

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: Last time, under the tutelage of Professor Blendi, you started to see some of the consequences, the more immediate consequences, of some of the properties we've been discussing over the last few lectures of binary liquid mixtures. And in particular he started showing you some of the colligative properties by which you see changes in the vapor pressure and the boiling point and so forth, that arise because you've got liquids in mixtures and this affects their chemical potentials. And today I want to start by going through the fourth colligative property that's commonly seen, what called osmotic pressure. And I mentioned this before. Today we'll go through a quantitative calculation of its effects. And then we'll start in on a different unit of the course. We'll start in on statistical mechanics. So let's continue with the colligative properties.

And osmotic pressure plays a super-important role, certainly in all of biology, because of course inside biological cells you have solutions. Water is in there, along with a whole bunch of other components. And so the chemical potential of the water is strongly affected by all the other constituents. And this has profound effects on what's involved in maintaining an equilibrium between what goes on inside and outside of the cell. And this is the case in all organisms. And it's also exploited of course in all sorts of other contexts. So what's happening with osmotic pressure is, you've got some situation where there's a semi-permeable membrane. It allows one of the constituents but not others to pass through. And typically that's water. Although, of course, it could be anything.

So, in the lab the way this typically looks is something like this. Here's your semi-permeable membrane. Outside, we can imagine that there's pure liquid A. Could be water, could be something else. Inside, you've got A mixed with B, one or more other constituents. And then there's some level of the liquid outside the container. Of course, A can pass between, pass through, the membrane. And so what ends up happening is, if you just look at the situation, what you immediately notice is that the liquid in here goes up higher than the liquid around it. It's at a higher level. It seems counterintuitive at first. And in fact it can be enormously higher. But that's what happens. So there's extra height, h , of liquid inside the container relative to what's outside it. And what that means is the pressure inside here must be higher, right? Of course, the weight of this liquid is pushing down on what's down here. So the pressure inside here is higher than the pressure outside. But of course, from what you know already about liquid mixtures, this should in fact make sense. Because what's happening is that if these things are at the same pressure, you know that the chemical potential of A is going to be lower in the mixture. It's always lower in the mixture than it is in pure A. So there's nothing to stop A from rushing in. So A does rush in. If the liquid level goes up, it raises the pressure down here. But as the pressure rises, now the chemical potential changes and

goes up. So eventually, equilibrium is reestablished. Because these aren't at the same pressure. The additional pressure means the chemical potential of A goes up. Of course, it goes down by virtue of being in the mixture. And those two things will eventually reach an equilibrium where the chemical potential of A is the same everywhere, as it has to be.

So what we'll calculate now is, OK, how much does the pressure go up. When that happens. So let's take a look. And so of course, there are all kinds of examples. In a typical one, A would be water. And B would be one or more sugars. All sorts of other potential constituents also. OK, so let's just consider the pressure at two places that are the same level. The same height. We'll label this one point alpha. Let's put some color in here. We'll have points alpha and beta. And we'll make this distance l .

So now let's just first calculate what the pressure is at both points. It's going to be higher here, because of this additional liquid on top. So, if we look at point alpha, well, let's start at point beta, actually. That's the simpler one since it doesn't have the additional liquid. And first of all, there's just atmospheric pressure. So there's an external pressure. Certainly, typically, it would just be atmosphere pressure. It's pushing down here. It's also there. So that has to factor in. So there's our external pressure. And then there is the product of this height times the density of the solution. Times the gravitational force constant. And I'll work through that momentarily so you see that that indeed is the pressure. That's exerted by the amount of liquid that's up above this level.

But before I do that, let's put the pressure at point alpha. So p_{α} . And that's $p_{\text{external}} + l \rho g$. But now, there's this additional height on top of that. So there's also $h \rho g$. There's an extra pressure. We'll label this π . So we can write p_{α} is equal to $p_{\beta} + h \rho g$. Which we'll say is equal to $p_{\beta} + \pi$. Now, just to make sure we're all on the same page about these quantities being the pressure that's exerted by the liquid, let's just look at that. So $l \times \rho \times g$. It should be a pressure which is a force per unit area. And in particular, this is going to be a weight per unit area. Remember the weight is the force that a mass exerts. And it exerts it because of the gravitational constant. So what happens? Well, this is just a length. ρ is a density, right? So it's mass per unit volume.

And of course, g is just our gravitational force constant. So it'll be in units of meters per second squared. So what's going to happen? If we multiply $l \times \rho$ together, this is units of length. This is units of volume. What we're going to get then is mass per unit area. And that's what it is, of course. It's distributed across the area of that surface where point alpha is. So these two together, $l \times \rho$, give us mass per unit area. So it could be in units of kilograms. Per unit area, per meters squared. And now we're going to multiply that by the gravitational force constant. So now we're going to have all together kilograms per meter second squared. In other words, it's force per unit area.

And so that's what we have. And now, of course, we have the additional component. Due to the height here. Now, what do we know about the chemical potential in both parts? The chemical potential has to be equal outside and inside the container. So μ_A at point alpha is just μ_A of the liquid at pressure p plus p_i and temperature T . And that has to be equal to μ_A at point beta. But point beta is just in the pure liquid. So that's just μ_A^* of the liquid at pressure p and temperature T . So here's the chemical potential at point beta. Here it is at point alpha. Those two have to be equal to each other. And now we're just going to use Raoult's law. And what does that tell us? We have $RT \log x_A$ plus μ_A^* of the liquid at pressure p plus p_i and T . So there is our chemical potential inside the container. And that's equal to μ_A^* of the pure liquid at point beta.

And so we can just rewrite this as $RT \log x_A$. Plus μ_A^* liquid at pressure p plus p_i of T minus μ_A^* at pressure p is equal to zero.

OK., now we need the pressure dependence of the chemical potential. But we certainly have an expression for that. So we know that dG is $S dT$ plus $V dp$. The temperature is the same on both sides. But we need to worry about what happens at different pressures. So at constant T , dG is $V dp$. So now if we look at the chemical potential, which is just the Gibbs free energy per mole, then $d\mu_A^*$ is $V_A^* \text{ molar } dp$. In other words, the difference in the chemical potential is, this changes as a function of pressure. Is going to be given by the molar volume of A under these conditions in the pure liquid dp . It's just the potential, the pressure dependence of G . It's not the kind of topic I would have thought would drive away a lot of people. But, you never know. OK.

For the rest of you who are willing to bear with me, let's continue. So now let's just integrate. We need to know the change in μ_A at a finite jump in pressure from inside to outside the container. So we're just going to integrate. So if I want μ_A^* at pressure p plus p_i minus μ_A^* at pressure p , then I just have to integrate from p to p plus p_i $V_A^* \text{ bar } dp$. And now I'm going to assume, and this is certainly a safe assumption, that this quantity this, molar volume, isn't going to change. In other words, the molar volume of liquid A isn't going to change significantly, going from the pressure out here to the pressure in here. On the grand scale of things it's a small pressure change. We can assume that the liquid is incompressible over that small pressure change. Which means that this is constant, since all we have then is simply $V_A^* \text{ bar } \times p_i$.

So then, substituting back here, so now we've calculated this difference, it's a simple result. So then we simply have that $RT \log x_A$ plus $V_A^* \text{ bar } p_i$ is equal to zero. Now, you saw last time, and I'll work through this quickly again, that the log of x_A can be written approximately as minus n_B over n_A if you remember, this came from writing that this is equal to the log of one minus x_B . But we're assuming that B is the minor constituent here. So x_B is a small number. So that this is roughly equal to minus x_B . And that's just minus n_B over n_A plus n_B . And since n_B is much smaller than n_A , this is approximately equal to minus n_B over n_A .

So we're going to make this substitution. Also, this quantity, this molar volume of A over the pure liquid, since the concentration of B is low, we can assume that this is just the molar volume of A in general. So, in other words, we can write \bar{V}_A^* . We can just write it as \bar{V}_A . That is, we're not going to worry about changes in the molar volume, either as a function of pressure or a function of concentration at the low concentrations that we're working. And then, note that n_A times \bar{V}_A per mole is just the total volume. Oops, not \bar{V}_A . Not the molar quantity. This is the molar quantity multiplied by the number of moles. So it just gives us V_A , the total volume, occupied by A. And using these two results, then we have that RT times negative n_B over n_A , that's this result. Plus V_A over n_A times p_i is equal to zero. n_A will cancel. And finally, since almost all the volume is due to A, again because B is the minor constituent, we can approximate further that V_A is approximately equal to V . So finally getting rid of the n_A on the denominators, we're left with a simple expression. p_i times V is $RT n_B$.

So we have a very simple expression which is called the van't Hoff expression, and look at how it resembles the ideal gas law. This is the pressure times the volume equals the number of moles times RT . Of course, really it's a change in pressure. And this number of moles is the number of moles of a guest constituent in a solution. But it has the same form that you're familiar with. We also, sometimes it's convenient to substitute, since n_B over V , that's just the concentration. That's the number of moles over the volume.

And in that case, the expression is rewritten as p_i is RT times c , where it's understood that this is the concentration of the solute. OK, any questions?

Let's work through a numerical problem. These things can seem a little bit, they can seem simple, but sometimes when confronted with a problem, sometimes it may seem less than obvious what to do. And also the results can seem a little bit surprising. So I just want to work through a simple numerical example to see how these play out. So let's try, let's imagine we have 10 grams of unknown solid. Dissolved in 1000 grams of water. The vapor pressure is 25.195 torr at 27 degrees C. And the pure water vapor pressure is 25.000. No, .200 torr at the same temperature.

So first, let's just calculate the molecular weight of the solid. Now, this you I think, saw from last time. So this is straightforward. It's going to come from the vapor pressure lowering not from the osmotic pressure, which we'll do next. So the Δp of H_2O is minus x_B times p^* of H_2O . Alright, this is from last time. The expression for the vapor pressure reduction. Due to the concentration of the mole fraction of B. And then x_B is 0.005 torr over 25.2 torr. Right? And that's 1.98×10^{-4} . So that's our mole fraction of B. And then x_B is n_B over n_A plus n_B . So, now we have this, it's 10 grams divided by the molecular weight, which is in grams per mole, over 1000 grams of water. Divided by 18 grams per mole plus 10 grams divided by the molecular weight in grams per mole. But since the solution is overwhelmingly water, we can get rid of this term and make the math simple. Can you see that? I guess barely.

So this turns out to be about 1.98×10^{-4} . And if we solve for molecular weight, then we wind up with 907 grams per mole. So that's using one of the colligative properties that you saw last time. Now let's, though, finish up by calculating the additional pressure that we'll find. So what's the osmotic pressure of the solution. So now we're going to assume that we have that solution inside, in a situation like this where there's pure water outside that can go in. And there's going to be excess pressure because of the water rushing into the solution. So what happens? And let's use a value of the density. 0.995 grams per centimeters cubed.

So, π is RT times c . That's 0.08314 bar per Kelvin mole. Times 300 K times c . c is n_B over V . So it's 10 grams divided by 907 grams per mole, over 1010 grams over 0.995 grams per milliliter, I'm equating with centimeter cubed. And that turns out to be 1.09×10^{-5} mole per milliliter. Let's get that into liters. So we'll multiply by 1000 milliliters per liter. So it's 1.09×10^{-5} moles per liter. OK, yeah.

STUDENT: [INAUDIBLE]

PROFESSOR: Let's see. Oh, don't forget that, right.

STUDENT: [INAUDIBLE]

PROFESSOR: Ah, yes, yes, yes. And I think that's the case in the notes also. Yep. Thank you. OK, so the concentration, just to calibrate us here, it's a pretty low concentration, right? It's about 10^{-5} moles per liter. The excess pressure is 0.27 bar. In other words, about a quarter of an atmosphere. Uh-oh, what's happening?

STUDENT: [INAUDIBLE]

PROFESSOR: Where are we?

STUDENT: [INAUDIBLE]

PROFESSOR: Sorry, where exactly?

STUDENT: [INAUDIBLE]

PROFESSOR: Oh, yes. Yes, sorry. Yeah. 10^{-2} moles per liter. All right.

And so the pressure that results is a modest pressure, right? It's about a quarter of an atmosphere of excess pressure. 2.7×10^4 pascals. So let's just calculate, given this small additional pressure, how high is the liquid going to go, above the level of the solution. Anybody want to guess? You know, a centimeter. 10 centimeters. 100 centimeters? Give me an order of magnitude. Who guesses one centimeter? You got three

choices. One, 10 and 100, and you have to make one. Who guesses one centimeter? Who guesses 10 centimeters? Who guesses 100 centimeters? OK. Now, fortunately we don't do problem sets. Or other scientific problems by popular vote. Rather, we just solve them. Whenever possible.

So, what's the height of a column of solution? Well, that ρ , remember, that's equal to ρ times g times h , right? And we know ρ now. So h is ρ over ρ times g . So it's 2.7×10^4 pascal over ρ , that's 0.995 grams per centimeters cubed. Times 1 kilogram over 1000 grams times 100 centimeters per meter cubed. I'm just getting this into units that will give me, that will be in meters, times 9.8 meter per a second squared. And when I multiply all that out, it's 2.8 meters. So that puny little pressure ends up being meters high. Taller than any of us, by far. Sorry, it's up to there. And that's actually pretty typical. And what that means is, it actually makes osmotic pressure a very, very sensitive measurement method. If you want to determine a molecular weight, for example. Because it's so high, that means, and of course you can measure that height pretty accurately. Certainly within about a millimeter. And so even relatively modest concentrations, it's still easy to see the effect and measure the impact. And use that to calculate the properties of the solute.

OK, any questions? Alright. Now we're going to start something new. So, so far what you've seen, and I hope come to appreciate, is the whole structure of thermodynamics is built on empirical macroscopic observation and deduction. We observe things. We formulate these broad thermodynamic laws. All macroscopic. It doesn't depend on any microscopic model at all. And in fact much of it was formulated before there was a good microscopic model of matter. Before the atomic theory of matter, and molecules, and so forth was well worked out. And it's incredible how powerful that whole formalism is, right? In some sense, although it may seem like neglecting what we now know to be an important part of nature, actually part of its power is its empiricism. Right? There are these very small number of fundamental laws from which everything else follows. And at this point, of course, there's just enormous confidence in that small number of laws. Because of their being verified in so many context. But, at the same time, we do know about the atomic theory of matter. We know there are atoms and molecules. So it ought to be possible, at least in principle, to start from a purely microscopic approach to nature. And just based on figuring out the microscopic properties. And then saying, well, OK, my macroscopic stuff is just a collection of those microscopic entities. I should also be able to figure out macroscopic thermodynamics. I should be able to start from my microscopic picture and get to macroscopic thermodynamic results.

And in fact, that is possible. And the theoretical formulation for it is what's called statistical mechanics. Called that because, of course, you won't be surprised to learn that it's going to require a statistical treatment. We're going to be dealing with moles of material. But now we're going to be trying to think about their microscopic properties. And so we'll be dealing not with n number of moles, but with 10^{24} number of molecules or atoms. And you know we won't be able to keep track of every one of their individual states, all the time. We may be able to do a

terrific calculation of the quantum mechanics or the classical mechanics that describe the states that they could be in. But there's no way in practice that we're either going to experimentally or theoretically keep track of all of that.

So we're going to, at some point, have to introduce statistics. To take what we know about the microscopic properties and try to go from there to the macroscopic results. The thermodynamics that we've seen so far. And that's what statistical mechanics is all about. So let's try to introduce a little bit of statistical mechanics. Where we're going to go from what we know about microscopic properties all the way to macroscopic thermodynamics. That's our objective.

So let's start by just trying to calculate energies of individual molecules, or individual particles. And what's the probability that some molecule, one of the oxygen molecules somewhere in this room, is in a certain energy state. Right? And what our strategy is going to be, is to determine what are the probabilities that molecules are in certain states with certain energies. And if we can determine that for all the possible states, then we can average over those. So without keeping track of every individual molecule, we could then calculate, an average energy, which is what you would measure thermodynamically when you look at the whole collection and measure what we call u , right? Of course, we're really averaging over disparate energies of lots of different atoms or molecules. They don't all have the same molecular energy. And we don't try to measure their individual energies. So that's what we'd like to calculate. And so we'd like to be able to know what are all these probabilities of different energy states. And then from that, statistically averaging, what are the macroscopic average energies. And other macroscopic quantities. So, let's start there.

Probability that a molecule, must be specific as possible is in state i with energy E_i . And right now we're not even going to specify the nature of the state. We could be worrying about translational energy. What state it is, how fast it's whizzing around the room. We could worry about its vibrational or rotational energy, or electronic state. For now, let's not even specify it. Let's just say the only thing we're really specifying about the state is it has some energy that we presumably know.

Now, what we'd like is to know a functional form for the probability. And we don't know one a priori. But let's just think about two molecules. And the probability that one of them has energy E_i , and another one has energy E_j . It also, let me say, would be sufficient to think about even one molecule and say, we'll let's think about independent parts of its energy. Let's say, translational energy in this direction. And in an orthogonal direction. What matters is, we're thinking about two independent energies. Could be separate molecules that aren't interacting. Or independent degrees of freedom on the same molecule. But let's just consider it.

So a molecule is in state i , and that another molecule is in state j , with energy E_j . So what we're going to calculate is the joint probability that the two molecules are in those states. So we want a probability. This one, we'll call

$P_i(E_i)$. This would be P_j of E_j . And this probability together will be $P_{ij}(E_i + E_j)$. I could just write E_i comma E_j , but the energy's add. And that's crucial. Now, since the molecules are completely independent, this should just be the product of the separate probabilities. Two completely independent events, or things whose probability we want to determine together. So, that suggests a really simple functional form. Because what we're saying here is we want a function of $a + b$, which is equal to the same sort of function of a times the same sort of function of b . In other words, when we multiply these functions, we get a function of the sum of those arguments. What functional form does that? Exponential. Sure. So in other words, e to the $a + b$ power equals e to the a times e to the b . So it suggests that there's an exponential probability. Suggests it.

Now, there's also some physical insight that we can apply to the problem. First of all, we expect that states with higher energy in general are going to be less probable than states with lower energy. If it's cold, not much thermal energy around, and you say, how much energy does a particle in an equilibrium collection of particles have, at thermal equilibrium, if there's not much thermal energy around, you don't expect it to be very likely that a particle has a huge amount of energy. That one individual particle has it. Could happen, right? From some random collisions that just happened to bulk up the energy of that one particle. But on a balance it's not going to be very probable. Lower energies will be more probable. And the other thing is, even in thinking about something like that, temperature comes in. Immediately. Right? If you raise the temperature a lot, surely we have to expect many more molecules will start to become energized. Will have more energy. So, lower probability for higher energy. But the probability of higher energy should go up if the temperature goes up. In other words, the ratio of energy to temperature should be involved here. So, P of E_i should go down as E_i goes up. And should depend on E_i over T . Right?

So, we can start with this functional form. And by the way, there's no reason we can't have, in general we would have some constants here, right? We haven't changed anything to do that. And now let's just use the little bit of physical intuition that we're thinking about here, to refine this just a little bit. Let's write, we expect that P_i of E_i is some exponential to the minus $C E_i$ over T , where C is some constant greater than zero. So this will have the property, then, that as energy goes up, the probability goes down. But it's scaled by temperature. If I raise the temperature, that makes the higher energy states get more and more likely.

Well, this is our functional form for probability of a molecule being in a state with energy E_i . And the only difference between this and what's written conventionally is the way the constant is labeled. So really what we have is P_i of E_i is proportional to e to the minus E_i over $k_B T$, or just $k T$, where k_B is called the Boltzmann constant. And it's just equal to R over Avogadro's number. It's the gas constant per molecule, rather than per mole.

One way to try to rationalize this is, you've probably seen, you've all seen Arrhenius kinetics, right? Arrhenius rate laws? If you remember what that looks like, you get this rate constant. Arrhenius rate constant is some constant A

times e to the minus E_a over RT . Remember that? You've all seen that before? And so what's happening here, E_a is what is an activation energy. Remember, the idea is that you've got reactants and products, and there's some barrier you've got to get over. The activation energy, before you can have reaction. So the rate depends on surmounting that activation barrier.

This is really coming from the same idea as this, which is the probability of one of the molecules having this much energy. Depends on e to the minus energy over RT . This is per mole, so this is per mole. The exact same relation. So the idea in the context of kinetics is that the rate depends on how many molecules can get up here. And what the probability is of their getting enough energy to go over the barrier. And that energy, that probability, is given by this expression. So in this form you've seen this kind of dependence before in a very explicit way.

Let me just take it one small but important step farther. Which is that proportionality constant. Couldn't we do better? Couldn't we say, couldn't we figure out what exactly it is, not just what it's proportional to? Well, let's try. So as we've written it, we've got P_i of E_i is proportional to e to the minus E_i over kT . We could write that as equals A e to the minus E_i over kT . This is just some proportionality constant. But we can determine what that constant is, because we know that if we sum over all the possible states, the molecule has to be in some state, right? So that sum of all those probabilities has to equal one. So sum over i of P_i of E_i , is equal to one. Because the molecule must be in one of the states. OK, now this is our expression for all those p_i 's. So let's just write that way. a is just a constant times the sum over all the i 's of e to the minus E_i over kT . And so there it is. a is equal to one over the sum over i , e to the minus E_i over kT . So now we can rewrite P_i of E_i is equal to e to the minus E_i over kT over, the sum over i , e to the minus E_i over kT .

Now, just so you're not confused, this i matters. That's the i we're talking about. What's the probability of it being in some particular state i with energy e_i . This i is just a dummy variable. So just to be explicit, we could write this just as well as e to the minus E_i over kT times the sum over j , e to the minus E_j over kT . It could be confusing, but this is usually the way that it's written. Even though this dummy variable here has nothing to do with the particular choice that's made here.

So now we know. And we can calculate what the probability is for a molecule to be in any particular state if we know the energy of that state. Now, notice we need to know all the energies of all the other states, too. Because this thing, in other words, this a , right, depends on summing over all of those. But if we know those, then we can do it. We can do the whole calculation. And for a good number of cases, we realistically do know it all. So from the proportionality relationship alone, of course, we can tell the ratio of chances that you're in one state or another. Even without this a constant, we already could've said well, if what's the ratio of, the probability of being in state i with energy E_i to state j with energy E_j . Well, it's just e to the minus E_i over kT over e to the minus E_j over kT . It's e to the minus E_i minus E_j over kT . But apart from just being able to get the ratio, we can get the absolute number.

What's the absolute probability that a molecule is in any particular state of a given energy.

Next time what you'll see is what can be calculated based on the results that you've seen, just the results you've seen so far. From these quantities alone, it turns out you'll be able to calculate every single macroscopic thermodynamic quantity. So you'll see some of that next time.