PROFESSOR: And last time, you got to see how you can derive macroscopic thermodynamic results from the microscopic point of view of statistical mechanics. So in contrast to the way we'd gone through the course from the beginning, starting from empirical macroscopic observation and ideas through macroscopic laws and then working out the consequences of them, in statistical mechanics you start from a microscopic model, build up the energy levels that are going to give you the terms in the partition function, determine the partition function, and what you saw last time is that given that, you can calculate everything. You can calculate all the macroscopic properties that ordinarily come from the thermodynamic laws that were based on empirical macroscopic observation.

So today, I just want to continue doing some statistical thermodynamics. Basically going through a few examples, just to see how it plays out when we calculate thermodynamic quantities, based on our microscopic picture. So let's just look at a few cases. And I'll start with a couple of examples that are entirely entropy driven. And we've seen them before from the macroscopic point of view. What we should hope, of course, is that we can derive those results from the partition functions that are appropriate here.

So, we'll look at entropically driven processes. And the first one, the simplest one, maybe, to cover, is just free expansion of a gas. So what happens if we've got particles in a gas that are enclosed in a certain volume. And now we let the volume expand into vacuum on the other side. Of course, you know what'll happen. The volume will be filled and we've derived the thermodynamics for it before. So let's see what happens in a statistical mechanical treatment.

So let's say, here's V1, there's our gas. There's vacuum. And now we'll open it up and just have the bigger volume, V2, and let the gas fill it. So I think last time you got introduced to basically a lattice model for translational motion. And I started in on this just a little bit the time before also. So this is a simple way to describe the statistical mechanics of filling an open volume. Filling a space with molecules. And all it does is say that we're going to divide up our available volume into little bits that are basically the size of an atom or a molecule, or whatever the particle is in the gas phase. And say, OK, there are maybe 10 of the 30th of those little volumes. The little volume elements are going to be on the order of an angstrom cubed if it's an atom, a little bit bigger if it's a molecule. And the available volume is on the order of meters cubed. So that works out, an angstrom is 10 to the minus 10 meters. And then the cube of that is 10 to the minus 30th. So if you look at the total number of little volume elements of this sort, it's on the order of 10 to the 30th.

So, the total volume, capital V, let's make that distinction clear. And in this case, all the states of the system have equal energies. So in other words, it doesn't matter we're dividing our volume up into imagined little volume elements. And for both the molecular and the system states, the energy is the same. It doesn't matter whether the molecules are here, here, here and so forth. They're not interacting in this picture at all. So it doesn't matter how close or far they are from each other. They can't be in the same volume. So all the energies are the same. And what that means is, in the partition function, which is a sum over all these terms with these Boltzmann factors. That have e to the minus energy over kT. But energy's all the same in every one of the terms. So in order to determine the partition function, to determine Q, all we have to do is count up the total number of states that are available.

So, our energy, our molecular translational energy, we'll just set it to zero. Same with our system translational energy. We'll set that to zero. All the microscopic available states, that is, if I take an individual particle and I say where it can be, all those states have the same energy. If I take a microstate of the system, you know, the whole collection of the particles, all of those states also have the same energy. So what that means is that little q is just equal to capital V over little v, right? How many states are there? Well, I've got my little volume, and then however many individual cells there are, that's the number of available states in this model. So it's big V over little v, on the order of 10 to the 30th. And then, big Q, the canonical partition function for the whole system, it's something that we've been through before. It's this process of saying, well, you take the first atom or molecule, and it has any one of these possible states. So it's on the order of 10 to the 30th possibilities. Then where does the second one go? Well, there are basically 10 to the 30th more possibilities. And the third one has 10 to the 30th.

Since there are so many fewer atoms or molecules then there are volume elements, when we're dealing with the gas phase, we don't have to worry about the fact that, well the first million of the atoms filled some of the sites. So the next ones have fewer sites available to them. It's true, but it's such a small fraction that are ever filled that we don't need to worry about it. So Q is just little q to the capital N power, the number of particles. And then we've seen you have to divide by N factorial to avoid the overcounting of configurations that are in fact not distinguishable. The whole idea that maybe there's an atom here and another one here. That configuration is indistinguishable from the configuration with those two atoms reversed if we're dealing with a mole of identical atoms.

So, OK. There's our capital Q. And that's also our system degeneracy. So the degeneracy, the little g for the molecular degeneracy, how many molecular states are there of a certain energy. Well, it's the same thing as q. And it's the same thing here for the system states. The capital Omega is just equal to that.

Well, now let's look at what happens when we do this. When we expand from volume V1 to volume V2. So that capital V is going to change. So capital V1 goes to capital V2. And we know it's entropically driven. Let's calculate

the change in entropy. So delta S is just k log capital Omega 2 minus k log capital Omega 1. Now, in the process of just seeing how the development goes, how you can derive all these thermodynamic quantities from the partition function, one of the really most central results concerns the entropy. How you can describe the entropy in terms of the probabilities that the different states are occupied. And how, in the case like this, where the states all have the same probability of being occupied, then you have this very simplified form for the entropy. Just, if S is k log capital Omega, where capital Omega is the degeneracy. The number of system states of that energy. An amazing result.

You told them about Boltzmann's tombstone and so forth. So it's on there. In fact, the ill acceptance of that result kind of led to the tombstone being erected when it did. Boltzmann, depressed over the lack of acceptance of his theories, put himself to an early death. And this was off a big part of the reason. But we, generations later, have come to accept the results that concerned him so deeply.

So this is our change in entropy, when we just have this expansion of gas from V1 to V2. So it's just k log omega 2, omega 1. So now let's just put in our results for the volumes. It's k log V2 over v, I'm exaggerating the size there. To the N, over N factorial, over V1 over v to the N divided by N factorial. So this is going to turn out to be a relatively simple case, because the factorials are going to cancel. So, then we just have N k, log V2 over V1. Terrific, right?

Now, remember, k, the Boltzmann constant, is just the ideal gas constant per molecule rather than per mole. So this is the same thing as little N, the number of moles, times R, times log V2 over V1. And that should be a familiar result. That's the change in entropy in expansion, free expansion of a gas, from V1 to V2. So now we've been able to derive that just based on this really simple molecular picture. Based on that plus the result that Boltzmann fought and paid so dearly for, so that we would have it and understand it. As before, of course note it's greater than zero. If the volume, if we're expanding into a bigger volume than before. The entropy goes up. Also notice one of the other results that you've seen is that you could relate the partition function to A, right? The Helmholtz free energy is minus k T log Q. And S is negative dA/dT.

And so this would immediately give us the same result, because omega's going to go in here. So we could get this result from another pathway, too, but a more simple and direct way is just to start directly from the Boltzmann result for entropy. Any questions?

Alright, let's go one step more complex. Let's now look at the entropy of mixing. So now we've got two species. Rather than one expanding into a free volume, and we're going to open a barrier between them and see them mix. And again, it's still entropically driven. And we should be able to calculate the entropy change that we saw before from a macroscopic perspective. So, let's take a look.

So now, we have N_A molecules of A in some volume V_A. And N_B molecules of B, in some volume V_B. And then we're going to open it up. Then we've got capital N, N_A plus N_B. And capital V. V_A plus V_B, right? And the whole thing is happening at constant pressure. Well, then let's look at the initial and final expressions for the entropy. So, S_1, at the beginning. It's just $k \log \omega_A$ plus $k \log \omega_B$. And that's $k \log V_A$ over little v to the N_A power. Divided by N_A factorial times the same thing for B. V_B over little v to the N_B power over N_B factorial. So that's our initial entropy. The sum of the two entropy contributions, from the two sides. And S_2, the final one, is just $k \log \omega$ for the whole shebang. And that's $k \log$. And now we have capital V over little v to the N power.

And now we have a combinatorics result that I think probably you're all familiar with from one context or another. That is, now we have to divide by N_A factorial times N_B factorial. In other words, the amount that we need to divide by to avoid overcounting is the product of those two factorials. Just in case it's been a while since you've seen stuff like that, I've got a couple of pages further on the notes at the bottom of the page. I've just broken that out for a really simple example where there are just two molecules of A, and two molecules of B represented by different little balls in lattice sites. And I just worked it out for the total where there are only ten total in the mixture.

The point is, now what happens is, when you start filling a lattice, but it's not all the same, right? So now you have A here. And A here. And B here. And B here. So, of course if you interchange A and B here, you don't have an identical state. An indistinguishable state. That's distinguishable and needs to be counted. It's only interchanging B with B that you need to correct for, and interchanging A with A. So of course you need to account for that in a way that's different from how it turns out if every one of the molecules is identical. So it's the product of these factorials. In the numerator.

So now delta S is then just this minus this. So it's $k \log V$ over little v to the N_A plus N_B over V_A divided by little v to the N_A power V_B divided by little v to the N_B power. And, let's see, maybe I'd better back up. Let me know, that's all. And what's happening is, again, luckily, these factorials are canceling. So we just have $k \log V$ to the N_A, V to the N_B over V_A to the N_A V_B to the N_B.

And now we're going to go back and use the fact that the pressures are the same, all the way through. And what that means is that the volumes must be in the same ratio as the number of molecules. And what that means is that, for example, V_A over V is the same as N_A over N. And that's just equal to X_A, the mole fraction of A. And the same for B. So we can substitute that. And know our delta S just becomes $k \log$, let's see, let me break this out from the beginning. And just take minus $k \log$ on X_A to the N_A power minus $k \log X_B$ to the N_B power. All I've done make these substitutions here. For the ratios of the volume. So it's just minus $N k$ $X_A \log X_A$ plus $X_B \log X_B$. Look familiar? Same thing we saw macroscopically before. So, again, we can still use our microscopic model. And continue to derive the macroscopic entropy changes. And, of course, for many of these we can still get our delta G

and so forth. Of course, it's only entropy that's going to contribute.

OK, let's look at just one more entropy driven problem. And that is, it's a little bit different from these. Let's look at the entropy mixing in a liquid. And the difference between the gas and the liquid is that in the case of the liquid, every single cell is filled. It's not like with the gas. And that makes a difference in how the states need to be counted. Because here, remember how I said when we were filling this lattice. Well, the first molecules goes somewhere. And there are 10 to the 30th possibilities. And the second molecule goes somewhere and there are 10 to the 30th possibilities. And there are always 10 to the 30th possibilities because there's so few molecules that you never have to worry about the fact that it's getting filled up, so there would become fewer possibilities for that last molecule. And the reason again is because there are so few molecules that essentially, all the cells are open and available even to the last molecule.

Because maybe there's a mole of molecules. Maybe there are 10 to the 24 molecules. And there are 10 to the 30th lattice sites. So there might be one in a million lattice sites occupied, even at the very end of the procedure. But of course for the liquid, that's definitely not the case. All the available volume is filled. So by the time you get to that last molecule, it has one and only one space it can go into. So you have to count for the diminishment of the available sites as molecules are placed in them.

OK. So it's a different kind of combinatorics problem that results. So, OK. Let's look at liquid mixture. So now, it's going to look like, and they're all filled up. I sure used fewer sites maybe, but that'll be easier here. We've got open circles. So these two are going to mix. And now we're going to have filled mixture. So we're going to have a bigger volume, but it's still going to be filled and all the positions are random. So we have to consider all the possible configurations that we have.

OK. So we still should go through the same part of the procedure. So let's call this A. And this, B. So at first, already the situation is different. If we say what's the entropy of system A. Well, in this simple model the entropy of the pure liquid is zero. Now, of course, that's not realistic. But it's going to turn out to be suitable. Because the change in entropy that we go from here to the mixture. The other contributions to entropy are going to more or less be the same from beginning to end. What's going to be significantly different is simply the fact that in the mixture you have the random positions that can be occupied. So that's going to be the term that matters, the contribution that matters. So entropy is just k log Omega A. There's only one system state. All of the lattice sites are filled, and they're all filled with A. So there's only one way to do it. And of course, it doesn't matter if you interchange them. They're indistinguishable. So this is zero. This is one. There's only one state. Same, of course, for B.

So that's our initial entropy in this simple model. Well then, delta S of the mixture is just S of the mixture. We don't

need to worry about the original parts. Now, this is going to be k log of N factorial over NA factorial NB factorial, like we've seen before. That's not different from before. Except that unlike before, we don't have the convenient and simple cancellation of the factorials. That happened before, because they were there. In other words, they were there in the entropy contribution in the gas phase as well, in the pure gas. They're only there here in the liquid. And that means we have to deal with them. But it turns out it's not too bad to deal with them. And we're going to use this Stirling's approximation that I introduced before.

So the log of N factorial is approximately equal to N log N minus N. So let's just break that out then, and use the Stirling's approximation for each of the factorial terms. So then we have S for the mixture is N k log N minus N k minus. And now we'll do the bottom, the numerator. NA k log NA minus NA times k, that's this part. And now we'll do this one. So it's plus NB k log NB minus NB times k. But NA and NB, of course, are just N. So these will cancel. So then we're left with NA plus NB times k log N minus NA -- And I just want to write this this way so that I can then separate the terms. Minus NA k log NA minus NB k log NB. And now I just want to combine the easily combined terms. So it's NA k log N over NA plus NB k log N over NB. Looks like we're home. Now we just have N k XA log XA plus XB log XB right? This NA over N is just XA. NB over is XB. And then I've divided the same thing to make XA's over here. And take the total number of moles out. Alright. Look familiar?

Again, the same thing derived now in a way that's a bit different from what we did in the gas phase. So now I've got the entropy of mixing even in a condensed phase the liquid. And again, the reason this simple model works is because although a real liquid, certainly a pure liquid has a finite entropy, a substantial amount of disorder, that's present, that kind of disorder. In other words, the fact that you don't really have the molecules sitting in sites on a lattice but there's disorder. There's also rotational disorder if it's a molecule. There are various contributions to entropy. But those are all comparable in the starting and final states. The thing that's changed is simply the interchanging of A and B molecules. That's introduced a new kind of disorder that didn't used to be present. A very big difference in the number of states at the end from the beginning. And that's what we've calculated and captured here. And that's why such a simple model still works OK. Any questions about these entropy driven cases?

OK. Now let's move on and talk about the cases which are more commonly encountered, where the states aren't all the same energy. Of course, here in the liquid, just like in the gas, we haven't treated interactions between the molecules. We've assumed that they're equal between A and B, as they are between A and A and B and B. So in this case then all the energy to the states are the same and it's just a counting problem. How many states are there available. As soon as the energies are different, then of course we need to account for all those Boltzmann factors. Those e to the minus energy over kT factors become part of the, they weight the counting. And we have to do that.

So, what I want to do is go back to this simple polymer configurational model that we introduced before. And this sort of model, in one form or another, we're going to see a few times in the rest of this treatment. The reason, really, is because it's a simple and also realistic way of portraying a system that has a finite number of well-defined energies. If you don't have that, of course you could do the treatment. In a, sort of, classical mechanic sense, a continuum of states, all with different energies. It's hard. Because then those sum in the partition function become integrals over all the different configurations and possibilities. It's much more complicated to do. It's much more straightforward if you can identify discrete states and add over them.

Now, someday you'll take quantum mechanics. And then you'll see that the states are discrete, even if we're talking about translation or rotation and so forth. So you'll have discrete quantum mechanical energy levels and so forth. You haven't had that yet. And so that's not the starting point that we're going to use. Instead, we're going to use a starting point that goes through the exact same formalism, which is just the discrete states available in molecules that have multiple configurations.

So we're going to have unequal energy states. And what we're envisioning is molecular or polymer configurations. So here we have our states. And just like before, the idea is that if you have two sub-units that are in proximity, there's some kind of favorable interaction. It could be hydrogen bonding. Could be different. But the point is, there's some sort of interaction there that reduces the energy.

And then there are other configurations that just don't have that. Because the sub-units aren't in proximity to each other. And it's not hard to work out that in this very simple model, these are the only configurations that are available. The only distinct configurations. So then, our molecular energy, E, we can define as zero here. Before we defined it as minus E int, but I'm going to, it's a little more convenient to make this the zero of energy, and then this is plus the interaction for each of these states. Of course, we can put the zero of energy wherever we prefer. And then we have the degeneracy is one. And in this case it's three.

So now we can just write out the configurational partition function for the molecules and also the canonical partition function for the system. So q configurational. And we're just going to sum over the states. So it's e to the zero over kT, for the lowest state. And then there are three states that are going to have this term, e to the minus E interaction over kT, where E int is a positive number. In other words, the probability of any one of these states is a little bit lower than this state. Because this state has lower energy. Remember what I mentioned earlier, which is although the probability of any one of these states is lower than the probability of this state, the probability of this energy is likely to be higher than the probability of this energy. Because there's only one of these states and the degeneracy here is three. So there are three possibilities in which the molecule could have this energy, and only one this. So if, basically, kT is bigger than E int, so in other words, if this term isn't very much smaller than one, then of course this will be bigger than this.

So of course this is one plus three e to the minus E int over kT. And now we have our capital Q, our canonical configurational partition function. And that's just little q configurational to the Nth power. And like you've seen so far, in various cases where the molecules are separate non-interacting molecules, this is just the molecular partition function taken capital N times. If all the molecules are behaving independently, then you sum over all of those states. And you see that each one of these is just taken again and again. As we've seen implicitly in all these treatments so far as well.

Now, in the translational case where you interchange particles. Like particles, you have to divide by N factorial. . But the different configurations don't do that. If I've got a system where a molecule over here is in this configuration and a molecule somewhere else is in this one, and now they change configurations, that's a distinguishable state. So there's no N factorial involved here. In the configurational partition function. So, fine, then it's just one plus three e to the minus E int over kT to the Nth power. That's Q. And once we know Q, as you've seen, we know everything. And so we can immediately start in deriving all of the thermodynamics, right?

And the place to start is A, Helmholtz free energy, or configurational free energy. Because that's the simplest relation, just as minus kT log of capital Q configurational. So it's minus N kT one plus three e to the minus E int over kT. I'm going to rewrite A configurational in terms of beta rather than kT, just because there are going to be a lot of these factors. So it's minus N kT one plus three e to the minus beta E interaction. And that's our A. Everything's going to follow from that.

Before I write some of the other results, one thing to notice is, it has N in it. In other words, there's a free energy per molecule. The total free energy is just something times N. And that's because all the molecules in this model are independently behaving. So whatever their average free energy is, that's going to add up and give us the total. They're all acting independently. And in fact, we could have gotten this directly from minus kT log little q configurational. And if we look at energy, regular energy, u, configurational, it's minus d log capital Q with respect to beta. V and N. And I'll just write the result. It's N times three E int e to the minus beta E int over one plus three e to the minus beta E int.

And once again, the feature I want to point out is that there's a factor of N here. Again, there's an energy per molecule. And so if we want, we can also just write the average E, little E, configurational. It's just u configurational over N. It's the same as this without the factor of N. Not only that, again, we could get this directly from the molecular partition function up there. From little q configurational. We'd have the same result. Exactly as it should be. So if we wrote E as we've seen in the past, it's just the sum over that energy times the probability for each state. And that's just the sum of Ei, e to the minus beta Ei over little q. And if you put in the terms, you immediately get this. Here's our little q, and this has the interaction energy brought out, multiplied here.

So the point is, in a case like this, where you have a bunch of independently behaving particles, the totals for these, for quantities like energy, are simply for the system energy, it's just the average particle energy times the number of particles. The number of molecules. And remember, before, I spoke a little bit about the fact that, well, if you look at one individual molecule and another and another, the energy will fluctuate. It'll vary considerably. Of course, this is a simple case with only two possible energies. But in a case where there may be many more possible energies, the molecular energies may vary quite widely. Still, there will be a well-defined average. And then the system energy would simply be the total number of molecules times that. Where they're all independent. Where all the energies are independent of each other. And, of course, the system energy fluctuates a great deal less than the individual molecule energies. OK, so that's our energy term. And I think I won't write out the results. They're in your notes for entropy. For chemical potential. They're all there. And again, the same point would hold. They all scale with N. But I want to talk a little bit about the heat capacity. The expression for it isn't so simple. So let me just right Cv configurational. It's du configurational / dT, at constant V and N. And again, the details are worked out in the notes. So I won't write out the whole derivation of it. And even writing out the results, I'm almost reluctant to do it. But I'll put it up. Three E int over k T squared, because there are some things I want to point out about it. N, and then one plus three e to the minus beta E int minus E int e to the minus beta E int. Minus e to the minus beta, I have to get rid of this. Minus e to the minus beta E int times negative three E int e to the minus beta E int. All over one plus three e to the minus beta E int, that whole thing squared. And this is just, of course, the kind of messy results that comes from taking this derivative with respect to temperature and doing the chain rule to get it with respect to beta and so forth.

So it looks like a little bit of an intractable, or at least a little bit of a complicated, function. And the detailed functional form is complicated. But what I want to emphasize is that it has simple limits that are very easy to understand physically. And that are important to understand for lots of systems. So I just want to look at the limits of the heat capacity at low and high temperature. And this is something that recurs in statistical mechanics, in an enormous number of systems where you have simplified limits. And they're really important. Because what's going to matter is this. Maybe I shouldn't have covered up those configurations. So there's a big difference in what happens when kT is much bigger than this energy. Where you know that under those conditions - let's say it's hot, and this is a small energy difference. Then the probabilities are essentially equal. That the molecules are in this state or in any of these states. Because the energy is so tiny compared to the thermal energy. The molecules are continually being kicked around between all the states, among all those states. So the high temperature limit, physically, is one that's simple to understand. Where now you're really just going to revert to the cases that we've treated before, where all the energies are effectively the same. Because compared to kT, they are. And in the low temperature limit, now go to the opposite limit. Let's say kT is much, much, smaller than the interaction energy. So that now this term is really small. Because this is much bigger than this.

So, what it means physically is all the molecules are in the ground state. The probability of this is basically one. The probability of being in any of these states is zero. And that's also a simple result. And there are lots and lots of cases where one of those limits really obtains. If you look at molecules moving around in the gas phase or in a liquid. And you say, well, OK, let's think about their rotational motion. Rotational energies are small. At room temperature, kT's much bigger than them. You immediately go to the high temperature limit. Now let's go to a small molecule and look at the vibrational energy. In most cases, the vibrational frequency is pretty, molecules are pretty stiff. And in many cases you can just say, look, forget it. All the molecules are in the ground state. And again, in the opposite, but also simple limit ends up holding. Or molecular electronic states, right? If you've got benzene or hydrogen atoms in, say, at room temperature, how many hydrogen atoms thermally are going to be up in the n equals two state, the 2p state or the 2s state? Forget it. There's not nearly enough thermal energy to do that. So these simpler limiting cases play a huge role in simplifying statistical mechanics and the calculations from them generally.

OK, so let's just see what happens. So, this we want, so this we can move down. So our limiting cases, when T goes to zero, not going to work out what happens to all the betas, beta gets big in that case. But the result is that Cv configurational goes to zero. That's the limit where, like I've described, here is E interaction, or E int. Here is zero. And your kT is barely above zero. So when all the molecules are here, none of them has enough thermal energy to be up here.

So why should the heat capacity be zero? The heat capacity is du/dT. It's zero because, here's kT. Let's say I increment kT up a little bit. I just heat the system a tiny, tiny bit. An infinitesimal amount, to look at the derivative. Once I do this, how many molecules are in this higher level now? Still zero, right? There's still not nearly enough thermal energy to have any molecules up here. So what, was the derivative of the energy with respect to the temperature? I changed the temperature. The energy didn't change a bit. And that means the heat capacity is zero.

Now let's look at the other limit. High temperature limit. Really, it doesn't need to be infinity. It's really just T greater than, much bigger than, kT is much bigger than E interaction. And in this case, kT is much less than the interaction energy. So it doesn't need to be nearly that extreme. Well, what you find out is the heat capacity is zero. Now, it's zero. Because we're in the following limit. Now kT is way up here. Compared to kT, E interaction is way down here. So, what happens? It means that this is so hot, that term, The e to the minus E int over kT, forget it, it's one.

And so the probability, in other words, that the molecules are in this state or this state are essentially equal. So now let's say I raise the temperature a little bit higher. What happens to those probabilities? What changes? Right. Nothing changes. So what happened to the energy when I raised the temperature? Of course, nothing happened

to it. Which means the heat capacity is zero. Now, the first limit that I described, this one is almost universal. For any system where you have quantized level, you can always eventually get to a low enough temperature that you're in the first limit. Where the kT is lower than, by far, than the first excited level. Everything's in the ground state. And so you reach this limit. This one, it's only cases where you have a finite number of levels. Then you have the same high temperature limit to it. In other words, the reason this result happens is because there aren't a bunch of other levels up here that eventually get to be comparable to kT. And not all kinds of systems or degrees of freedom are like that. This one is, because there's a finite number, four in this case, of configurations. So there's just nothing above there.

There's another really important kind of degree of freedom like that. And that's spin. Think of proton spins. It's plus 1/2 or it's minus 1/2, and that's it. You can't put any more spin energy into it. Just like you can't put any more configurational energy into the system than to be in this state. That's it. So in other words, the maximum possible energy is finite. Of course, lots of other degrees of freedom are different from that. If you think a molecule's rotating, they could always spin faster. Vibrating, they can always vibrate harder, translate faster. Those degrees of freedom won't have this limit. They also will have a simple high temperature limit, but not zero, because, of course, if there are always more levels, and I keep increasing kT, then I'll have thermal energy to go into those higher and higher levels. It'll still go up. But for systems with a finite number of possible levels, and a finite amount of total energy, for degrees of freedom like that, once I get to a temperature higher than any of that stuff, then forget it. You can't change the energy anymore thermally. So your heat capacity is zero. You can change the temperature and nothing further happens.

OK, next time we'll see what happens when you do have continuing basically unbounded possible energies.