

[SQUEAKING]

[RUSTLING]

[CLICKING]

**JONATHAN
GRUBER:**

All right, let's get started. Today, we're going to continue our discussion of factor markets. If you recall, last Monday, we started talking about the labor market. And we talked about how workers make the decision between work and leisure. And we talked about the implications for setting the wage rate in the labor market.

What I want to do today is return to that labor market equilibrium and talk about the important case of the minimum wage. So today, I want to talk about the labor market equilibrium and how it's affected by the minimum wage because it's an interesting case which allows us to introduce some complications as to how we think about the labor market.

So let's go back and think about the labor market. So let's go to figure 16-1. The labor market, like any other market, has a price and a quantity. The quantity is the amount of labor supply. That's on the x-axis. The price is the wage. That's on the y-axis.

The supply curve that's upward sloping-- typically we'll assume an upward-sloping supply curve. But as we discussed last time, that doesn't have to be true. If income effects dominate substitution effects, which they very well may, you could actually have a backward-bending or downward-sloping supply curve. So we talked about that last time.

Having taught that interesting case, typically, we'll assume supply is upward sloping or at least not backwards bending, not downward sloping. But remember, that's an assumption. So this upward-sloping supply curve is not necessarily as obvious as a downward-sloping demand curve is. Downward-sloping demand will almost always exist unless there's a weird Giffen good, whereas upward-sloping supply is a little more questionable.

So we have the equilibrium, and we have this equilibrium at L1 workers at a wage

W1. So now we know where this comes from. So basically, going all the way back to producer theory where we just gave you a W , now we're telling where the W comes from. We're telling you where the wage comes from that you then plug into the firm's optimization for them to produce goods.

Now, let's imagine that we have a minimum wage. So let's go to figure 16-2. So this is a regulation which says that you're not allowed to pay workers below some minimum level. And let's say we set that minimum wage at the level W_2 above the market wage W_1 .

Quick question. What would happen if we passed a law and set a minimum wage that was below W_1 ? So there'd be a regulation which insists you couldn't pay workers below W_2 , but W_2 is below W_1 . What would that do to the labor market? Nothing.

And here's the key point. Markets in economics will always endeavor to avoid government regulations if they can. So if a government regulation is not binding, it won't matter. Markets will just avoid it. So the interesting case is only where the minimum wage is binding, as in the figure 16-2.

So what happens? Well, if you set a minimum wage at W_2 , workers at that high wage would love to work a lot. That's a high wage. They're high in the supply curve. They would like to work $L_{sub s}$ hours. They would like to supply $L_{sub s}$ amount of labor supply to the market.

Firms, however, if forced to pay a high wage, W_2 , are going to say, wait, I'm only going to pay that high wage if the marginal revenue product of labor is sufficiently high. Remember, we talked about the marginal revenue of product last time. It's the marginal product of labor times the price.

So if you're going to raise the wage I'm going to have to pay workers, unless that affects the market price, I'm going to need to have a higher marginal product of labor, right? The demand equation was, I said, the wage equal to the marginal product of labor times the price.

Well, if the price hasn't changed with the minimum wage going in, I'm going to need a high-- if the wage is forced up by the minimum wage, I'm going to need a higher

marginal product of labor. How do I get a higher marginal product of labor? By hiring less workers because the marginal product of labor's diminishing. So if you're going to force me to pay a higher wage, you're going to force me to only hire workers until the point where the marginal product of labor justifies that higher wage, which means I'm going to hire fewer workers.

So firms demand only L sub d . Well, workers can't get jobs firms don't want to give. So the equilibrium is L sub d jobs at a wage W sub 2 , OK?

What does this do to welfare? We can see before, before the minimum wage was in place, the market featured a consumer surplus that-- here, consumers are firms, right? But there was a consumer surplus of A plus B plus C . That is, firms were willing to pay what was on the demand curve. They only had to pay W_1 . So their surplus was A plus B plus C .

Workers were willing to work at a wage that's given by the supply curve S sub 1 . They were paid at W sub 1 . So they got a surplus of D plus E . So here, the firms get the consumer surplus. The workers get the producer surplus because the workers are now the producers.

Now let's say you roll in a set minimum wage. Well, two things have happened. One thing is you've then transferred some resources to workers. That's the area B . You've taken the area B that firms used to get, and now workers get it. That's the idea. You want to make workers better off. So you transferred to workers the area B .

On the other hand, you've created a deadweight loss of the area C plus E . You've created deadweight loss in the area C plus E because now there are fewer jobs. There are workers who would happily work at a higher wage who are not being allowed to work by the limited demand that comes from the minimum wage.

So the bottom line is you end up with fewer workers, a higher wage, and ambiguous welfare implications. Clearly, social welfare goes down. Whether worker welfare goes up or not depends a bit on the size of area B versus the size of area E . It's not clear if worker surplus goes up or not. It depends on size of B versus E . In this diagram, workers are a net better off, but it doesn't have to be true.

What's clear is that social welfare has gone down. Because remember, as I talked

about, the cheat, the shortcut I talked about when we talked about oligopoly, is, roughly speaking, welfare is proportional to the quantity in the market. Essentially, the further you deviate from the perfectly competitive quantity, the bigger the deadweight loss. So that's what happens if you put in a minimum wage. Questions about that? OK?

Well, that seems pretty straightforward, and that's what I learned growing up as a kid in economics class. But then some empirical economists, some very famous empirical economists, started doing a series of articles that actually studied, gee, what happens when the minimum wage does change. They did things like, for example, comparing what happened when New Jersey raised its minimum wage but the state of Pennsylvania next door did not, and looked at fast food workers in New Jersey, where the minimum wage went up, compared to fast food workers in Pennsylvania where the minimum wage didn't go up.

And what they found was there was no difference in employment, that jobs didn't fall in New Jersey even though the minimum wage went up. And a series of follow-on studies continue to find that, actually, higher minimum wages didn't seem to cause jobs to fall, which is directly in contradiction with this graph.

So what's going on? That led to a big question and revision of what's going on in these markets that leads to that. And there's really three possibilities for what's going on.

Possibility one is that the minimum wage wasn't binding. Maybe New Jersey set a minimum wage below the market wage. But actually, empirically, that's not true. We can look at what workers were paid before the minimum wage. It was well below where the minimum wage was set for restaurant workers that were studied in that most famous study. So this is not true. The minimum wage was binding.

There's a second possibility that's absolutely consistent with a perfectly competitive market. What's a possible answer for why I could impose a minimum wage in a perfectly competitive labor market and have employment not go down? Yeah?

AUDIENCE: Price goes up.

JONATHAN The price that the firm charges goes up. But in a perfect competitive labor market,

GRUBER: that still wouldn't happen. You might see some price adjustment, but you'd still see some adjustment in the marginal product of labor. But what else about this diagram? Yeah.

AUDIENCE: The firm's demand for labor is perfectly inelastic.

JONATHAN GRUBER: The firm's-- actually, you're close. It'd be the worker's supply of labor is perfectly inelastic. It's the right idea. If workers are perfectly inelastic in their supply of labor, then the same amount of workers will work no matter what the wage.

So basically, you're just going to essentially end up-- you'd also, in fact-- that's a good point-- also get inelastic demand, the same thing. If either supply or demand is inelastic, you'll end up with no effect of a minimum wage. So that's another possibility.

But in fact, we've done a lot of studies. So you could have inelastic supply or demand. But in fact, we've done lots of studies of supply and demand in these markets, and that's not true. Remember, supply was largely inelastic for men, but it was somewhat elastic for women. And these low-income markets have a good mix of men and women working in them. Demand has been shown to be somewhat elastic. So neither supply nor demand's very elastic, but they're sufficiently elastic that that rules out as zero.

So the third possibility and the one economists have focused on is that we're not in a competitive labor market. They're focused on a noncompetitive labor market. Just like we discussed noncompetitive markets for goods with a monopoly and oligopoly, you can have noncompetitive markets for labor. It's the basic same idea.

So now let's look at-- so when we thought about-- let's go back, think about perfect competition, the basics of perfect competition. We thought about perfect competition. The basic idea was, remember, I talked about laying out a bunch of rugs in a market where you could literally shop costlessly across all the people selling their little fake Eiffel towers, little statue Eiffel towers. And you could perfectly shop. It was easy to go from carpet to carpet. There was full information. The prices were posted.

And so basically what you ended up was perfectly elastic demand facing any given

firm. Any given firm, if they tried to charge one cent more for their Eiffel tower, no one would buy it. If they charged one cent less, they'd immediately run out. Everyone'd buy it.

Well, when we are modeling labor markets-- and I discussed this last time, but not very well. So I want to come back to it. When we're modeling labor markets, we're thinking about the same feature of perfect competition. But here, it's not consumers shopping over where to buy their goods. It's workers shopping over where to work.

It's workers saying, gee, in a perfectly competitive labor market, the idea is I know what I could earn at any firm and I can easily shop across firms, see where I'm going to work. So if any firm tried to pay me one cent less than the market wage, I'd never work there. And if they tried to pay me one cent more than the market wage, every worker in the world would want to work there.

So in a perfectly competitive labor market, any given firm faces a perfectly elastic supply of labor. So we can see that in figure 16-4, which we actually showed-- and I'll let you skip this since we covered it-- 16-4, which I actually showed in the last lecture. Remember the last lecture. I was focused on this downward-sloping demand curve, but I casually threw in this flat labor supply curve and botched explaining it.

Now I'm explaining it, hopefully more clearly, which is to any given firm, the labor supply curve is perfectly elastic because workers can perfectly shop across job opportunities. So if that firm tried to pay less, they'd get no workers. So they faced a perfectly elastic supply of labor.

But just like, in reality, there's no such thing as a perfectly competitive product market, in reality, there's no such thing as a perfectly competitive labor market. In fact, we can't shop easily across all possible jobs and know what every job could pay. And the fact that we can't means that firms on the labor market side will have market power. Just like we talked about monopolists and oligopolists having market power over consumers through barriers to entry, firms will have market power over workers because workers can't perfectly shop across their job alternatives.

So as a result, firms may be able to get away with paying you less than what you might earn elsewhere. In a perfectly competitive labor market, a firm could never

pay you less than what you're worth elsewhere because you'd just go work somewhere else. But now, if McDonald's wants to pay you less than you might get at Wendy's, but it's hard to go find out what Wendy's going to pay you-- you have to go a distance down the road, and you have to ask them, and you're shy and it's embarrassing-- then McDonald's might be able to get away with paying you less than you might earn at Wendy's.

So this is very much parallel to monopoly. In fact, we call this a monopsony. A monopsony is a labor market where firms have market power over workers just like a monopoly is a goods market where firms have market power over consumers. Now, this is not so crazy. And in fact, it applies very much to me.

Think about my situation at MIT. I've been here 25 years. I just got my 25th year rocking chair, although actually it's not a rocking chair because it comes in the box with the rockers off it. And it arrived in my office, so it's sort of a short chair. My wife's 5 foot, and she always complains how chairs are too big for her. So she sat, and she's like, it's a perfect chair for me. So now I have a nonrocking rocking chair in my office that she sits in.

But anyway, I've been at MIT for 25 years. It's going to be really hard for me to move. I like my house. I like my colleagues. I like my friends. Kind of, I like my view out the window. It's going to be kind of hard for me to move.

Moreover, it'd be pretty hard for me to figure out what I'd get paid if I moved. I can't go to other universities and say, hey, what would you pay me if you hired me? That's be awkward. I can't really ask my colleagues what they make. That's awkward. So at the end of the day, MIT has market power over me because I don't really want to move and I can't really figure out what I'd get paid if I did move.

And MIT will exploit that market power over me by paying me less than I might earn elsewhere. And we know this as a fact because in academia, the only way to get a raise is to go get an offer from someone else and have them say how much more they'll pay you, and then you take that to your boss and they say, match this. But if you're not willing to do this, as, frankly, MIT knows I'm not willing to do, then MIT can essentially underpay me.

So basically, any responsible profit-maximizing or even nonprofit employer will

exploit this market power and they'll pay me less than my market wage. And that means that MIT will earn surplus on me. In a perfectly competitive labor market, the firm earns no surplus on the worker. They pay the worker their marginal revenue product.

So if you go to this figure, what am I paying the worker? What I'm paying them is exactly the marginal revenue product just like, in a competitive market for the goods, a firm is selling at exactly their marginal cost. So just like a firm makes no surplus in a perfectly competitive goods market, a firm hiring workers makes no surplus in a competitive labor market. But in a monopsony market, the firm makes surplus over me. They pay me less than they'd have to because I don't shop and find a better opportunity.

Now, are there questions about how that market works? I'm not going to do all the math and graphs. It's all the same as monopoly, just flipping demand and supply curves. It's a pain in the ass. I'm not going to do it. I just want you guys to understand the intuition. So please, since I went through this, are there questions about this or how it works? OK.

Now let's take this noncompetitive labor market and let's throw in a minimum wage. Well, as before, if the minimum wage is below what the firm was already paying, there's no effect. So let's assume it's a binding minimum wage. Now, let's say the binding minimum wage is above what my true market wage would be, what my wage would be in the perfectly competitive market.

So in a perfectly competitive market, my wage would equal my marginal revenue product of labor, right? That's in a competitive market. In this noncompetitive market, my wage is below my marginal revenue product of labor. Firms are exploiting me because I can't effectively shop for a better job. I don't want to or it's hard to do so.

Now, in this noncompetitive market, if we set a minimum wage that's higher than the marginal revenue product of labor, then the analysis is just like it's a competitive firm. Once that marginal wage is higher than the marginal revenue product of labor, it's just like a competitive firm. So it's not that interesting.

The interesting case is, what if the minimum wage comes in and it's above the wage I make but below the marginal revenue product of labor? So let's say McDonald's, someone working there yields a marginal revenue product of labor of \$10, but they're only being paid \$7. Let's say you roll in minimum wage of \$9-- so above what they're being paid now, but below their actual marginal revenue product of labor. Will the firm fire that worker? Why not? Yeah.

AUDIENCE: They're still paying them-- they're still making a profit off of that worker.

JONATHAN GRUBER: They're still making surplus, which is as long as the marginal product of labor's bigger than the wage, they love that worker. So before-- so let's write down the numbers as an example. So imagine my marginal revenue product of labor at McDonald's is \$10, but my wage is \$7. And then you come and you set a minimum wage of \$9.

Well, 10 is still greater than 9. So the firm has no desire to fire me. So all you've done is just given me money. And where'd that money come from? The surplus the firm earned. So all you've done is shifted the surplus from-- you've shifted producer surplus to consumer-- I'm sorry, consumer surplus-- consumers are the firms-- to producer surplus, the workers.

So in a monopsony market, a minimum wage doesn't cause deadweight loss. It just shifts surplus around. And that's a really important outcome because that, once again, says the government isn't always bad here.

This is just like-- if you want to think about this graphically, go back to exactly the analysis we did of regulating monopolies. Remember we talked about regulating monopolies. We talked about, if a regulator comes in and sets a price below the monopoly price but above the competitive price, it reduced the deadweight loss of monopoly. It's the same thing. And if you set a minimum wage above the market wage but below the marginal revenue product of labor, then you simply transfer surplus to workers without causing deadweight loss.

Now, that raised the question, of course, is the minimum wage in between the wage of the marginal product of labor? Well, we don't know, but let's go back to the studies that motivated this. The very fact that the minimum wage doesn't seem to cause unemployment suggests we are hitting the sweet spot, suggests we are

hitting the sweet spot, that we're basically managing, with the minimum wage policy, at least to date, to essentially just find a way, without the government spending any money, to shift resources from businesses to workers.

So what does this mean? Well, it means that around the level of current minimum wages, we can raise the minimum wage by a small amount pretty costlessly. It doesn't necessarily mean that a \$15 minimum wage is OK. So in some sense, the existing-- this is the important thing about empirical economics. You only learn the answer in the range that you study it.

So for example, there've been studies that have looked at what happens if you have a \$10 minimum wage, and those show no unemployment. There haven't been studies that show what happens if you have a \$15 minimum wage. Now, Seattle just actually put in a \$15 minimum wage about two years ago. So we actually can run the experiment.

And the early evidence is the Seattle \$15 minimum wage did lower employment, that the Seattle \$15 minimum wage actually went above the marginal revenue product of labor. And once it's above, you're back in the competitive case. You're back in the case where you're lowering employment. Yeah?

AUDIENCE: How can you increase competitiveness in the market?

JONATHAN GRUBER: Well, that's the other question, is how could you increase-- so you tell me. How could you increase the competitiveness of a labor market?

AUDIENCE: You make it easier to tell how much money you would get at each place.

JONATHAN GRUBER: So Norway has a day every year they call Envy Day, which was yesterday, I believe, where they literally can go online and look up anybody's income in Norway. They literally make public every single person's tax return in Norway. And you can go online and look at what everybody makes. That would do it. So you could provide more information.

You could make it easier to move between jobs. For example, there's a lot of restrictions in our labor market, like noncompete clauses, which say that if you work for one firm, you can't ever go work for another firm in that industry for x years.

That gives some monopsony power to firms, et cetera. So we could do things which try to loosen the flow of the labor market, and that would close this gap between wage and marginal revenue product of labor.

Now, let's go back to Seattle, just to conclude this. This doesn't mean the Seattle policy was a bad one. The bottom line is what we learned from Seattle was that basically, employment fell a small amount and a bunch of workers made a bunch more money.

So is that good or bad? Well, it depends. If you're one of the people that lost their job, it's really bad. If you're one of the workers who got a raise up to \$15 an hour, it's good. How do you weigh them against each other? That's exactly what we'll talk about in a couple lectures.

So once we start talking about normative economics, about is a policy good or bad, there's typically trade-offs. And this is a classic example.

What we're learning here is, is the minimum wage in the range we are now, right now, the federal minimum wage at \$7.25-- the evidence suggests it could easily rise without causing that trade-off. The evidence suggest we could increase the federal minimum wage by some nontrivial amount, at least up to \$9 or \$10, without causing much of a trade-off. But once you get too far ahead of that, there starts to be a trade-off. Question about that? Yeah.

AUDIENCE: Are there any states where it's actually still that low?

JONATHAN GRUBER: Oh, yeah. Many states don't have their own minimum wage. Massachusetts is at \$11, but we're pretty unusual. We're one of the higher ones. A number of states have \$7.25 as the minimum wage, OK?

And the evidence seems to be, from states like Massachusetts and others which are on the \$10, \$11 range, it doesn't seem to lower employment. It seems like we could clearly-- we'd be safe raising that federal minimum wage. We would simply be transferring resources and not causing unemployment. Yeah?

AUDIENCE: Is there anything about the cost of living in areas where the minimum wage is more expensive? Is it possible that if a McDonald's worker makes more money in this state, McDonald's is more expensive in that state?

**JONATHAN
GRUBER:**

That's a great question. So what I assumed was I assumed firms would just say, oh, you got me. I'm going to throw some of my profits at workers.

Firms don't have to do that. Firms could say, well, if you make me pay workers more, I'm going to raise my price. Now, if it's a competitive output market, that shouldn't happen, right? Because in a competitive output market-- well, no. Marginal cost goes up.

It's not clear. It's not clear whether that would happen or not, and the evidence is that it's unclear whether higher minimum wage causes higher prices or whether it just comes out of profits. We don't know yet, OK?

All right, so that's what I want to say about labor markets. Now I want to move on and talk about capital markets. Now, as confusing as our discussion of labor markets was, that's easy compared to capital markets. Capital market's a lot harder to understand.

And that's because capital itself-- labor's something you get your hands around. It's the time you spend at work. Capital is this sort of amorphous thing that I've kept pushing off defining. So I'll define it now. We talk about capital as this vague collection of buildings and machines and the other stuff that goes into production.

And we know where labor comes from. It comes from our work. But where does capital come from? Well, capital is a harder concept, but there's one unifying thread that all elements of capital have, which is they represent the diversion of current consumption towards future consumption. Capital is about diverting consuming today towards consuming in the future.

In fact, the original concept of capital came from farmers. Farmers, every year, when they would pick their grain, they had a choice. They could eat all the grain, or they could save some to plant for next year's grain. Now, the more they saved, the more they'd have next year, but the less they'd have today. So farmers faced a trade-off-- literally, consumption today or consumption next year. That's what we mean by capital.

In other words, in today's market economy, the link is not that direct, but it's the

same basic idea-- that firms have a choice, firms and their investors have a choice. They can take what they make and eat it now, or they can invest it in having more in the future.

So basically, when we think about capital, we're not going to think about capital as physical capital. We're really thinking about capital as financial capital. What links all types of capital is their financial aspect. What links machines and buildings is all the aspect that, by putting money into them today, you have less you can spend on fun stuff today, but more you'll be able to spend tomorrow. And it's this financial aspect that links all forms of capital.

Now, how do firms get the money to invest in machines and buildings and stuff like that? They get it through going to the capital market. Where do firms get this money that they invest? They get it through going to the capital market, which is basically the pool of money that firms can draw on to make their investments.

So think of it literally as I'm a firm. I want to build a building and buy a machine. I literally go over, and there's a big pool of money. And I have to take the money out of there to go buy my machine or build my building. And where does the money in that pool come from? It comes from household savings decisions.

So the capital market is a market where the demand for capital comes from firm's interest in investing and having more in the future. The supply of capital comes from people's decisions to save. And essentially, the money firms use to buy stuff is borrowed from people. And that's the bottom line of how capital markets work.

So just as the supply of labor that determines how many workers a firm can hire comes from your decision of how hard to work, the supply of capital that determines how many machines a firm can buy comes from your decision of how hard to save. So let's look at figure 16-5, equilibrium in capital markets.

Let's start with the demand. We already talked, last lecture, demand for capital. The demand for capital comes from the marginal revenue product of capital. It's the marginal product of the next machine. So the demand comes from the marginal product of the next machine times the price the firm can get for its output, which is the marginal revenue product of capital. So it's the same logic as for labor. There's nothing interesting there. Same logic as for labor.

The supply's what's more interesting here. Where does supply come from? The supply comes from household savings, how much money is around for firms to actually get to get these machines.

And how do they get it? They borrow. And what do they borrow at? They borrow at the interest rate i . So i represents the rate that firms pay households to get their money.

So think of this as-- we'll talk about how it really works. But in theory, the idea is think of literally a marketplace in the center of town. Downtown Boston, Haymarket, there's this marketplace. And a firm comes and says, I need to borrow money to buy a machine. And a person's there with their savings and they say, well, I'll loan you some money. What interest rate you going to give me? And that's the market for capital.

So where the supply of capital meets the demand of capital yields the interest rate. So basically, what this means is as the interest rate's higher, what that means is I have to pay people back more to borrow their money. So an interest rate of 10%, if I borrow \$10 from you, I pay you back \$1.10 next period.

If I borrow \$10, I pay you back \$1.10 next period. If the interest rate's 20%, if I borrow \$10-- if I borrow \$1-- I'm sorry. If I borrow \$1 from you, I pay you \$1.10 next period. If I have 20% and I borrow \$1, I pay you back \$1.20 next period, et cetera, OK?

So basically, that is essentially how the transaction works. And the key point here is the reason the supply curve is upward sloping is the more you're willing to pay me for my money, the more I'm willing to lend you. So if you come to me and say give me \$1 and next year I'll give you back \$1, I'm like, I don't know. Why would I do that?

If you say, give me \$1 and next year I'll give you back \$1.10, you're like, OK, now I'm interested. \$1.20, I'm very interested. \$1.50, for sure. Literally, I just give you my money and, next year, I get back 50% more? Why not? So basically, the higher the interest rate, the more I'm willing to loan the firm and, therefore, you get an upward-sloping supply curve.

Now, of course, in reality, people don't actually-- we don't sit in Haymarket, downtown Boston, and give money to firms. In reality, this transaction happens through capital markets. And essentially, there are three mechanisms by which implicitly I loan money to firms.

The first is I could literally buy corporate debt. I could literally loan the money to firms. I could literally go and the firm could say, I, General Motors, am issuing a bond. This is through bond, issuing a bond.

And the way that bond works is I promise that for every dollar you spend buying my bond, you'll get 1 plus I dollars back at the end-- or next year, say, depends on how long the bond is. So literally, you're loaning the money to the firm by buying-- you're buying their promise to pay you back.

Now, a second way you can loan money to the firm is through investing in their equity. You can buy their stock. The way this works is GM says to you, buy a piece of me and you'll get paid back not some fixed interest rate, but you get paid back according to how well GM does.

So with corporate debt, I get paid back something that's predetermined. When I buy stock or equity, I don't get back a predetermined amount. I get back some-- it depends on how well the company does. But it's the same basic idea. I'm giving the company some money today in return for my getting more money, I hope, tomorrow. That's the diversion of consumption from today to tomorrow.

And the third thing I could do is I could put it in the bank. Now, how is that loaned to companies? Because the bank then loans it to companies. Why do banks say they'll pay you interest on your money?

Why did banks going crazy-- I'll give you 1-- it used to be interesting. Now it's 1%, 2%. When I was a kid, I was like 10%, 12%. We'll give you lots of money. And we'll talk later about why it was so much higher when I was a kid.

Why are banks so eager to do that? It's not out of the goodness of their heart. It's because when you give them dollars, they turn around and loan them. They add a bunch to the interest rate and loan them out to firms. So those dollars you're giving the banks and they're paying you 2% interest, they loan to firms at 6%. And that's

why bankers are rich.

So basically, the reason a bank exists is because it's a way-- corporate debt and equity markets are hard and complicated. It's much easier to put your money in a bank. You put your money in a bank. But when you put your money in a bank, you're essentially loaning it to companies. That's essentially what you're doing.

So through these mechanisms, we have a capital market where essentially, by my putting money away and diverting from today's consumption, I'm loaning to a firm. They'll produce more, and they'll pay me back more in the future. Questions about that?

OK, so let's talk about where the supply curve comes from. We know where the demand curve comes from. It just simply comes from the marginal revenue product of capital.

Where does supply curve comes from? The supply curve comes from what we call intertemporal choice. As I said, economists like putting fancy names on things. That helps us get paid more money. It just means choosing over time, intertemporal choice. Intertemporal choice is essentially about how do you decide how much to save. What's going to determine that is going to be your decision of how much you value money today versus valuing money tomorrow.

So for ease, let's imagine I'm considering two periods, this year versus next year. When I talk about periods, I'm talking about days and years and whatever. It's the basic logic. It's about now versus the future. Whether I say days or years, it doesn't really matter right now. The point is I'm just talking about today versus the future. So let's talk about this year versus next year.

And let's imagine prices aren't going to change. I'll come back to prices next lecture. But let's imagine the price of goods aren't going to go up. There's no inflation in this economy, which is roughly true today.

And let's suppose I'm going to take next year off to care for my children. Lord knows why I'd want to do that when the youngest one's 19, but imagine they still need my care. So let's say I'll take next-- this example gets dated. Let's say I take next year off to care for my children. And let's say my income is \$80,000 a year.

Now, here is my-- but I'm going to take next year off unpaid. So I'm going to work this year for 80k. Next year I'm going to take off unpaid. So I have a couple of choices. I could work this year, earn my 80k, spend my 80k, and have nothing next year to live on.

I could work this year and eat nothing and save all of the 80k to live on, or some combination in between. And we could illustrate-- but the key difference is every dollar that I don't consume this year that I save to consume next year earns interest. And that's where the trade-off comes.

So let's look at figure 16-6. This is a familiar-looking optimization diagram. Now my optimization is not over pizza versus cookies, but my optimization is over consumption this period versus consumption next period. It's a bit mind-blowing. We're a little science-fictiony here, right?

We're now not talking about choosing between two goods, like leisure and consumption or cookies and pizza. Now I'm talking about two time periods, consumption today versus consumption tomorrow. But that's the key thing about the tools we learn with consumer choice. Those tools are incredibly powerful. You just need to shove your problem into that framework. And we're going to shove our problem into this framework.

The problem we're facing is how do I decide how much to save. Well, savings is a bad just like labor's a bad. What do we do when we have a bad to model? We don't model the bad. We model the complementary good. So our choice is, how much do I consume today? My choice is, how much do I consume today and how much am I going to save?

Well, saving is a bad, but the other way to think about it is, how much am I going to consume today versus how much am I going to consume tomorrow? Then that's two goods and I can model them against each other. And that's what I do in figure 16-6. I model consumption today versus consumption next year.

So here's my choices. As I said, if I consume everything today, I'm at the x-intercept at 80,000. I have 80,000 to consume today, nothing next year.

If I consume everything next year, what do I get? Well, let's say the interest rate is 10%. What that means is then I'll have \$88,000 next year. Why will I have more next year? Because by saving, I earn interest. By diverting my consumption to the future, I earn interest. At 10%, that means I would have \$88,000 next year.

So my budget constraint is the line with the slope minus 1 plus i . My budget constraint is the line with the slope minus 1 plus i . In other words, the price of consumption today in terms of consumption tomorrow is minus 1 plus i . OK, let me think about it. Let me say that again. It's really confusing.

The price of consuming today instead of consuming tomorrow, assuming no inflation-- so prices are the same in the market-- is minus 1 plus i . Think about that. I find it useful to think back to the labor case for parallel. In the labor case, what did we say was the price of leisure? What was the price of leisure? Someone raise their hand and tell me. In the labor-- yeah?

AUDIENCE: The wages.

JONATHAN GRUBER: The wages. Why?

GRUBER:

AUDIENCE: Just because that's the opportunity cost of not--

JONATHAN GRUBER: Right. So by that same logic, can tell me why is the price of consuming today 1 plus

GRUBER: i ?

AUDIENCE: Because if you choose to save, then we're effectively richer.

JONATHAN GRUBER: Exactly. The opportunity cost-- remember, we are an annoying discipline with a dismal science. We're telling you, hey, enjoy that cookie, but by the way, if you weren't eating that cookie, you could have 1 plus i cookies tomorrow. So just like we nag you for sitting around watching TV, we nag you for eating today by saying, hey, the more you consume today, the less you can have tomorrow.

And in fact, that trade-off is that for every cookie you consume today, you forgo 1 plus i cookies tomorrow. So that's the budget constraint. The slope is the opportunity cost of consuming today in terms of tomorrow's consumption or next year's consumption, which is 1 plus i . That's the slope of the budget constraint, is the

opportunity cost.

And then, then we say, OK, well, that's the opportunity cost. That's the budget constraint. Well, how do I decide? Well, then we know how to make these decisions, which is go to utility function. You can write down the utility function, which is a function of C_1 and C_2 .

Now, what is C ? C is all my pizza and cookies, but we're aggregating it up. Just like our utility function last time was a function of leisure and consumption-- we said consumption was the bundle of goods you eat and leisure is this thing. Now we're saying, OK, our utility function now is a function of this trade-off.

Now, you might say, wait a second. How can both those be utility functions? And the answer is you have some meta-utility function that includes consumption today, tomorrow, leisure, pizza, cookies, et cetera.

But we can think about this in sequential steps. First, we decide how we're going to split our income. Then we can decide what to spend it on each period. Then you can do a separate consumer maximization decision. But our first question is simply how am I going to split my income. Well, that's going to be a function of my taste for consumption in this period versus next period and the price the bank will pay me for delaying consumption till next period.

Now, what happens? Questions about that? Now, what happens in the scenario when the interest rate goes up? What do you think happens if the interest rate goes up? Yeah?

AUDIENCE: There's [INAUDIBLE].

JONATHAN GRUBER: Right. So what do you think you should-- what do you think will happen to your consumption pattern? Yeah?

AUDIENCE: You should spend less today.

JONATHAN GRUBER: Spend less today and save more because it's rewarded. And why is that not necessarily true? Yeah?

AUDIENCE: Because you might only need a certain amount of money to live. So you don't have

to save as much today because you'll make--

**JONATHAN
GRUBER:**

Because of what two effects? Income and substitution effects. You gave exactly the intuition that the substitution effect gives you. The substitution effect is exactly right. If the interest rate goes up, that's like the price of consumption today going up. And if the price of something goes up, the substitution effect says you do less of it.

But if interest rate goes up, you're richer. And if you're rich, you do more of everything, including consuming today. The income effect goes the other way. It's like labor. Once again, income and substitution effects is why we bothered telling you so. Because income and substitution effects, in these cases, go against each other.

Let's look at figure 16-7, OK? In figure 16-7, we start at point A. Now imagine the interest rate doubles to 20%. Now imagine the interest rate doubles. As you said, that pivots the budget constraint upwards. You could still consume only \$80,000 this year, but now for every dollar you save, you get \$1.20 next year.

That has two effects on your decision. The substitution effect, we get by drawing an imaginary budget constraint-- that's the dash line-- tangent to the original indifference curve but at the new slope. By definition, that means you consume less today. You consume less today by definition. If the price of something goes up, the substitution effect always says you do less of it. You consume less today, which means you'll save more.

Remember, savings is just income minus consumption in period one. So just as labor was 24 minus leisure-- and so if we just solve for leisure, we could get labor. Savings is just income minus consumption in period one. So if we solve for consumption in period one, we get savings. People see that?

So basically, the point here is the substitution effect says, well, gee, the price of consumption in period one just went up. It's more costly in terms of future consumption. I'm going to do less, but then my savings is going to go up. Substitution effect says you save more.

But the income effect says, wait a second. You're now richer. Every dollar of your

savings you are doing now yields twice as much in interest. If you're richer, you'll consume more of everything, including period one consumption. So the income effect takes you back the other way.

Now, whether the income effect dominates or not, we don't know. In this case, it doesn't dominate. In this case, you still, on net, end up consuming less in period one and saving more. But we don't know what's going to dominate.

And in fact, the evidence here is incredibly weak. I won't spend a long time on the evidence because it's not nearly as interesting and strong as labor supply. The evidence is incredibly weak even about the sign.

And let's come to the intuition that was given for why. Well, think about how people make savings decisions. Lots of people have savings goals. I want to have x by the time I retire. Typical way if you ask people about their savings-- if you ask them, they typically say I want to make sure I have x in the bank in case I'm in an accident. I want to make sure I have y by the time I retire.

Well, in those models, if the interest rate goes up, savings rates go down. Because after all, to hit a target with a higher interest rate, I can save less. So it's actually not that surprising that you'd have a higher interest rate leading to less savings. It's kind of intuitive, actually. If people have savings targets, a higher interest rate would lead to less savings because they can get to their target more easily.

So actually, we don't even know which way this goes. It's, I think, one of the great unsolved mysteries in economics empirically, is, once again, we typically assume-- and with a gun to my head, I would say it's probably true that higher interest rates leads to more savings. But the evidence on which that rests is pretty weak. And the key point for you is to understand it's uncertain and it depends on whether income and substitution effects dominate. Questions about that? OK.

So now let's step back and put it all together and think about you making your decision about life. You can think about your decisions about your life in three steps. Step one is you decide how hard to work. Step one is you decide, how much money do I want to make? Well, that's about maximizing utility over consumption and leisure.

Step two is, having decided how much you're going to make-- and that yields your labor. Step two is, deciding how much you're going to make, you decide, well, how do I want to spread that over time? How much do I want to consume today versus tomorrow? Well, that's about intertemporal choice. That's about deciding on C_1 versus C_2 , and that's going to yield your savings.

Step three is, now that I know how much I'm going to consume each period, now I want to maximize utility across all my goods I might want to consume-- x_2 , across all the goods I want to consume. That was our original cookies and pizza example. So you could think of it as a hierarchical set of consumer optimization problems that you're going to solve.

Now, you might say, well, gee, Jon, that's sort of confusing because, in fact, the interest rate and how much am I saving could determine how hard I work, right? Let's say the interest rate goes way up and I have a savings target. I have to work less hard to hit that savings target. And I'd say to you, good for you. Take more advanced economics. More advanced economics, we recognize this is one integrated whole and we allow these systems to affect each other.

But for here, just think of them as separable steps, independent steps. But in practice, I hope you can see the steps will be integrated and they'll affect each other. Think of it. If the price of a good you really want to buy goes up a lot, not only will you buy less of that good; you might save more to buy it and work harder. So you can imagine how these things are integrated. But for now, we'll keep them separable, OK?

Questions about that? OK. Next time, we're to come back and talk about all the interesting stuff in capital markets and how we make decisions about how much to save and things like that.