

Lecture 15B • 11/01/11

Le Châtelier's principle; Q versus K

Let's go ahead and dive into this topic of equilibrium. At first, my example's going to be a physical example; I'm going to use vapor pressure. For vapor pressure being my example, the forward process will be evaporation; the reverse process will be condensation. First let me ask, why is it that vapor pressure forms in the first place? What is vapor pressure? Boiling is when the vapor pressure equals the surrounding atmospheric pressure, but what causes vapor pressure? This is that energy distribution graph; at any temperature you're going to have a distribution of molecular energies. Some of those molecules, at whatever temperature that you're at, are going to have enough energy that they'd be able to evaporate. The example that I had given is if you put some water out on the desktop here, that, over time, even below the boiling point, the water was going to evaporate; that's just because there's a fraction of molecules that have that energy. Conversely, though, if you have vapor – there is water vapor in the air here – but if that vapor builds up to a certain concentration, then we get dew or other forms of condensation. There's always going to be some subset of molecules that are moving slowly enough that that vapor's going to be able to recondense. So we've got two opposite processes going on; we've got evaporation and we've got condensation.

Let's talk about the interplay between the two, because the same could then be said about a chemical reaction. Let's imagine that we have a container, and at some time $t = 0$, water is just magically transported into that container. Let's say that the container has absolutely nothing in it to begin with, so in other words, it's at a vacuum. At time equals zero, water is suddenly introduced into a completely empty chamber; it's at first under vacuum. If we only have water liquid around, and if truly was at a vacuum, there's no condensation going to be possible, because there's no water vapor, so the only process that's going to happen, in fact, is the forward process of evaporation. I could say that the concentration of products is zero, and therefore the reverse rate is equal to zero; we have only evaporation. Let's say that we wait just a minute even. That evaporation's going to happen because we have a fraction of molecules able to evaporate, but as soon as they start to evaporate, that means we're going to have a little bit of vapor that's around, and a few of those molecules will be moving slowly enough that they're able to condense, so we're start to have a reverse rate of reaction. I'll say just after time $t = 0$, we could say there is a reverse rate, but the forward rate is greater.

But, the more evaporation that keeps occurring, at some point you build up enough vapor that that condensation rate is going to increase. At some point, those two processes balance each other out. The way I'll say it here is $t =$ equilibrium, that's the point at which the forward and reverse rates will equal each other. So there's no more observable change in either the amount of liquid or the amount of vapor, because the rate of conversion between the two is exactly equal. After we've achieved equilibrium, let's say we pushed equilibrium. Let's say that, just as in the beginning we were able to magically transport some liquid water into there, let's say we transport some vapor into there. The rate of evaporation of the liquid won't have changed, because you haven't changed the amount of liquid that's in there; but if you suddenly added more vapor, that means the rate of condensation is going to be greater than that of evaporation. So that means that, overall, some of the vapor's going to have to condense, until you get rid of enough vapor where the rate of condensation slows down again enough, where it again balances the rate of evaporation. In other words, if you push the equilibrium by putting in too much vapor, it responds by having some of that vapor condense. That's exactly type of thing that's explained by Le Châtelier's principle. We'll say past equilibrium, the products (which in this case would be the vapor), that concentration's too great, and so the reverse process, condensation, is going to be a greater rate than the forward process.

Summarizing this whole example again, starting out with an empty chamber, only liquid water, since there's only liquid water, only evaporation's possible. So you build up some concentration of vapor, but as you do, some of that vapor starts to be able to condense. That process keeps going on until the two rates equal each other, at least if we're still in a closed container, but if you push that equilibrium stress that equilibrium, add extra vapor, then that vapor's going to end up recondensing.

A similar thing is going to happen now that we start talking about chemical reactions. Let's get to Le Châtelier's principle itself. Le Châtelier's principle is, when a system at equilibrium is stressed, which means it's changed somehow, the system will respond in a way that relieves that stress. To quantify it, we have this value called Q, which stands for the reaction quotient – quantify the position of equilibrium, I should say. Q is an expression that has exactly the same mathematical form as K. Q is essentially products versus reactants, except that you can measure the value of Q at any point in a chemical reaction, where K you only calculate using equilibrium concentrations. Since K is only calculated at equilibrium, Q will tell you do you have too many products or do you have too many reactants; that's how we use Q. Q is mathematically identical to K, except K is only calculated using equilibrium concentrations, while Q can be calculated at any point in a reaction. Q would be the amount of products versus the amount of reactants.

Let's say how we could use Q and K together in an example of Le Châtelier's principle. Let's start with an exergonic reaction. Since I'm talking about equilibrium, I'm operating in free energy. We saw a free energy graph, this reaction progress diagram. Remember the difference between this and a reaction coordinate diagram is a reaction coordinate diagram shows just how one set of molecules is changing as energy as it goes from being reactants to products.

A reaction progress, this is the sum of all of the different chemical reactions going on in a solution, so it's focusing on the average energy. Exergonic means we have a negative change in free energy. There's a term 'spontaneous' that we use to describe processes that have negative free energy changes, because they're the ones that are gonna happen just as is; if you have a release of energy, it's a favorable reaction, it wants to move forward, and that's called spontaneous. On this graph, then, the lowest-energy point would, of course, correspond to equilibrium. Let's say that we have a Q such as this, to the left of equilibrium – left will mean to have too many reactants. If you have a value of products over reactants, that means that Q is going to be too small. In cases where Q is less than K, not enough products have formed; or, you could say that there's too many reactants around. Conversely, what if I put Q to the right of K, which means you've gone closer to products; that's what would make Q greater than K, products over reactants. So if Q is greater than K, not enough reactants are present; or, you could say too many products have formed.

To give you a chemical example and see how we can use all of this information, I'm going to show you a reaction that we will learn next quarter. If you take a compound like a ketone, something that has a carbonyl in it, it is possible to react that ketone in the presence of acid and water. What can happen first is that the oxygen of the carbonyl get protonated; we can then show the carbonyl bond opening in response to relieve that positive charge; technically, that is just a resonance structure. That ends up giving us a carbocation, which can be attacked by water; like other processes we've seen involving water, since water begins as a neutral compound, once it forms a bond, it's going to have a positively-charged oxygen. As a response, one of the hydrogens comes off, and we formed this product. It could be classified as a geminal diol – two alcohols located on the same carbon, but even more frequently, it's called a hydrate, because you've reacted water, you've pushed water onto this ketone. Notice that every step of this process, I've written reversible arrows, because this really is a reversible reaction. In fact, this is not a favorable reaction; this would be an example of an endergonic reaction, one that would rather stay as reactants. Let me draw a reaction progress diagram for an endergonic process. As we go from reactants to products, we'd be going uphill, so $\Delta G > 0$; ΔG would be my energy scale. Equilibrium lies to the left on the energy curve. The reaction that I've shown you is reversible and non-spontaneous.

What if we wanted to force this reaction to happen, though? What could we do to make us get products? Remove the products once they form; that is one approach that is often taken, and it's one that we're going to take in one of our experiments. As the products forms, we keep on removing it. Why does that help things? If products are removed, as the system tries to reach equilibrium, we will suddenly have made Q less than K, because Q is products over reactants; if you remove products, then you're dropping the value of Q. When Q is less than K, we say that the reaction's going to back forward to make more products, to make Q get large enough again that it equals K. Why does that occur? Because if you removed products, and if it was at equilibrium when you disturbed it, that means the rate of the forward reaction will be greater than the reverse reaction, because products were removed, but not reactants. On a graph, that will again put us to the left. Either direction from equilibrium is uphill in potential energy, so by removing products, we've effectively put the system at higher potential energy, and it's going to make more products to get back down to that lowest-energy equilibrium state. The reaction will move forward to make more products to reestablish equilibrium – to again make Q equal K, to again make the rates of the forward and reverse reactions equal to each, and again to reach the lowest, non-changing energy point of this reaction.

To put it in practical terms, if products are continuously removed as they form, then the reaction can be forced to completion – completion meaning 100% products, even though the reaction is non-spontaneous, even though it's uphill in energy. In some cases, we could use similar logic to do it the backwards way, because instead of removing products to tease more products to form, sometimes you can stuff a reaction with a whole excess of reactants. If you've got a huge concentration of reactants, that again makes Q less than K, and that encourages the reaction to move forward. But that might not always be as practical, either because of cost, if you reagents are really costly, or if you don't get good mixing of the reagents. For example, with some ketones, they wouldn't be water-soluble, so no matter how much water you threw at it, you might have a mixing problem. Depending on a particular reaction, if you're trying to force it forward, you might add reactants, or, as I've shown in this example, we would remove products.

We're talking about moving equilibrium around, but really we're not talking about moving the equilibrium point around; we're talking about messing with product or reactant concentration. You can mess with those concentrations, but then they're going to try to reestablish equilibrium. What I'm trying to say then is any change in concentration never affects K; it only affects Q. There's only one thing, really, that can affect K – temperature, because temperature will change the rates of both the forward and reverse reactions. That will change the balance, therefore, between reactants and products, so you're going to change your equilibrium constant. K, equilibrium constant, is only affected by temperature, not on concentrations. You can mess with concentrations, it will mess with the position of that reaction relative to equilibrium, but it won't move the equilibrium point itself.

Delocalization

Why do we need to discuss delocalization? It has a big structural impact on the way that molecules react. When we write a mechanism that involves resonance structures, there isn't really a reaction going on, because resonance structures are just something that we play with on paper. Let me go back to this allyl cation example. Let's regenerate the SMOG for this.

Each of the carbons in this case is sp^2 -hybridized, which, fortunately in this case, is pretty easy to see – no matter how we draw it, each carbon only has three bonds explicitly coming from it. Since each carbon is sp^2 -hybridized, that means in this system, each carbon also has a p orbital. Should the left two or the right two of these p orbitals be connected? The answer is yes, meaning that, in fact, all three p orbitals are connected. That means that neither one of these resonance structures, individually, properly conveyed the structure of the molecule. The only reason these exist, though, these resonance structures, is because that's the only writing system that we've got, if you want to count the number of electrons. There are these ways of kinda showing resonance where you can use dotted lines, and you'll see textbooks use that sometimes, but I'm extraordinarily hesitant to teach that, because with a dotted line, you can't count electrons. The whole fault with resonance structures is that they come from Lewis dot structures that assume a bond only occurs between two atoms. Resonance structures only exist because the system used to write structures – which comes from Lewis dot structures – is unable to unambiguously represent bonds that cover more than two atoms. That also means that no individual resonance structure correctly shows the true structure of the molecule. Therefore, the only clear picture of a molecule's true structure could be obtained by averaging all of these different resonance structures.

There really is only one main rule for writing resonance structures: only lone pairs and pi bonds are able to move. Single bonds can never be formed or broken. Why is this true? Because single bonds are sigma bonds, in which orbital overlap is only occurring between two atoms. The whole point of delocalization is that you've got the interaction of orbitals from multiple atoms involved. Lone pairs are able to delocalize, and systems of p orbitals can be extended, which means we can move double bonds around by choosing which of the pairs of p orbitals we connect. Let's see a sneaky example, though, of just pushing lone pairs around.

We've seen the allyl cation, but how about now the allyl anion? If the rule is correct, I should be able to push the lone pair over to where the double bond is, which, because carbon already has four bonds, would cause the double bond to break, and I'd end up with a new resonance structure. On this second structure I just wrote, where the negative charge is, what is the hybridization on that carbon? If you're following the rules, you'd say sp^3 , right? Because, on that position, you have a lone pair, there's still two more hydrogens that would be left over since it's a methyl group, and then we have the carbon-carbon bond, so that gives us four pairs of electrons. But that's not the structure of this molecule. That carbon, regardless of the fact that I've written a negative charge there, is sp^2 hybridized. Why? Because if it weren't, you wouldn't have this delocalization. This is exactly [an example of] a resonance structure [not giving] you a true picture of the structure of the molecule. In order to correctly draw a SMOG for this, you'd have to know that it is undergoing resonance. It undergoes resonances only because we can't write the real structure; it's delocalized. This carbon is sp^2 -hybridized. It appears that it is sp^3 -hybridized, in this resonance structure, but that's because resonance structures fail to show delocalized bonding. Delocalized – local means right here, delocal, not right here; delocalized means to make not right here, in other words to spread it around. I won't draw a SMOG for this because it's identical to the allyl cation.

What would the molecular orbitals look like in this system? When we were talking about pi bonds, just two p orbitals interacting with each other, how many molecular orbitals exist when we have two p orbitals interacting with each other? In other words, when we made the pi bond, what else could exist in theory, but is not filled if you really do have a bond? The pi antibond. Because you have two nuclei interacting with these electrons, you end up with two molecular orbitals that we can estimate the shapes of by using p orbitals. It's not really that the p orbitals are adding and subtracting, we just do that on paper to make picture of what the molecular orbitals really look like. One of those molecular orbitals looks like what you'd get if you add p orbitals with matching sign, and that was the bond, and then you'd get another molecular orbital if you took the two p orbitals and added them oppositely to each other, and that was the antibond.

But now we have three orbitals interacting. Imagine this: think of the two ends of the molecules as endpoints for a larger wave that covers all of those atoms. That's going to be much like this phone cord demonstration that I did. Part of the reason I did that demonstration is exactly for this kind of situation. If we end up with three nuclei, with three p orbitals interacting, we're going to end up with three molecular orbitals, just like when we went from the bonding to the antibonding orbitals of a pi bond, and that had zero and one nodes, you go up one node each time you go up an energy level, then in this system where we have three orbitals interacting, we're going to have three molecular orbitals that have zero, one, and two nodes, respectively. In an idealized way, these possibilities can be shown like this. If we wanted to do the same thing that we did with pi bond and antibond, which is to come up with a molecular orbital representation of p orbitals that could just overlay p orbitals on this wave representation. In order to get a lowest energy wavefunction, where the sign of that wavefunction is the same going across the molecule, I'm going to use three different p orbitals that all have the same phase. Why was it that a bonding orbital was lower in energy? Because that wavefunction for that molecular orbital put electrons between nuclei. Notice that this is doing exactly the same thing, so this ends up being a bonding orbital. Now the opposite extreme. If between every nucleus you end up having a node – and the maximum number of nodes here is two, which is the number of atoms minus 1, since you can't be between the two outside atoms. $n - 1$ gives us that maximum number of nodes of two, which we could get if every other p orbital had opposite sign. In a simple two-atom system, this would be an antibond. It still would be an antibond in this case. Why? Because if you're preventing electron from being between two nuclei, then that means the nuclei are exposed, so you're going to be higher in energy.

But what about this middle case? Because if I drew in the p orbitals here, it's going to turn out that one of the nodes passes right through that orbital. What that means is that when we estimate that molecular orbital, we don't even use that p orbital. An initial response might be: if we're not using that orbital, how can we ignore it? Remember that that orbital does not even exist. This is our way, on paper, of adding and subtracting orbitals to make things look like the real molecular orbitals that do exist. It's ok to write it this way. But they're so far apart now, these two different [p] orbitals, that there's really no interaction, either favorably or unfavorably, they're just too far apart to interact. If there's no interaction, that means it doesn't change energy of what it would have been if you had just left the atoms alone in the first place. So, we've seen antibonds, we've seen bonds, but this is our first example of a non-bonding orbital. Non-bonding would mean zero energy, relative to the uninteracted atoms.

If we see that we have a double bond and a lone pair, that's a total of four electrons, we have to put those electrons somewhere – two of them will go into the bonding orbital, which is exactly what happens with the allyl cation. But this allyl anion can exist because if you put two more electrons into it, they don't add or subtract to or from the energy of the system. They don't effect it because it's a non-bonding orbital, so no problem with the electrons being there, and that's why they can delocalize. This is just the molecular orbital explanation to back up what's going on with resonance. If you had four atoms, then you would have four molecular orbitals: two would be bonding, two would be antibond. In generalization, if you have a linear system of p orbitals, if it's even, then half are bonding, half are antibonding, and if you have an odd number, then one orbital in the middle is non-bonding. This is just another energy level, but what you're not used to is energy levels that don't, in a way, have energy in them. They're not higher or lower in energy than the individual atoms. A non-bonding orbital is the same energy as the atoms would be if they had never undergone bonding.

Better resonance structures generally have full octets when possible, they lack charge separation, and have charges that match electronegativity. Two examples. Let's take the molecule acetone. That's a molecule in which all atoms have octets (or all hydrogens have duets). I could technically, though, write a resonance structure, where I have the carbonyl bond opening; but, this is not going to be a very good resonance structure, because I took something that has octets and I went to something that didn't. On top of that, I've moved a negative charge away from a positive charge, which is resonance is happening, why would the charges move away from each other? So this is unfavorable due to charge separation. But, it's not a 100% wrong structure either. The oxygen is more negative, the carbon is more positive, so put these two structure together, and it gives you a little better picture of what a carbonyl does look like. But I could make an even worse situation if I take that same double bond and open it the opposite direction. Technically possible, but it would have given us this odd situation, where oxygen would be positively charged and carbon would be negatively charged. This is even worse, since oxygen is now positively charged. It's not favorable, even less favorable than moving charge apart.

forward – evaporation; reverse – condensation

@ t=0, water is introduced into a completely empty chamber (under vacuum)

Le Châtelier's principle – When a system @ equilibrium is stressed (changed), the system will respond in a way that relieves that stress.

Q = reaction quotient

Q is mathematically identical to K, except K is only calculated using equilibrium concentrations, while Q can be calculated at any point in a rxn.

Q = products / reactants.

If $Q < K$, then not enough products have formed, and/or there are too many reactants present. If $Q > K$, not enough reactants are present and/or too many products have formed.

The rxn above is reversible and non-spontaneous. If products are removed as the reaction tries to reach equilibrium:

- 1) $Q < K$, since Q is [products] just decreased
- 2) $R_f > R_r$, since products were removed (but not reactants)
- 3) Potential E of the system is effectively increased.

Therefore rxn will move forward to make more products to re-establish equilibrium.

– If products are continuously removed as they form, the rxn can be forced to completion, even though it is non-spontaneous

K is only affected by temperature, not [].

– Resonance structures only exist because the system used to write structures (L. D. S.) is unable to unambiguously represent bonds that cover more than two atoms.

– No single resonance structure correctly reflects the true structure of a molecule.

– The true structure of a molecule can be estimated by averaging all possible resonance structures together.

Rules for writing resonance structures:

– Only lone pairs and pi bonds are able to move; single bonds can never be formed or broken.

This carbon is sp^2 -hybridized. It appears that it is sp^3 hybridized in this resonance structure, but that is because resonance structures fail to show delocalized bonding.

max # of nodes in a pi system is = # of atoms – 1

A non-bonding orbital is the same in energy as the atoms would be if they had never undergone bonding.

– Better resonance structures generally have full octets (where possible), lack of charge separation, and have charges that match electronegativity.

Structures

Identical to those from lecture 15A (10/31/11)