near enough energy to populate even though lowest, you know, the first level above the ground state. Change the temperature a little bit, still not enough thermal energy to get up there. So the energy doesn't change with temperature in that low temperature limit. So you can immediately see what's going to happen at low temperature. Any questions? Yeah?

STUDENT: [UNINTELLIGIBLE] will still be $n k t$, and then it's just -- [UNINTELLIGIBLE]

PROFESSOR: So, of course, that's a limiting case here, right? That happened because in the limit of high temperature, then this exponent is really small, right? So then you can Taylor expand it. In the limit of low temperature, this is really big. This is nearly zero. So $e^{\epsilon_0/kT}$ is much bigger than kT in that case. This is big. This is negligible, right? Oh. Wait a minute. Something's making me real unhappy. I can't have this right. Ah. Yes. It needs to be zero, is what it needs to be. What am I thinking? Of course. It's one over a huge number. It's zero. OK. So of course you can't expand it. It's just, this is much bigger than this. This is an enormous number at that point. So in other words, what it's saying is the configurational energy is zero, because everything is stuck in the ground state. And it stays zero if you vary the temperature. So you get the zero limiting value for the heat capacity, and the energy itself is also zero. Other questions? OK.

Let me just make a few comments about the entropy. So I didn't write it out before. And I'm tempted not to do it again, but I suppose I'll do it. So it turns out to be Nk , minus \log of one minus $e^{-\epsilon_0/kT}$. Plus ϵ_0/kT . Over $e^{-\epsilon_0/kT}$, minus one. And this comes from combining the terms for A and u . It comes from $-A/T$. Plus u/T . And what I want to do is look at its limiting cases also. In particular, what happens here in the limit of high temperature. And what happens turns out to be $k \log kT$ over ϵ_0 to the N th power. Or put the N over here. OK.

And what this is telling us about is the number of available states. Roughly, how many states are there that are accessible to the system at some temperature. So in other words, think of it as -- let's put the N back there. $k \log kT$ over ϵ_0 to the N th power. And think of it as $k \log \Omega$. Where that would be the degeneracy. Now all the states are in equal energy, but remember for the whole system, remember we discussed this before. How you have a very, very narrow distribution of system energy states at equilibrium. So you can think of this as the degeneracy of the system states that are actually going to exist at a particular temperature.

So if we look at the limiting value of the partition function, it's just kT over Ω . Or the same thing for capital Q . kT over ϵ_0 , sorry. To the N th power. So you have a very simple expression. And so again, what this is doing, is it's giving us a measure of about how many states are available. And it's particularly informative to look at that for the molecular partition function. What it's telling us is, if I look at kT over ϵ_0 . And these levels,

remember, are evenly spaced. So here's ϵ_0 , two ϵ_0 , three ϵ_0 , and so on. It says, you know, if kT is about ten times bigger than ϵ_0 . So this is ten. It's telling us, roughly how many states does the system have thermal access to. The molecular state. It has about ten states. So going over to our picture of the set of structures, you could have anywhere up to about ten bonds broken. Now, the individual molecules are going to be in a whole range of states. Some of them will have fewer than that, and some of them will have more than that number of bonds broken. But on average, it's going to be about that number.

And then, if you look at the whole system, the number of states available, of course it's astronomical. Because you have to take each molecule, and say, well it could be in any one of something on the order of ten states. And then whole set of N other molecules can be in the states they might be in. Change the first one, and do it again. So you have this astronomical number of system states. But remember, like we discussed once before, it'll turn out that although the individual molecule states vary considerably with energy, the system states, which are averaging over some astronomical number of molecules, where capital N is something like ten to the 24th. Once you average over that many individual molecules, you find that there's very, very little fluctuation in the energy. In the system energy. The individual molecule energies vary considerably. Realistically the variation of the molecular energies -- that variation is going to be comparable to the energy itself. To the average energy. So if the average energy is roughly you know, the ten ϵ_0 , you say okay, how much might it vary? Well, there are going to be some molecules that have only a couple of bonds broken, and some that might have 20 bonds broken. The variation will be on the same order of magnitude as the average itself. But then you average over capital N of them. And then you immediately discover that there's very, very, very little variation.

And in particular, what happens then, is, you know, for molecular average energy, ϵ_0 , and variance, or standard deviation about the same magnitude. System average energy is u , and it's going to be capital N times. Well let's say that -- average -- Oh sorry, let me not put the zero There. It's just average energy. And the system average energy is just N times the molecular energy. But the system variance is going to be on the order of the square root of N times ϵ_0 . If you've done statistics, then you've seen that sort of result. You do a bunch of samplings a capital N number of samplings, and the variation looks like the square root of that number. So what ends up happening then, if you look at the relative variation, you might look and then say, well it's pretty big. This is ten to the 12, right? It's still a huge variance.

But let's compare it to the average. So this is ten to the 24th. This is ten to the 12th. It's on the order of ten to the minus 12. An incredibly tiny fractional variation in the system energy. You could never -- There would be no practical way to measure it. And of course that is consistent with what you would expect. If you say let's measure the configurational energy of a bunch of molecules in a liquid solution, or molecules in a gas floating around. And it's a mole of them, that total average energy is not going to fluctuate significantly, even though individual

molecules that you pick out of that whole system might have quite widely varying energies. And that's the point. So it's an incredibly small variation.

Now you can derive this. I didn't derive this, of course. I asserted that this is the case. And it's probably familiar to some of you if you've seen some statistics. But the way you would derive it is, you know in addition to calculating the average energy, the average of E , you can also calculate the average of the energy squared. And you can calculate the standard deviation that way. You calculate the average of E squared. And then you minus the average of E , the quantity squared. Take the square root. That's your root mean square. And that's what leads to this result. So it's pretty straight forward to calculate it also.

Alright. So any questions about just the extent of variation of the individual molecule energies or the system energies?

Okay. Then, now what I'd like to do is look at another kind of energy that's going to turn out to have exactly the same set of levels that we just derived from this simple model. And some of you may have seen this before. Many of you may not have. And for right now, I'll just assert it. But it will illustrate why this is so useful. It turns out that if I've got molecular vibrations -- you know, there's nitrogen and oxygen in the air. If I look at those nitrogen vibrational energy levels. Or oxygen energy levels. They look like this. Evenly spaced, non degenerate energy levels. So the model that we've constructed here, based on this simple two-chain polymer, actually gives us a set of energy levels that maps directly onto the vibrational energy levels of a molecule. So all the results that we've just seen apply, not just for conformations of a polymer, but for vibrations of a molecule. So everywhere where it says configurational, you can just write in vibrational. And you'll still be right. because of course, what matters is what are the states that are available to the system, and what are their energies? After that, the formalism of statistical mechanics takes over, and calculates partition functions and thermodynamic functions. The input into that is states. And their energies. Well, we have the exact same set of states and energies.

So we immediately arrive at a super important case and its results for molecular vibrations. And not only molecular vibrations, but vibrations of a crystal lattice. Acoustic vibrations of a glass. Or even a liquid. So there's an enormously wide ranging set of results that we've derived by starting at this simple picture. And when you take quantum mechanics, you'll see that indeed you do get this set of evenly spaced energy levels. You may well have seen the results before, and you'll see it derived there.

OK. Given that, then I want to talk a little bit further about the heat capacity. So we've seen the limiting cases for the heat capacity. Namely, the low temperature limit is zero. That's the case whenever you have quantized levels. You can always get kT lower by far than the lowest excited state. So that everything is stuck in the ground state. How low that is depends on how far apart the states are spaced. But there's got to be some temperature

somewhere that's down there. So we've seen the heat capacity limiting cases. Let me just sort of draw them again. So here's kT much less than ϵ_0 . And up here, kT is much bigger than ϵ_0 . Now I'm just going to sketch the full temperature dependence. In other words, I'm going to connect those two limits that we've seen. So here's what that looks like. And I'll write it as vibrational. So we know it's got to be zero at the low temperature limit. And we know that it's this constant value in the high temperature limit. It's Nk . And in some way, it smoothly connects. And if you look at the actual scale, here -- so here's kT over ϵ_0 . Here's the limit of low temperature. Actually, you don't have to be very high above ϵ_0 to be in in this high temperature limit. So if you look at the plot that's on your notes, this is already -- it's not quite leveling off precisely. But it's already pretty close. When kT is just two times ϵ_0 . It's already almost at the limiting case.

Now this has tremendous significance. So a long time before quantum mechanics was developed, people made measurements of the heat capacities of materials. And they were familiar with the fact that when you got to ordinary temperature -- room temperature, the heat capacity was temperature independent. And that was understandable for reasons we'll discuss shortly. Basically, in that case, the heat capacity -- it's not just Nk for one vibrational mode. Of course, that's what I describe for a single vibrational mode of, say, a molecule. If you have a crystal lattice with N atoms in it. Let's say it's an atomic crystal. Each atom has three degrees of freedom. It can move in each of three independent directions. And there are N of those atoms. There are $3N$ total degrees of freedom. That's how many vibrations the lattice has. So in fact, you have to multiply this by $3N$. So what would be seen is, you'd have $3R$, the gas constant, per mole, in other words. And if you looked at C_v over $3R$, then this value would be one. Very useful to see that. To figure that out. And people understood it.

What people didn't understand is this stuff. Why did it do that? That was a great mystery. Why did the heat capacity go to zero at low temperature? And the reason they didn't understand it is because the model for vibration was really simple. The lattice is a bunch of masses and springs. I know the vibrational energy. And I can calculate it. That is, I know the classical vibrational energy. Remember, the reason it goes to zero is this. It's all because of the fact that the energy levels are discrete, quantized levels, with gaps in between them. If I've got classical mechanics, then that means it's vibrational energy. The energy just gets bigger if the amplitude get bigger. It's not discrete, it's continuous. There are always energies in there. In that case, there's never a situation like this. Where kT is lower than the first available excited level, and everything's in the ground state. That never happens in classical mechanics, because there are always levels there. But it does happen in quantum mechanics. And Einstein recognized that this was a way to explain this low temperature limiting heat capacity. So actually if you look at early development of quantum mechanics, really it was all predicated on statistical mechanics. And the idea that, well, that you could then do the statistical mechanics with quantized levels, just the way we've done it. And what you immediately discover is, gee, it all makes sense. You go to low temperature, everything's stunk down here. And suddenly you've got zero heat capacity, because you changed the temperature

a little bit, and everything is still stuck back here.

So that's the vibrational heat capacity of a solid. That's what it looks like. And at moderate temperature, this doesn't have to be very much bigger than ϵ_0 . You're already in, essentially, the limit. The high temperature limit. And now for molecules, that limit isn't usually reached at room temperature. Here's a kind of calibration. kT at room temperature is about equal to 200 wave numbers. So molecular vibrations, you know, you've taken IR spectra they're typically on the order of 1,000 wave numbers or so. kT isn't bigger than the vibrational energy. But a crystal lattice, you know the vibrations are the acoustics vibrations. And those are much lower in frequency. If you have an atomic crystal, it just has the sound vibrations at all the different wavelengths that are available. They're never very high. So it's easy to get into the high temperature limit, in that case. Where you basically see a temperature independent heat capacity.

OK. By the way, this was actually used commonly to determine the molecular weights of molecular crystals. Because you've got a factor of the number of moles in there. If you ask how big is the heat capacity? Well it does depend on how many moles of material you have. Because it depends on how many atoms. That control how many modes there are in the crystal. How many vibrational modes. Because each atom has $3N$ degrees of freedom. So if you just weigh the whole crystal, and then you measure the heat capacity, you know how many moles there are, and now you know the weight, you can figure out the molecular weight.

OK. So that's the low temperature limit. And now I want to talk a little bit further about the high temperature limit. And in particular, I want to talk about both the heat capacity and the energy that we've seen. Namely, this thing. Of course it's really the same result for the energy and the heat capacity. They're obviously intimately connected. So these are the high temperature limits. Now it turns out that that high temperature limit doesn't just obtain for vibrations. But it's also the case for molecular rotations, translations. All these low energy degrees of freedom. So let's see why.

It has a name, that result. It's called the equipartition of energy. Sometimes it's called the classical equipartition of energy theorem. And what it says is one half kT per degree of freedom equals energy in the high T limit. And in particular, for translation, so $E_{\text{translational}}$ is $\frac{3}{2} NkT$ for N atoms. Or molecules. $E_{\text{rotational}}$ is, now it depends how many rotational degrees of freedom there are. If I've got a linear molecule, there are only two. It can rotate this way, and it can rotate this way. If I've got a non linear molecule, parts all over the place, it has three unique axes. It can also spin this way. And that's a rotational degree of freedom. Of course the linear molecule wasn't like that. Nothing is moving when you do that. So it's either NkT , linear. Or $\frac{3}{2} NkT$. Non linear. Great. And then $E_{\text{vibrational}}$ is NkT per vibrational mode. That's because vibrational energy is potential and kinetic energy, and it's $\frac{1}{2} kT$ each. Why does all that stuff happen? Turns out, you can see why pretty easily. All those degrees of freedom. Those classical degrees of freedom. If you say, let's write out an expression for the energy. Classical.

You know it's $\frac{1}{2} m v^2$. Or it's $\frac{1}{2} k x^2$. Or for rotational energy it's $\frac{1}{2} I \omega^2$. You see a functional form emerging here? That looks similar in all these cases? You know, $\frac{1}{2}$ times some constant, times some variable squared. Well, now let's look at our expressions for the energy. Average energy. We're going to sum over all the energies of ϵ_i , $e^{-\epsilon_i/kT}$. Over sum over i , $e^{-\epsilon_i/kT}$. That's just how we originally derived the average of energy. In other words, it's the probability of each state times the energy of that state. Summed up over all the states.

Well, now we have an expression for the energy. It's one of these things. It's something times a variable squared. So, we'll use some general functional form, $a y^2$. And now, let's assume we're in the high temperature limit. And we've seen what that means. It means kT is big, compared to the separation between the energies. When that's the case, there are lots of these terms in the sum. We can convert them to integrals. We can forget about the fact that the energies are discrete. We can say look, they're so close together, compared to kT . That we can turn the sums into integrals. So then we have integrals instead of sums. And here's our energy. It's $a y^2$. Whatever that is. It could be this, it could be this, it could be this. And it's not going to matter. $e^{-a y^2/kT}$ over, integral over $e^{-a y^2/kT} dy$.

OK. And here's what's going to happen. If you do this integral by parts. This one. What ends up happening is it gives you this integral times a certain number. This is going to come out of the integral. And so it will turn out. And it's straightforward to do it. It's in the notes. It goes through. But I'll just write the result here. The result is that you get, in this high temperature limit where you've gone to the integral form, you get $\frac{1}{2} kT$. And it will always be the case. All you need to know is the form of the energy. As long as that's the case -- Remember y is just a variable of integration here. That's not going to be preserved. This comes out because of what happens here.

So what that's telling you is whenever you have an energy of this form, and you're in the high temperature limit, then you get this classical equipartition of energy result. So, for translation, of course there are three separate degrees of freedom. For velocity in the x , y , or z direction. For rotation, if it's a linear molecule, let's say there are two separate degrees of freedom. You have to keep track of how many degrees of freedom there are. But that's all you have to do. That's enormously powerful. It means that without doing anything, I know the average translational energy of the molecules in this room. Because of course I'm certainly in the high temperature limit. With respect to translational energy levels, they're really closely spaced. Same with rotations. Now at room temperature vibrations, forget it. I have essentially no vibrational energy. I'm at the low temperature limit for the molecular vibrations of nitrogen or oxygen. Those are high in frequency. So much higher than 200 wavenumbers. But for each molecule, I have $\frac{3}{2} kT$ of translational energy. Linear molecules, I have kT of rotational energy. Without doing any work at all. And since that applies to most molecules at room temperature, it's an incredibly useful, very, very general result.

OK. Next time we'll do a little bit of chemistry, and look at phase transformations.