[SQUEAKING]

[RUSTLING]

[CLICKING]

IAN BALL: So today, we're going to talk about signaling. And what we did last week was kind of a simple example of signaling, and today, we're going to do a more involved application of signaling.

So what do we mean by signaling? What's the key idea here? Well, we're going to look at games where we have two players. We have a sender and a receiver. And the sender has some private information. And then the sender, knowing their private information, is going to make some choice about an actions. They're going to then choose an action.

And we often call this private information, maybe the sender's type. This is what the sender knows that the receiver doesn't know. And then the receiver is going to observe the action that the sender chooses, but they're not going to directly observe the sender's private information. So the receiver observes the action, but not type.

So they see what action the sender takes, but they can't observe the private information that the sender has. And then based upon that, they're also going to choose an action.

So in the examples that we're going to look at, the receiver, when they're choosing their action, cares a lot about the sender's private information. And the receiver is going to try to learn or infer something about the sender's private information based upon the action that the sender chooses.

And conversely, the sender cares a lot about what the receiver believes about them, and the sender is going to try to convey information about their private information through their choice of action. Now that's a bit abstract, so let's think of some examples.

I don't know if any of you come across-- heard of the term signaling used colloquially and heard of examples where this term is applied. So a few examples. I guess a classic one would be luxury goods. So here, the story would be that the private information that the sender has, their type is maybe how rich they are, how much money they have. The action they're choosing is what brand of bag to buy.

The receiver, then, doesn't directly know how rich the sender is, but they do observe what brand of item they choose. And the receiver is going to make an inference about how rich the person is based upon how luxurious the item is that they hold. And the sender, wanting to be perceived as rich, will then choose to buy something luxurious.

Now, why does this work? Well, your concern might be, well, look, if everyone wants to be perceived as rich, why doesn't everyone just buy the luxury item? And the issue is that, well, they're expensive, and only some people can afford them.

So a crucial aspect for these signals to be credible is that the action people are taking is generally going to have a different cost depending on their private information. If everyone was able to buy luxury items because they were so cheap, that wouldn't credibly convey that I was wealthy. But for a very expensive item, because only sufficiently wealthy people can buy it, that's going to convey information.

Another classic story from the animal world, I guess, is a peacock's feathers. Any ideas about what's being conveyed here and what's the role of the peacock's feathers? Why do peacocks grow very, very ostentatious large feathers? Well, it's not because people really care about the feathers. The idea would be that this is a signal of health. So if we map this into our example, the sender, one peacock privately has private information about how healthy they are. They're trying to attract maybe a mate, and the mate cares about how healthy the peacock is.

If the peacock grows these enormous feathers, that's extremely costly. It takes a lot of energy. You have to-- to grow this. So by growing very, very large, ostentatious feathers, they're conveying to the receiver that I must be very, very healthy. And the receiver doesn't care about the feathers, they just care about underlying health, and the feathers provide a signal of that.

One classic example, which we'll go through today, would be college. So this is a classic signaling application. How does college fit into this sender/receiver model, or how could it? Yeah?

AUDIENCE: --show a job recruiter that you are aspiring--

IAN BALL: Right, exactly. So here, your private information might be how smart you are, how diligent you are, how hardworking you are, we could have various kind of positive attributes. You're choosing whether to go to some difficult university and get some demanding degree. The receiver might directly care about what you learn in your degree, but they might not.

It might be that your degree is actually not that relevant for what you're doing on the job, but they care about your degree because it's a signal of how smart you are. The idea would be that some people are able to get this degree and pass, and other people are not, and therefore, your ability to get this degree and take this action reveals information about your ability.

So let's dig in a little bit deeper into this college example because that's going to be kind of the main application we're going to look at today, and this is often goes under the name of job market signaling.

So this was introduced in this famous paper by Michael Spence in 1973-- I think it was actually his PhD dissertation, and he ultimately won the Nobel Prize for this idea, which we might go over today and you might say it's pretty simple. How did you get the Nobel Prize for that? But I think things after the fact seem simple, but coming up with them is often quite difficult.

So let's just observe some facts about the world, about college education. There's something called the college wage premium. And this refers to the fact that at least all over the world, but let's focus in the United States, people who go to college or complete a college degree tend to earn more than people who don't.

And this is something labor economists study a lot and try to give different estimates. You'll see a lot of estimates floating around depending on how they measure it or which cohort or which kind of person they're looking at, but maybe a rough estimate in the US might be something like maybe 500K to a million is maybe kind of the middle-of-the-road estimates, and this is in lifetime earnings.

So what do I mean by this? I mean, if you just look at people in the US, you look at those who have college degrees, and you compute the average earnings of those people who have college degrees, and then you look at people who don't have college degrees, and you look at the average earnings of those people over their lifetime, and you compare it, you see that people who go to college, on average, earn maybe 500K to $1 million more over their lifetime. As I said, there's a few different estimates about this.

Now the question is why. Any thoughts about why this is? Why do people go to college earn more? You're here at college. Maybe you partly want to earn more by coming to college. Any thoughts about why this is? So maybe-- we often separate this into two different stories.

So one is what we might call a causal story. I'll just say causal, we don't have to put the word story. So is it causal? Is it causal, what I mean is, let's look at someone who doesn't go to college. If they chose to instead go to college, would their earnings go up? That would mean it's causal.

Or maybe it's not causal. Maybe someone who doesn't go to college, if they did go to college, wouldn't actually increase their earnings. Now I think a lot of people would look at this fact and just immediately say, oh, it must be causal. If going to college-- if people who go to college earn more, then it must be that if I choose to go to college, that's going to increase my earnings, and this is kind of a classic fallacy-- or potential fallacy when people haven't thought so carefully about this data, and that's possible. It could be causal.

But one explanation is it's not causal. How could that be? How could it be that people who go to college earn more money, yet going to college wouldn't actually increase my earning? How could that story make sense?

Well, the story here would be that, well, the people who go to college are very different from the people who don't go to college. You have two pools of people. You have people who go to college and people who don't. People who go to college are just smarter, more ambitious on average or have some other trait that is hard to measure. They do earn more, but it's not because they go to college, they just happen to be going to college.

And this is, I guess, the classic story of this. This is a problem with all these health studies where they say people who eat five almonds a day live five years longer. Eating five almonds a day doesn't make you live five years longer, but people who eat almonds are different in a lot of different ways. They eat healthily in other ways, they run, they exercise, they do all these other things.

So it could be that though there's a big earnings gap, it's not because you go to college, it's just a spurious correlation. So that's one potential explanation.

Another would be that it is causal. That if I, as an individual, choose to go to college, that's going to increase my expected earning. And I think the best evidence suggests that both-- there's an element of both of these at play. There's-- maybe this we might call selection, the people who go to college are a selected group. They're different on average than people who don't go to college.

And I frame this as causal versus not causal. I think the best evidence suggests it's a bit of both. That when we look at this earning gap between people who go to college and people who don't, partly it's attributable to selection. Those people just have different skills, and partly it's a causal effect of college.

But I want to drill down within causal, and I think there's actually two different mechanisms that could be at work here. So let's focus on the causal component. Let's just suppose that it is true that if I go to college, my earnings go up. Why might that be the case? Why my going to college cause me to earn more? Yeah?

AUDIENCE: You're able to learn skills that employers might want.

IAN BALL: Great. So. Certainly the story that people who run colleges would like to focus on is this story, that it's about learning skills. Sometimes this is called human capital formation, that you go to school, you learn how to code, you learn how to communicate with people, you learn knowledge, information about the world, you mature, various things, and these make you a more valuable employee. They allow you to be more productive. So this is sometimes-- we talk about this might be about productivity.

But another possibility, which we're going to focus on today-- learning, is signaling. So the other possibility could be that it is causal. Going to college makes me earn more, but it's not because I actually learn anything valuable in college. It's simply because going to college proves to the employer that I'm smart, and therefore, the employer pays me more. It's still causal.

If I didn't go to college, I wouldn't earn as much. So it is true that going to college changes my expected earnings, but it changes it not because I've actually acquired valuable skills, but because I've demonstrated that I could do a very difficult task, and that makes employers want to hire me.

So one way of separating this, you could imagine that a school-- in the extreme case, you could imagine a school that teaches skills that really have no value, but are just really, really hard to learn. And you might have-- people who graduate from that school might actually still do well in the workforce.

And maybe an example of that-- I mean, I think athletics are a valuable pursuit, but you might think that people who succeed in athletics are often hired by firms to work in finance, say. It's not that being good at running makes you-- or being good at weightlifting makes you better at finance, but it shows that you've done something very hard that reveals something maybe about your diligence, your work ethic, and that might be a reason why firms want to hire you.

So in reality, I think both are probably present, and I think the best empirical evidence at the moment suggests both of these mechanisms are at play, but to highlight the signaling aspect, we're going to focus all in on signaling, and we're going to consider a model of education where education conveys no valuable skills. I want to say up front, this is not my view of education, but just to highlight this-- I think we all understand how learning skills increases your wage. I think this is the novel element of Spence's model, and this is what we want to look at.

And I will say, there are people who argue that this is almost all of education. It's a debate, but there are some people who say this is really all that's what's going on. So let's go through this.

So one-- I'll just maybe make one observation, one way that people try to argue it's all about signaling, one kind of interesting empirical test is they look at people who go to college for three and a half years and then drop out. Now if it's really about skills, you'd think someone who goes for three and a half years has acquired almost all of the skills as someone who goes for four years.

But if it's about signaling, you might think, well, the fact that they dropped out reveals a lot about who they are as a person, and indeed, people who drop out just before graduating earn a lot less. And that's one argument for why signaling is at play, but we can debate about these things.

OK. So now let's go through our formal model of signaling. So in our model we're going to have a sender and a receiver, just as I said before. And we're really going to simplify things. We're going to imagine the sender just has two types. They privately know they're type t, which is either l or h. They're either a low type or a high type.

And we're going to think of this-- maybe we can think of this a bit abstractly, but a low type is someone who is going to find education more costly, and a high type is going to find education less costly. We can interpret this in a lot of ways. It could be about your intelligence, it could be about your work ethic, it could be about the quality of the school you went to before. There's a lot of things that can go into this, but we're just going to quite abstractly call it a low type and a high type.

And what they're going to choose is they're going to choose a level of education. And in reality, education kind of usually comes in in units. Like, you can't-- we don't really think you get 862 days of education. We usually think you either get-- graduate high school or you graduate college, but to make things simple, we're going to make education continuous, e, and we're just going to say it's a number between 0 and 1.

So you can think of this-- just a stylized model where you can get any continuous amount of education. This could be something like the fraction of your life you spend being educated or something.

Now the receiver, we're going to think of this as maybe a firm or a market of firms. And they're going to observe e, but not t. When a firm chooses to hire you, they see what education you got. They don't directly observe this type t, whether you're a low or high type person. And then they're going to choose a wage.

In reality, I think a lot of individual firms, the wage is roughly set and they're choosing whether to hire you or not. There's maybe some flexibility over wages, but that's one way we can think of-- I think it's helpful to think of this kind of as a market. So it may be that each individual firm is making a binary choice about whether to hire you, but then you go to the firm that gives you the best offer. So in the end, the outcome of the firm-- of the labor market is that you get hired at some place for some wage, and depending on how well you do, that wage is determined by which firm hires you.

So now let's talk about payoffs. So we're going to assume the sender, their payoffs are like this. So the sender wants to get paid more. Their payoff is-- their utility is increasing their wage. They don't like going to school. Again, another maybe stylized assumption, some people enjoy going to school and that's why they go, but we're not going to have that in our model. We're going to assume that people just don't like school.

And notice, the more they get educated, the more costly it is. So they'd always rather have less school than more school, other things being equal. But now their type t appears in the denominator here-- and maybe I should be more precise. Here, we have b, a natural if we did it any other way, the low type is lower than the high type.

So by putting t in the denominator here, what I mean is if you're a high type, if t is higher, then the coefficient in front of this cost is lower, and therefore, the cost, or more importantly, the marginal cost of education is lower for the high type. Why might that be true? Well, maybe they don't have to study as hard to get into the school. Maybe once they get in, it's easier for them to get good grades and graduate, there's many stories we can tell, but, again, we're going to think of your type as representing how costly it is for you to get education.

And then the receiver, we're going to use a bit of a weird payoff specification. The receiver's utility is going to be like this. Now, this may at first seem a bit strange. What we're saying is the receiver's utility is the negative of the difference between the wage they actually charge and the type of the applicant of the sender. And really, the reason we make that assumption is the implication of this is that the optimal wage w is equal to the expectation of t.

And this is just an algebraic fact about the quadratic loss function, you can check this with algebra. So if I don't know what t is, with some probability, it's low; with some probability, it's high. The optimal wage for me to set is exactly my expectation of t. So if I know the person is low ability, it's optimal for me to set a wage of l. If I know they're a high type, h, it's optimal for me to set a wage of h. And if I think it's a 50/50, then I want to pay them exactly 1/2 l plus 1/2 h.

Now you might say, wait a second, that's not what firms do. They pay the lowest wage they can subject to getting the worker, but again, we can think of this as the outcome of competition. So if there's many firms competing, if I charge-- if I pay a worker less than their expected ability, then maybe I'm going to lose out to another firm and the other firm is going to hire them. If I pay them more than their expected ability, then I'm overpaying them. So we can think of this as a reduced form representation of a competitive labor market.

And I should say something about the prior here. Let's assume that h occurs with probability pi and l with 1 minus p. So in the population, fraction pi of people are of type h, are the high types, and fraction 1 minus pi of people are the low types.

OK, so now we're ready to analyze this game. And the first step is always to write down what a strategy is. Notice, this game is a little more complicated than what we looked at before because now, education is something continuous. Before in the class on Thursday, we looked at this problem where you chose, whether to submit an essay or not, and that was a binary choice.

I'll make one other observation. Last Thursday, we looked at something where your private type affected how much you wanted to get into a school, but not how costly it was to submit an essay.

Now we're doing an opposite assumption where your type doesn't affect how much you want to get paid. Everyone wants to get paid more, but your type only affects how costly it is for you to get educated, so it's kind of the inverse of what we did last week. And I think this is probably a pretty good model of the labor market where most people want more higher wages, that's not really what distinguishes them.

OK, so what is a strategy, or more precisely, what is an assessment? So remember, an assessment is a strategy profile together with a belief system. So we need to first say, let's look at the sender strategy. What is the strategy for the sender? Well, the sender chooses an action, a level of education after knowing or learning their private type t. So we need to specify what education the sender chooses when they're the low type and what education they choose when they're high type.

So we'll represent a sender strategy as el, comma, eh. It's two numbers. el says if you're a low type sender, this is how much education I'm going to get. eh says if I'm a high type sender, this is how much education I'm going to get. And yeah, I say sender, you can think of this as the student.

OK, what is the strategy for the receiver? This is a bit harder. Remember, a strategy is always a complete contingent plan. It says what action do you take as a function of what you observe. So in this case, the receiver observes your level of education, and they're going to choose what wage they set as a function of your education.

So formerly, the receiver is choosing a function w from 0, 1, maybe I'll say to 0, infinity. So what this says is it assigns to each level of education e a wage w of e that the receiver or firm will pay you if they observe that your education level is e. And they have to do that for every single education level 0, 1.

So that describes a strategy profile, but we're not done yet because an assessment also specifies beliefs. So now we also have to specify a belief system. Well, the sender already knows their type, so we don't need to say anything about the sender's type. We have to specify the receiver's beliefs about the sender's type. And once again, the receiver might learn about the sender's type through the level of education e.

So a belief system is going to say, what do I as the firm believe about you as a function of what education I see? If I see that you've got a lot of education, what do I believe about your ability? If I see that you didn't get very much education, what do I believe about your ability? So a belief system formerly-- maybe if we write it-- let's start formally, and then we'll be a little less formal.

So if we're formal about it, a belief system says for every level of education I see, I'm going to form a belief which is a probability distribution over your ability, either l or h. So the notation here just means this is the set of probability distributions over your ability being l and your ability being h. But we can specify a probability distribution here just by saying, what do I think is the probability that you're a high type? Because once we specify the probability you're high type, then the probability that you're low type is just 1 minus that.

So we're going to think of this-- maybe we'll write b h e-- so we'll just use this notation. And this says, if I see that your education is e, this is the probability that I think you're a high type. And that fully specifies my beliefs because the probability I think you're a low type is always 1 minus that. So maybe I'll just write that because b l given e, the probability that I assign to you being a low type given e, is just 1 minus b h e.

So technically, we have to specify both of these things, but we're only going to specify one of them because it's pinned down by the other. Great. So now, we have to fill all these things in. We have to specify the sender strategy, the receiver strategy, and then the belief system, which is a system of beliefs for the receiver. Any questions on this? OK.

Now what do we need to check? So now we're going to look at our equilibrium conditions. And remember, we're looking for the conditions, and these are for pbE. Our equilibrium notion is perfect Bayesian equilibrium. And we have three conditions or, depending on how we write it, one is sequential optimality or sequential rationality. You might see both terms used. Sequential optimality, sequential rationality, they mean the same thing. And then belief consistency.

And it's helpful to separate this into sequential optimality for the sender, and also for the receiver. So let's first just say what each of these things mean in words. Sequential optimality for the sender just means if I'm the sender, and I know my type t, my equilibrium level of education is optimal for me given what everything else that's happening in the game.

So for the sender, this means for each t where, remember, t is my type, and it could be either low ability or high ability. So for each t, if my type is t, choosing et-- remember, this could either be el or eh depending on the value of t, is optimal given w.

Remember, what is w here? w is the wage that the receiver sets. And which education level is optimal for me depends on w because what I'm thinking, what if I go to school a little bit less or a little bit more, what if I deviate? Well, to see whether that's worth it, I have to trade off the cost of, say, additional education against the effect that would have on my wage. And that's going to be governed by the strategy w that the receiver uses. So that's going to be relevant in determining whether et is optimal.

And then for the receiver-- well, with the receiver, this is actually-- we can write this maybe a bit more easily. So for each e-- and e is in 0 or 1, the receiver's wage must be optimal given their beliefs. So we need choosing w of e is optimal given b h e.

So what we're saying is I'm the receiver. I observe an education level e. That causes me to form beliefs b h given e about your type. And given those beliefs, the wage that I actually choose as the receiver is optimal.

Can we write that out in math? I mean, given what we said above, we said that it's always optimal given my quadratic loss utility to always charge a wage equal to my updated expectation about your type. So can we write this condition in words? What does it mean for w of e to be optimal given b of h given e? What must w of e be? Yeah?

AUDIENCE: Expected value 2.

IAN BALL: OK. And how can we write that in terms of b of h given e?

AUDIENCE: Um--

IAN BALL: So you're right.

AUDIENCE: Wait, are all-- other numbers. So then the l top were-- Oh, is it h times b of h plus l times minus 1 minus--

IAN BALL: Exactly right. Yeah. So indeed, it is the expectation of t, but we have to be careful because our expectation of t is influenced by the education we see and our beliefs. So it's not just our prior expectation. And maybe I should be clear here, when I write this, this is the updated expectation. This is the expectation of t that reflects the new information that I have given your educational choice.

And indeed, if seeing e caused me to believe that this is the probability that you're a high type and 1 minus that is the probability that you're a low type, well, now I can compute the expected value of your type.

And, yeah, I should emphasize here that l and h are actually numbers. Sometimes we think of high type and low type as just labels, but here, a low type is like 0.7 and a high type is like 2.3, and these numbers actually have meaning. Yeah. OK. So indeed, we get h times p of h given e plus l times-- well, we could write b of l given e, but using the notation we're using, we'll just write 1 minus b of h given e. So that's pretty easy to write down.

And now, belief consistency. Well, remember, belief consistency says the receiver has to be applying Bayes' rule wherever possible. And the receiver can only apply Bayes' rule when they see a level of education that someone is actually supposed to choose with positive probability. So we can separate-- sometimes this is called on-path and off-path. Let's say e prime.

So if they see an education level e that's on-path, meaning that either the low type or the high type is supposed to choose this with positive probability under their equilibrium strategy, then we apply Bayes' rule.

But if we see an off-path education, we see a level of education that is not supposed to be chosen by either type, and then there are no restrictions. And let's be a bit more precise here, which education levels are on-path given the strategy that we've written down for the sender? Well, there's only two that are on-path. It's exactly the numbers el and eh. Those are the only two education levels that the receiver expects to see in equilibrium. So these are precisely the cases where e is either el or eh.

If the sender was using a mixed strategy, this would be a bit more complicated. There may be more education levels they might choose. But I, as the receiver, know that, well, if you're low type, you're supposed to choose this education level; if you're high type, you're supposed to choose this education level; and if I see any other education level, well, I wasn't supposed to see this in equilibrium, so I'm not able to apply Bayes' rule.

Of course, this is quite a simplification. In reality, we see a lot of education levels, people's types are not binary, and we'll do extensions later, but for now, this is where we are. And then we have e is not an element of el and eh. OK.

So to start dealing with Bayes' rule, it turns out that it's helpful to distinguish two different kinds of equilibria, and that's what we're going to look at now. So we're going to consider two kinds of equilibria what are called separating equilibria and a pooling equilibria.

So what defines a separating equilibria is that el is not equal to eh. So let's think through this. What does this mean? What does the receiver learn? If we're in an equilibrium where the low type chooses one education level and the high type chooses a different education level, what can the receiver conclude when they see the sender's education level? What do they learn?

Well, they perfectly learn the sender's type because in a separating equilibrium, the implication of that-- let's write it this way. If I'm the receiver, and I see that you chose a high level-- that you chose eh, I know for certain that you must be a high type. And conversely, 0 here.

Why? Well, if I see eh, the low type never chooses eh. They always choose a different education level, el. So if I see eh, I know that you must be a high type, and therefore, I assign probability 1 to you being a high type.

Conversely, if I see a low education level, then I know that the only kind of person who would choose a low level of education, at least in this equilibrium, is the low type, so I assign 0 probability to being the high type. And then what about b of h given e for all other e? What about for e that aren't equal to el or eh? Well, there's no restriction on these. So these are unconstrained. And we'll fill these in as we want, but this is going to be unconstrained.

OK, now let's look at the pooling equilibrium where el equals eh. And sometimes it's nice to give a label to this. So let's call this e star. So here, both the high types and the low types choose the same level of education, e star. So now, what is b of h given e star?

Well, if everyone choose the same level of education, their choice of education hasn't revealed any new information to me. So my belief about them should be the same that it was, the same as my prior. And my prior belief was pi that they were high ability. So this must be pi. Because-- I mean, you could write out Bayes' rule, but I think the basic idea is if I thought that the probability that your high ability was pi, and then I know that whoever you are, you're going to choose the education level e star, then the fact that you chose e star doesn't give me any new information about your ability. So it's-- this must still be pi.

And then again, we have flexibility is unconstrained for all other e. And part of constructing equilibria is going to be to think a bit about what those beliefs have to be. So let's start with pooling. I think pooling is actually an easier case, so let's start there. So let's try to construct our equilibrium. So case 1 is the pooling equilibrium.

And for the sender, we know that they're choosing el equals eh equals e star. And we're going to try to solve for this value of e star, but we're just going to put it in as a placeholder for now. And now let's look at the receiver. The receiver needs to choose w of e-- let's think through as the analyst what we want to do here. We're trying to construct an equilibrium where both the low type and the high type choose this education level e star. And now we're going to try to choose what this wage is to make e star optimal for the sender. So any ideas about what might be--

Well, let's think a little bit about this wage. Let's separate into two cases. What is w of e star have to be? We're going to argue that if we're looking for an equilibrium where both the low type and the high type pool on this education level e star, and now we're going to try to fill in, in this equilibrium, what the receiver's wage strategy is. And I argue that here, we basically know what it has to be by the combination of belief consistency and sequential optimality.

Well, we know what the beliefs are here. We just said the beliefs are that your ability is high probability pi. And if those are our beliefs, well, over-- where did we write it? Here, we said if those are your beliefs, we computed what the wage has to be. So let's plug in pi for this and 1 minus pi for this. And indeed, we get that this is h times pi plus l times 1 minus pi. So that's pinned down.

And I think this is a key observation when analyzing these signaling games. The question is, I said that your beliefs off-path are unconstrained. So if you observe-- when the receiver observes a level of education that's different from e star, we can fill in whatever beliefs we want for that kind of receiver, but we can't fill in whatever wage we want. What is the range of wages we could possibly put in here that would be consistent with receiver optimality? How high or how low could these wages be?

Well, what's the most pessimistic belief the receiver could form? What's the worst belief they could have about the-- they could think you're certainly a low type. And if they think you're certainly a low type, what is the optimal wage to charge? Just l. And if they think you're certainly a high type, then the optimal wage is h.

And any belief they form is going to make it so that the optimal wage for them is going to be somewhere between l and h. It can't be lower than l because there's no belief that would make the expectation lower than l. It can't be higher than h because there's no belief that would make the wage the optimal wage higher than h, but we can actually choose anything between l and h. Because whatever choice we make between l and h, we can then construct some belief that justifies that wage choice.

So let's try to think this through. If you're trying to construct a wage that makes-- we want to make this an equilibrium. So let's maybe think of it graphing it down here. Here's e star. Here's our wage w. The wage here is pinned down to be this. Maybe I'll give notation for this. I'll just write it down. I'll just-- that's here.

And now the question is, what choices of wage should we make off-path to ensure that the sender actually wants to choose an education level of e star? So here, we're trying to construct an equilibrium. We don't want the sender to have any profitable deviations. So any ideas about what wage we would set off-path to encourage the sender to choose this education level e star?

Well, what's the sender tempted to do? They're supposed to choose e star in equilibrium. They're attempting to choose a lower level of education because if they choose a lower level of education, that's less costly for them. So the only way to make sure that that's not a profitable deviation is to say, well, if you choose any less education, your wage is going to drop a lot.

And let's look at the extreme case where the wage drops all the way to l. So we're going to think of a wage function that says, if you don't get enough education e star, you're going to get paid a very low wage of l. But if you go to college, if you achieve this certification, if you go to school for long enough, now you can get a job that pays a higher wage of this. And then what happens if e is even higher than that? Well, we want to make sure that no one wants to get even more education than necessary, so we might as well just keep the wage flat.

And if we keep the wage flat, then we know that, well, if someone can get the same wage by choosing e star as choosing a higher level of education, they would never pay more to get an even higher level of education if the wage that they're going to get is exactly the same. So this is, I think, a good conjecture for what this equilibrium might be.

So let's write this out more formally. What we have is w of e equals l if e is less than e star. Strictly less than e star. And it's h pi plus l1 minus pi. And then if you wanted to fill in beliefs, let's just be really precise, our beliefs are-- well, if e is less than e star, the firm charges-- pays a wage of l. The only way it's optimal for them to charge a wage of l is if they believe you're certain to be a low type.

So we must have 0 here. And then if they see a level of education that's at least e star, they're paying exactly this wage. So what must their belief be if they see a level of education at least e star to make that wage optimal given their belief? Yeah? Let's think.

If you're a high type, what's the optimal wage? If I know you're a high type, now my optimal wage is h. But the wage they're charging is this, which is h times pi plus l times 1 minus pi. So I need to find the belief that makes my expected-- so that makes the expected type equal to this. Pi, exactly.

If I ever-- if there was some education level that made me believe you were a high type, then I'd want to charge you a high wage there, and now there might be a deviation because now someone would want to get that education to prove that they're a high type. So let's think through what's happening here.

The firms are saying, look, there's some minimal education level e star that you need to get. If you get less than that education level, I don't think you're very high ability, I think you're certain to be a low ability, and I'm going to charge you a low wage, namely l. But if you get enough education-- let's say you graduate from college, then I'm going to form a more optimistic belief about you. I'm going to believe that you're just like anyone from the population. You're a high type with probability pi and you're low type with probability 1 minus pi. And because I believe that, I'm going to charge you a wage that's equal to this.

So we have this two-tiered wage system where if you graduate from college, you get a wage premium, but at least in this equilibrium, further education doesn't further increase your wage. And that's consistent with some settings, we could also construct other equilibria.

But we're not quite done yet. We're going to have to be careful about the value of e star to actually make this in equilibrium. So let's check-- let's see what we've chosen, what we've satisfied here. Belief consistency is good because the only time belief consistency has any bite here is when e equals e star. That's the only on-path level of education, and we're consistent with Bayes' rule, so we're good with belief consistency.

And by construction, the receiver is always paying the wage that's equal to their updated expectation of the sender's type. So we're good for the receiver. The one thing we still have to check is sender optimality. So what could go wrong in this equilibrium? How could sender optimality potentially be violated if we choose e star wrong?

Well, the problem is, if e star is really, really big, I might say, look, yes, I get a higher wage if I get at least e star of education, but that education is so costly to me that it's not worth me going to school to get that higher wage. So we need to check that e star is intuitively-- e star is going to have to be low enough that it's actually worthwhile for people to go to college. Just because there's a college wage premium, if the college wage premium is smaller than the cost of education, people aren't going to go.

And in fact, if we look over here, the college wage premium is not too far from the cost of education. So that is an issue. We have to make sure that the wage premium exceeds the cost of education. So we have a final condition here. So we need to-- it remains to check sender optimality.

So at first, it's tricky. We have to say, for both types, t equals l and h. Choosing e star is better than any other education e in 0, 1. Now this is a hard condition because it has to hold for both types. And we're comparing e star to any other level of education, and there's infinitely many other levels of education someone might choose.

So the key to simplifying things is to say, well, instead of saying it's better than any other education, let's try to find the best deviation. I could deviate from e star to anything, and we have to make sure that I don't want to deviate from e star to any other level of education. But some of these other levels of education, it's clear that I'm never going to choose.

So let's focus on the most attractive deviation, and as long as the most attractive deviation is not profitable, then no other deviation can be profitable either. So if I'm looking at this wage chart, this is the wage function I'm facing, and I'm choosing how much education to get. I'm already here at e star. What would be the most attractive deviation? I'm not saying it's profitable, but if I'm going to deviate, how should I deviate? What's the best possible deviation? Yeah?

0 is the only one that has a chance, because clearly, deviating above e star doesn't make any sense because I'm paying more in education and I'm not getting any wage benefit. So the only deviation that could be profitable-- and I should put "best" in parentheses. It could be that-- yeah, anyway. If I-- so I don't want to choose any education higher than e star. If I choose education strictly below e star, well, my wage is the same no matter what education I choose.

So since I hate school, according to the model, I might as well choose the lowest possible level of education. So the only thing we have to check is a deviation of e equals 0. And now we have to check that both for the high type and for the low type. So let's write down the two conditions we need.

So let's look at t equals l and t equals h. We want to compare what utility that person get if they choose e star versus 0. So if they choose e star, their utility is going to be-- I really need notation for this. Let me call this prior mean. Let me call this mu, mu for mean. So let me use-- just call that mu.

So if I choose e star, I get a wage of mu minus 1 over l e star squared. And that must be better than l. Maybe, just so we see what's going on here. OK, this is the key condition, let's go over it slowly. I'm a type l.

If I do what I'm supposed to do, I get paid a wage of mu. I choose a level of education e star. My cost of that education is 1 over l times e star squared. That comes from the formula for my utility up here. And we've plugged in for my wage mu because that's the wage I get if I choose e star. So let's remind ourselves here. This is w of e star.

If instead, I choose zero education, well, my wage is l because we see over here-- and maybe I'll write that down here. This is w of 0. This is the wage that's offered to me if I choose no education. And if I choose no education, this is the cost, but this cost is just 0, because if I choose no education, the cost is 0. So we don't even have to worry about this term. 0 squared is 0, I can just cross this out.

And now let's go down to l equals h and write down the same formula. Now I think a lot of people are tempted to put an h on the right-hand side, but that's not correct. Because remember, what is l representing? l is the wage the firm charges. If you get no education. If I deviate and get no education, I'm offered a wage of l whether I'm a high type or low type. So this is coming from w of 0. It's not about my true type, it's about what the firm is actually doing.

So here are the two conditions we have, and we need both of them to hold. It turns out, one of these conditions is going to be more demanding than the other. Which of these is harder to satisfy? We have two inequalities. Yeah?

AUDIENCE: The one on top.

IAN BALL: OK. And why?

AUDIENCE: Because the 1 over l times e star squared term is greater than the 1 over h times the e star squared term.

IAN BALL: Example right. So because l is smaller, 1 over l is bigger. So we're subtracting a bigger number, and that's harder to satisfy. And this makes sense. Let's think about what each of these inequalities say.

If we move these things around, what we're saying is, let's think of moving the cost to this side and subtracting l on this side. Let me write that down here. This is the key condition. So the key condition is mu minus l is greater than or equal to 1 over l e star squared. And this, I think, is a very natural interpretation.

Mu minus l is the, quote unquote, "college wage premium." This is the difference in wage I get if I choose an education level of e star versus education level of 0 because if I choose e star, I get mu; if I choose 0, I get l. So this is the wage premium.

And this condition says the wage premium is higher than the cost of education. And indeed, it's worth it for me to go to college if the wage premium I get from going to college-- that is, the additional wage I get from going to college more than outweighs the cost I have to bear of going to college.

Now, we also need to check that the wage premium exceeds the cost of education for the high type. But for the high type, going to college is less costly. So as long as we check that the wage premium is worth it for the low type, well, it's certainly going to be worth it for the high type because they get the same wage premium, but they find education less costly, and that's exactly the observation that was made at the end.

So this is our equilibrium condition. And as long as this is satisfied, then this is indeed going to be an equilibrium. So we actually have constructed many different equilibria parametrized by this level e star. There could be an equilibrium where we say if you go to school for 0.3-- for e equals 0.3, then you get a high wage. There could also be an equilibrium where if you go to school for 0.4, you get a high wage. And these are both conceivable equilibria, and there's a huge range of equilibria that we can construct here. Any questions on this?

OK, so now, let's look for a-- let's look for a separating equilibria. So, so far, we've just done pooling. Let's look for a separating equilibria. Now let's look for a separating equilibria.

So what that means is we have el is not equal to eh. So instead of just solving for e star we have to solve for el and eh. And now, let's try to think of what w is going to have to be.

So I think on a problem set, if you're trying to construct equilibrium, it's always very helpful to separate what the firm-- or what the receiver does on-path from what the receiver does off-path. Because on-path, Bayes' rule applies, and we really have precise control over what the wage has to be. Off-path, Bayes rule can't be applied because we divide by 0 when we apply Bayes' rule, and we have a lot of flexibility in what this wage is.

OK. So if we're looking for an equilibrium, a separating equilibrium, with these two levels of education, eh and el, what is the firm's wage going to have to be if they see an education level of el? Or maybe the first question is, what is their belief going to have to be, and then what is their wage going to have to be?

Well, what do they learn if they see someone choose education level el? They learn that they must be a low type because the high types always choose eh. Notice, this is a key difference with the pooling equilibrium. If all the types choose the same level of education, I don't learn about your type by seeing your education level. But if different types choose different levels of education, now, your choice of education reveals a lot of information to me as the receiver.

So when I el, I know that you're a low type and I must set the wage equal to l. And maybe in parentheses, this is because eh of el is 0. What about if I see an education level of eh? What do I conclude? Yeah?

AUDIENCE: I want to say type h.

IAN BALL: You must be type h. So now this is h. And now we have b of h given eh equals 1. And now the off-path wage, there's kind of a trick that we sometimes can use. If we really want to make this an equilibrium, we want to make it as unappealing as possible to deviate to any off-path level of education.

So how could-- if we have flexibility here to choose this wage, how could we choose this wage to make deviations unappealing for the sender? And this may be the most realistic equilibrium, but it's kind of a good trick for computing equilibria. So 0 would really deter deviations, but we can't do 0. Why can't we necessarily-- why can't we do 0? That would violate one of our conditions.

The problem is, it would violate r, and you may not have been here when we went over this, but this would violate r, and I think I just erased it, our receiver optimality condition. So off-path, Bayes' rule doesn't impose any constraints in our belief. But whatever our belief is, we still have to be behaving optimally given our belief. And the worst possible belief we could form off-path, worse from the standpoint of the sender, would to be certain that you're a low type.

So if we're certain you're a low type, the wage we would charge is l. So you're right that a wage of 0 would be a stronger punishment for the sender for deviating, but it's not a credible punishment because even if the firm thinks you're always a low-ability person, they'd rather charge you a wage of l, which is strictly positive than a wage of 0, but that's a helpful comment.

So here we are. Now I think we have some trouble because if I'm a low type, well, why should I choose el and get l if I could just-- well, any other education level I choose, I'll get l as well, so why don't I just deviate to education level equals 0? If I'm the low type, I think deviating to equals 0 looks pretty good.

So in fact, we immediately see a strong constraint on what happens in a separating equilibrium. So first observation, note, must be 0. If el was not 0, then the low type would deviate to choose an education level of 0. And let's think through why this is.

In a separating equilibrium, put yourself in the position of a low type. When you follow your equilibrium strategy and choose el, you're fully revealing to the firm that you're a low ability or that you're a low type sender, which is the worst possible thing they can believe about you. So it never makes sense to pay a positive cost to go to school, only to be believed that you're the worst possible type. So if I'm going to induce the worst possible belief about myself, I might as well pay no cost and not go to school at all.

So now el equals 0. Good. And let's try to graph this. Let's do our graph again, and I think it's actually helpful to try to represent the sender optimality conditions on the graph, so let's try to do this. So we're going to look at a wage here. el is going to be 0. eh is not equal to el, so eh has to be strictly positive. So let's just put a guess of where eh is. And we said that the wage is going to be l If you deviate-- if you do something off-path. So here's l. Here's h. And to here, the wage is going to look like this.

Now, at eh, the wage jumps to h. So this is an open circle, this is a closed circle here, and then the wage continues down. So this is, if we want to be really safe, what we would do. But to be a little more realistic, well, we want to make sure that no one wants to deviate to this high level of education. But this seems like-- I mean, it could be an equilibrium, but it seems kind of weird to say, if you get a really high education, we believe you're a low type. I mean, it's possible, but it seems a little weird.

So I think often what we'll do is we'll choose the off-path beliefs like this so that actually, if you choose any education level above eh, we'll still pay you a wage of h. Now why is that OK?

Well, if you can always get an education level of eh and get a wage of h, no one's going to ever want to deviate beyond eh to get the same wage that they would get at eh. So that going off your comment, to be really safe, we'd set we to bl, but really, the important thing is that it's l down here.

Beyond here, we can make it h and we're not going to introduce any deviations, so let's-- bl for e less than eh, that's kind of the important thing. And let's try to do it like this. So notice that this graph looks a lot like what we had in the pooling equilibrium, but there's a key difference. So like the pooling equilibrium, if your education was too low, we paid you l, but in the pooling equilibrium, if you got to a high enough level of education, we paid you the average ability in the population.

Now we pay you a lot more. We pay you h. Because now, if I see a high level of education, I draw a different inference. As the receiver, I say, oh, if you got a high level of education, you must be a high type person. So this wage is much, much higher, and we'll find that, not too surprisingly, eh is going to have to be pretty high as well to ensure that no one wants to deviate.

OK. So we have our wage function written down. The wage function satisfies receiver optimality by our construction. Let's just check in words. Bayes rule applies only at these two points because at this point, at e equals 0, that's the level of education chosen by the low type. We believe that you're the low type, and we charge you a wage of l. That's great.

Here, eh, If we see that, we believe you're a high type, and we charge you that-- we pay you the wage that's optimal given that belief. So Bayes' rule pins us down here. Everything outside of that is off-path, we had a lot of flexibility in our choice, and we made those choices to try to deter deviations by the sender. OK. So now, there's two things we have to check. We have to check that the low type-- so maybe I'll just put it over here. The low type prefers el to all other e, e prime, and the high type prefers eh.

So this is what sender optimality means. The type l sender in equilibrium is choosing el, which is 0. And we have to check that that's better for them than deviating to any other level of education.

Similarly, the high type, in equilibrium, is choosing eh, which is strictly positive. And we have to check that they prefer eh to any other level of education. Now I think this can be overwhelming at first because there's so many other levels of education that people might choose, but just like before, given this wage function, I think there's only one really candidate deviation. I think most of the deviations are clearly not going to be the optimal deviation.

So let's say I'm a low type. Here I am. I'm down here. I'm getting a wage of l. I see that this is what the wage function looks like, and this is the choice that people actually face when they design their lives. What level of education might I be tempted to deviate to? Yeah?

AUDIENCE: e of h?

IAN BALL: eh, right? None of the other deviations make sense. If I deviate to here, I'm paying more and I'm not getting any wage return. If I deviate to here, well, I am getting a wage return, but I'm getting the same return I could have gotten by just going to eh. So if I want to deviate, that's really the key deviation.

And similarly, now let's put myself in the position of someone of a high type who's choosing eh. They're here, and they're looking over the graph. What is the only reasonable or tempting deviation for the high type? Yeah?

AUDIENCE: el?

IAN BALL: Exactly. el, which, in this case, is 0. So el, for the same reason. I'm already getting a wage of h. There's no reason for me to go to school more. I don't like school, and it's not going to raise my wages, according to this strategy. If I go to less school, well, I'm going to get a big cut in wages, but I can go all the way down to here and I'm still going to get the same wage of l. So I might as well choose the lowest education level that ensures that I get that wage. So now let's write down these conditions.

So really, what we need to check is-- and sometimes I write it like this, in shorthand. So what I mean by this is, this, I say, this means type h prefers. This is a weak. Maybe I'll say something like this. So I need to check is if I'm type h, I prefer choosing an education level of eh to an education level of el. Of course, I have to prefer it to any other education level, but we said this is really the deviation I'm tempted by.

And then we have to check if I'm a low type l, I weakly prefer my own education level el to deviating to the education level eh. Now in general, I have to worry about every other deviation, but because of the special structure we've used when we define this wage function, we only have to check these two conditions. So let's just write this out in math.

If I'm a high type and I choose eh, then I'm going to get a wage of h, but then I'm going to have to pay 1 over h eh squared. Maybe let's go down here. What if I'm a low type-- no, we'll just go through it. Let's see. Now if I deviate to el, well, then I get l. Maybe I'll write it out. It's 1 over h el squared, but that's just l because el is 0.

And I think as usual, let's simplify this a little bit. That is, h minus l is greater than or equal to-- and this has the same interpretation. Now the wage premium is different in the separating equilibrium. The wage premium is now h minus l.

And what this condition says is the wage premium is higher than the cost of education for the high type. So going to college and getting an education of eh is worth it for the high type because their cost of doing that is weakly smaller than the wage premium they'll get from doing that.

Now let's look at the same condition for the low type. So for the low type, well, if they choose their equilibrium education level of 0, they don't pay anything and they get a wage of l. So we get-- and again, maybe I'll-- it's l, but if we really write it out, it's this. Where this just second term disappears because el is 0.

And this has to be greater than or equal to what? Well, if they deviate and choose a high level of education, they're going to get the wage h, but then they have to pay 1 over l eh squared. And again, let's simplify. And what does this say in words? Can we interpret this? Yeah?

AUDIENCE: The wage premium is smaller than the cost of, I guess, getting education eh for the low type.

IAN BALL: For the low type, exactly. So-- and this is-- if you recall, it's similar when we talked about the essay problem last week, and we said the cost of the essay has to be cheap enough that the excited students take it, but expensive enough that the less excited students take it, that's exactly what's going on here.

In order to sustain a separating equilibrium, because it's a separating equilibrium, the wage premium is going to be h minus l, because high education fully reveals you're a high type, low education fully reveals you're a low type, so that defines-- the fact that it's a separating equilibrium defines what the wage premium is going to be.

But then we have to check that it's actually optimal for each type to follow what they're supposed to do. In this equilibrium. The high type people are supposed to get educated, so the level of education eh must be low enough that it's worth it for high ability or high type people to choose to get that education because the wage premium outweighs the cost for those types.

At the same time, the level of education must be high enough. If it was too low, then even the low types would also choose that education, and we would no longer have a separating equilibrium. So the level of education must, at the same time, be high enough that for the low types, it's not worth going to education-- going to school because the cost exceeds the wage premium that they get.

And if we do this, we can get a formula for what eh has to be. So maybe this is helpful. Let's pull the h over here, and what we get is eh needs to be between h, h minus l, and it has to be above l times h minus l. And let's just check that this works. Well, because l is smaller than h, this inequality does indeed-- is indeed satisfied for a certain range of types.

I think it's helpful to compare. So this is-- let's compare what happened in the pooling equilibrium. So this is separating. So in separating, the low types got no education and the high types got eh. In the pooling equilibrium, both types got e star. And e star satisfied the following property.

So how do these compare? If we look at these bounds, how does the pooling equilibrium compare? How does e star compare to eh if we look at these inequalities? Let's look.

Mu is an average between h and l. So mu is certainly smaller than h. So this is smaller than this. So we see a fundamental difference between the two kinds of equilibria. In the separating equilibrium, the low types get no education and the high types get a lot of education. In the pooling equilibrium, both types get this lower intermediate level of education. That's e star is must be strictly lower than eh.

So we have two worlds. We have one world-- and they're both equilibria. There's one world where getting educated conveys a lot of benefits, and people-- or I should say, getting very highly educated conveys a lot of benefits. People pay a huge cost. The high types pay a huge cost to get highly educated, and they get a huge return.

There's also another equilibrium we could be in where people don't get as much as educated, there's not a huge benefit to education, but everyone gets it. And I think these actually are consistent with-- you might imagine in different countries, we might kind of reach different equilibria where education has a different role in society.

Both of these are consistent. Everyone's behaving optimally given what everyone else is doing, but here, the beliefs that the market forms about your education and this equilibrium are different-- sorry, the beliefs, the market forms about your ability given your education are different than the beliefs the market forms in this equilibrium, and then people are responding accordingly, and it's mutually consistent in each case.

Let me maybe just do one final thing. I think it's sometimes helpful to visualize these constraints geometrically. So I've written out the constraint for the separating equilibrium up here, but I want to try to visualize this. And I think a good way to visualize this is to say, well, we want to look at something that is called an indifference curve. And if you've studied other economics classes, you would have seen this.

But what we want to say is, well, we're in the plain here where the horizontal axis is your level of education and the vertical axis is your wage. And we want to ask, what are the pairs of wage and education that give you the same utility?

Well let's write the formula down. So for the low type, or for type t, an indifference curve would have the form, say, u bar equals wage minus 1 over t e squared. So what I'm saying is I want to look at all the pairs we that give type t the utility u bar. And if I solve-- if I write this out, what I get is w equals u bar plus 1 over t e squared.

So what this tells me is w is going to be an increasing quadratic function of e squared. So let's try to write this down. It might look something like this. What this indicates is the set-- what I'm saying is if I moved along this curve, the utility that I would get-- say I'm the low type here-- is the same along this curve.

Because if I move a little bit to the right along this curve, I'm getting more educated, that's more costly, but my wage is also increasing at just the right rate to ensure that my utility stays constant. And then we could draw another utility curve. By shifting the value of u bar, we could draw another utility curve like this and another one like this.

So what this says is, these all give me the same utility, these all give me the same utility, but that utility is going to be higher than this utility. And similarly, these all give me the same utility, but a higher utility still than the previous case.

Now, let's say these are the indifference curves for the low types. Now I want to draw the indifference curves for the high types. And this is tricky. Are the high types indifference curves going to be steeper than the low type or shallower than the low type? This is-- we can do the algebra, but I think it's good to think it through.

So they're also going to be quadratic. They're going to have a similar shape. Is that shallower or steeper? What do you think? You have a half-chance.

Well, let's think it through. What does this curve mean? It tells me, if I get this much more education, I have to get paid this much more to compensate me for it and keep me indifferent. So it is a high type need to be compensated more or less for the same level of education.

AUDIENCE: More.

IAN BALL: Actually less because a high type finds education less costly. So if I'm really good at school, and I go to school a little bit more, I don't need to get paid too much more to keep me indifferent. If I'm really good at school, I can go to school for 20 years and you just need to pay me a little bit.

So here, I'm going to get curves like-- say something like this. No, these should not intersect. Sorry. Something like that. Let me-- I forget this last one. So these are the indifference curves for the high types.

And in fact, using these curves, we can exactly represent graphically the levels of education that could support a separating equilibrium. So it's exactly going to be this range here. Do I have-- and let's just see it in words.

This dotted line here is the indifference curve that passes through 0l for the low type. So these are the wage-education level pairs that make the low type indifferent with 0l, with no education and a wage of l. So if I want to have a separating equilibrium, I better make this eh high enough that the low type doesn't want to deviate to eh. And that precisely means eh has to be farther to the right of this point because if eh is over here, that tells me the low type strictly prefers to get no education and a wage of l than to get this level of education and a wage of h.

On the other hand, the high type must weakly prefer to get high education than no education. Ah, I misdrew this. It's going to go all the way over to here. Sorry. So the high types indifference curve that passes through 0l goes all the way up to here. So just to maybe go through it one last time, let's make sure we understand.

If I could get a level of education here and get paid a wage of h, that would be better than no education and a wage of l for both high and low types. So that couldn't be part of a separating equilibrium. If I chose eh to be way over here, then getting this level of education and a wage of h would be worse than 0l for both types.

The only way I can have a separating equilibrium is if I lie exactly in this range. And if I lie in that range, I ensure that this additional education is worth it for the high type, but not worth it for the low type. And the crucial point here is I'm looking at the indifference curves that pass through the point 0l. So these are considering the pairs of education and wages that--

So let me just-- as I'm concluding, let me say it smoothly in words one last time. These are the pairs of wages in education that make the high type indifferent with this point, and these are the pairs of education and wages that make the high type-- that make the low type indifferent with this point. And everything in between is going to ensure that the high type finds it worth it to get educated and the low type finds it not worth it. We did the algebra, but it's just good to see the graph, and I think on the problem set, you'll be asked to illustrate this graphically.