

[SQUEAKING]

[RUSTLING]

[CLICKING]

**JONATHAN
GRUBER:**

All right. Why don't we get started? Today, we're going to continue our discussion of producer theory. Once again, to remember to put this in context, the first few lectures were working consumer theory to help us derive a demand curve. Now we're working on producer theory to help us come up with a supply curve.

We started last time by talking about how producers profit maximizes, and the profit maximization implies cost minimization. Therefore, to maximize profits, you're going to want to produce as efficiently as possible.

And basically, to do that, we need to understand how your costs vary with your output. If you're going to produce at an efficient level, you need to understand how what your costs are going to vary with your level of production. So essentially, our goal of this lecture is to develop a cost curve-- develop a curve which tells you how the cost of your production varies with how much you produce. And that's what we're after in this lecture. OK?

So we're going to start with the short run and then turn to the long run. So start with developing a short-run cost curve and then turn to a long-run. OK? And to make this lecture sort of mathematically coherent, throughout the whole lecture, we'll work with our favorite functional form-- a production form of the function q equals the square root of L times K . Remember, firms produce goods-- little q -- using two inputs, labor and capital. Labor is a variable input. That means you can change it in the short run and the long run. Capital is a fixed input, which means that you can only change it in the long run. OK?

So in the short run, we're going to have two kinds of costs. We're going to have fixed costs, which are going to come from a fixed level of capital, and we're going to have variable costs, which are going to come from our labor. So we're going to have two kinds of costs: fixed costs-- they're fixed, because in the short run, you can't change

the level of capital-- and variable costs, which are the costs of our labor. And then we're going to have total costs. They're simply going to be fixed costs plus variable costs. OK?

So that's how we think about costs. The costs of a firm's production in the short run is the sum of their fixed costs-- i.e., their capital costs-- and the variable costs-- i.e., their labor costs. Now, we're going to show you how you can turn a production function into a cost function. And to do so, you simply need to recognize that cost-- the costs of firms' production-- are simply the amount of capital that uses k times the price of that capital, which we'll call r , plus the amount of labor it uses, times the price of that labor, which we'll call W , the wage. OK?

This is the easy part. Think of the amount of hours of work you use times the wage per hour, or the amount of workers times the salary per year. In any case, this is easier to understand. Every additional unit of labor comes with a cost that is the wage of that unit of labor. Think about an hourly model. Every hour you work at your convenience store, they have to pay you the minimum wage for that hour. OK? That's the cost of that hour of labor.

Capital is harder. We call r the rental rate. And the reason is because we don't think of buying capital. Don't think of buying a machine. Think of renting a machine. The reason we do that is to make the periodicity work. OK?

You don't buy a worker, thank god. You rent that worker. And you rent that worker at a price, w . So when the firm uses your time, they're renting your time an hour at a time at a price, w . When we get machines, think of a firm as renting machines for a price, r per machine per time period. OK?

So I understand firms usually don't do this. They usually buy machines. And we'll come back to even if you buy a machine how it effectively is like renting it, but for ease of thinking about this, you want to think about flows, not stocks. Think about the firm's decision as renting a worker at the price, w , or renting a machine at the price, r . Yeah.

AUDIENCE:

Kind of like also we consider the gas and electricity a machine can use in productivity?

JONATHAN
GRUBER:

All that would be in there, and we'll come back to that. That's right. This is sort of the per period cost of using a machine. This is the per period cost of using a worker, is the wage. This is the per period cost of using a machine, which will include all the costs of running the machine as well as the costs of renting the machine itself. OK?

So later, we'll talk about how to own the machine. And we'll come back to the fact that you can actually use r as a representation. It's not bad. But for now, just think of renting a machine. Or if it's a building, think of this as the rent you pay on that building, OK? Not the cost to build the building.

Now, armed with our production function-- and let's also say, to make life easy-- I did this wrong on my notes. So I hope the other teachers figured it out. To make life easy, let's say that the rental rate, r , we're going to say is \$10, and the wage rate, w , is going to be \$5. Now, armed with a production function, if you simply have the production function and these two prices, you can derive the short-run cost function.

How do you do that? Well, just look at the math. We know that q equals square root of L times k bar in the short-run. L times k bar. OK? So inverting that simply means that L equals q squared over k bar. L equals q squared over k bar. OK?

And that means that cost can be written as $10 k$ bar-- the price the rental rate is 10, we have k bar amount of capital-- plus $5q$ squared over k bar. That's our cost function. I just plugged in for L and multiplied by the wage rate, w , which is 5. OK?

So for example, for a fixed level of capital, this is cost. So for example, let's imagine our short-run level of capital is 1. Let's imagine there's 1 unit of capital in the short run, just to make the math easy. Then that simply says that cost equals 10 plus $5q$ squared. And that's our cost function. 10 plus $5q$ squared. We've just derived the short-run cost function.

10 is the fixed cost component. That doesn't vary with the amount you produce. So there's no q part of this. $5q$ squared is the variable component. That varies with how much you produce. So the cost function is a fixed part, which comes from that one fixed unit of capital, and the varying part, which comes from the fact that the amount of q drives the amount of labor we need. OK? Questions about that math?

All right. Now, armed with this cost function, let's write that down again here. So C

equals 10 plus 5q squared. That's going to be our short-run cost function we're going to work with. And remember, that short-run cost function came directly from that production function. To derive this equation, all I needed was that production function and those two prices, and I derived it.

OK. Armed with that, we can define some of the key concepts that will drive our entire analysis of firms. And the single most important concept is marginal cost, which is what it sounds like-- the derivative costs with respect to quantity. OK? So in this, the marginal cost is $\Delta c / \Delta q$. OK? That's marginal cost.

OK. We'll also care about average cost, which is just c over q . It's very important in this class to keep our marginals separate from our averages. The average is simply over the entire range of production. What is the average cost to produce each unit? The marginal is what's the cost of producing the next unit? And since production functions are nonlinear, those will not be the same, generally. OK. Average will not equal marginal in general, because a nonlinear function $\Delta c / \Delta q$ is not the same as c/q . OK?

So we can actually graph these in figure 6-1. Figure 6-1 shows the cost curves for this cost function I just wrote down, which comes from that production function. OK? So you can see that the marginal cost, as I said, is $\Delta c / \Delta q$. Well, that's $10q$.

So the marginal cost, the cost of producing the next unit, rises with the number of units. Which makes sense. The cost has a q squared term in it, so obviously, the marginal cost is going to have a q term in it. So basically, the more you produce, the higher your marginal cost. The more unit you need to produce, the more the little q , the higher your marginal cost.

Average cost is this sort of funky shaped thing, where it's-- which is 10 over q plus $5q$ -- I just divided this by q -- where it's first declining and then increasing. Why is that? Why is average cost first declining and then increasing?

We've seen what-- just intuitively? Why is that? Why in general in the short run would we expect that? Average cost first to fall and then increase. Anyone have ideas? Yeah.

AUDIENCE: [INAUDIBLE] start up [INAUDIBLE].

JONATHAN Well-- no, but it's falling first. So why is it falling first? It's about that-- yeah.

GRUBER:

AUDIENCE: They have a really high average-- or first fixed cost is not [INAUDIBLE] the more you make, the lower that is.

JONATHAN Right. The first units are paying off your fixed costs if you think about it. Well, the
GRUBER: first unit you sell, basically you start with this huge fixed cost. So actually, by selling two units, yes, you get the variable cost, second unit, but you get to pay off the fixed cost, the first unit.

So look at here on this graph. We show average fixed costs and average variable costs. Average fixed costs are $10/q$. If you only produce one unit, your average fixed cost is \$10. To produce two units it's \$5. With every unit you produce, your average fixed cost is falling. You're paying off that fixed cost.

Average variable cost rises. Every unit you produce, you're getting more and more variable cost. You put those together, and you get a function that first declines and then rises. You first pay off your fixed costs, so your average costs are falling, then your marginal cost-- then you start to rise, because you've got marginal costs that increase with quantity produced.

And critically, the marginal cost intersects the average cost curve at the minimum of average cost. And that's just mathematical. If you have any function, and you take the average, then the minimum is going to be a derivative that basically-- before you get to here, 1.5 units, before 1.5 units, average cost is above marginal cost, because you're paying off your fixed costs.

Once you get beyond 1.5 units, average cost is below marginal cost. So average cost hits marginal cost at the minimum of average cost. Yeah.

AUDIENCE: In a relatively large company, eventually doesn't really worry about their fixed cost, because if they're having a lot more workers, their entire cost is going to be considered basically--

JONATHAN In the short run. It depends on this function. You said large, but large can be defined
GRUBER: in two ways. But absolute true. Certainly, if we take this very function, for large

enough q 's average fixed cost asymptotes to 0. 10 over infinity is 0.

So for large enough q 's, the average fixed cost goes away-- in the short run. Remember, we're in the short run with these fixed costs. OK. Now, so other questions. Good question. Other questions about that?

So that's our basic intuition, the short run-- is that at first our costs are super high, because you've got to pay-- you had to build the plant. But then over time, that plant cost falls, and then your only costs is basically the fact you've got to hire more workers if you want to produce more.

Now, what we want to notice is that in the short run there is a really close relationship-- one was the key relationship between marginal cost and the marginal product of labor, which we defined last time. Remember, the marginal product of labor was dq/dL . How much-- remember digging a hole and diminishing marginal product? That each additional worker, for a fixed level of capital, is less and less productive, right? We talked about that last time.

Well, the marginal cost of production is, as I said, equal to $\Delta c / \Delta q$. OK? Well, we know from last time that $\Delta q / \Delta L$ we defined as the marginal product of labor. Plugging those and-- so we know marginal cost is $\Delta c / \Delta q$.

And we can write this-- if you take the derivative of the cost function, so our general cost function up there-- see the cost function at the top there? Take the derivative of that with respect to the amount of labor. Well, the first term drops out, because it's fixed.

So you can rewrite $\Delta c / \Delta q$ as w times $\Delta L / \Delta q$ -- w times $\Delta L / \Delta q$. Right. I just rewrote $\Delta c / \Delta q$ as w times $\Delta L / \Delta q$ -- god. Brutal. Sorry, guys. w times $\Delta L / \Delta q$. OK. A little bit better.

OK. I read that, because I just took the derivative of the cost function. First term drops out, we take the derivative. Second term, I just took the derivative here, or the discrete derivative, and I just said, $\Delta c / \Delta q$ is w times $\Delta L / \Delta q$.

Well, we know the marginal product is $\Delta q / \Delta L$. So we can rewrite marginal cost as w over the marginal product of labor. The marginal cost is equal to w over the marginal product of labor. And that makes sense.

The marginal cost of the next unit will be higher the higher the wage and lower the more productive the worker. Making another unit with a super productive worker is cheap. Making another unit with a very unproductive worker is expensive. So essentially, the more you pay them for each hour of work, the higher your marginal cost. But the less they get done each hour of work, the higher your marginal cost.

So that's why, roughly speaking, firms might want to pay a lot to people who are high skilled. So you might say, gee, why are they paying my friend twice what I'm paid? Well, maybe your friend's twice as productive as you. That would make-- or two and a half times as productive as you. And that would make sense.

So we can't just say that it's a mistake to pay someone higher wages. We have to consider their wages relative to how productive they are. And that's a key relationship we'll come back to. OK? Question about that?

All right. That's the short run. Now let's go to the long run, which gets a little more interesting. Actually, let's write in that side. I'll switch to this side today. Switch things up.

This way, this side, I block you less, probably. So I should do this side. OK. Long run cost curves. OK. Long run cost.

Now, here what gets interesting, is now K is no longer fixed. Now we get to choose our input mix. And now our goal is going to be, how do we choose our input mix to minimize costs? That's our goal here. How we going to choose the mix of workers and machines to produce a given quantity most efficiently?

Now, that optimal mix may change with the quantity. So we're going to start. We're going to do this in two steps. First, we're going to say, for a given quantity picked out of a hat, what's the right mix of labor and capital that minimizes the cost of producing that quantity, given our production function? Then we're going to say, as the quantity varies how does that change the optimal mix of L and K ? And does it?

So two steps. First say, for a given quantity, what's the right L and K to minimize costs of producing that quantity? Then ask, well, as we vary the quantity how does L - how do L and K vary optimally? OK.

So basically, we want to find the economically efficient combination of L and K which is a combination that produces goods at minimum cost. And to do this, we are going to write down-- to derive this we're going to write down what we call isocost curves. Isocost curves.

Remember last time we did isoquants, which felt a lot like the indifference curves? Isocosts are going to feel a lot like budget constraints. Isocost curves are essentially the firm's budget constraint. They're essentially mappings of the function c equals wL plus rK -- essentially, mappings for different amounts of K and L of the function c equals rK plus wL .

So if you look at figure 6-2, here we see our isocost lines. OK. So let's talk about this for a second. So for example, take the middle one, the \$100 isocost. This is saying, what combinations of labor and capital cost you \$100?

Well, with a rental rate of \$10 and a wage of \$5, that means you can have 10 machines and no workers or 20 workers and no machines, or some combination in between. This is just a budget constraint. It's just saying, given the amount of money you want to spend, given your cost, how many machines and workers can you have?

The difference is, we don't start-- I didn't start this example by saying, your parents give you x dollars. That's why firm theory is harder than consumer theory. I pin down the consumer theory problem much more easily by saying, your parents give you x dollars, which told you which line to derive-- graph.

I don't have that here. I haven't told you that here. So you have to graph a series of isocost curves, because you don't know what the optimal cost is going to be. That's to be pinned down later. That's what makes supply theory harder. You have to draw the series.

So you draw these series of isocost curves-- different combinations that represent different amounts, different totals of cost. And of course, the slope of that isocost curve is $\Delta K / \Delta L$, or minus w over r . That's the slope. Or in this case, minus 0.5.

Now, those of you thinking ahead-- I know you guys are very insightful as a class, I'm

sure many of you are thinking ahead-- might think, gee, that slope might change as the number of workers and machines change. Could you imagine the relative price of capital labor changes in different costs-- and it might, and we'll put that aside for now.

For now, assuming for every relevant quantity these prices \$5 and \$10 are fixed-- let's ignore where those prices come from. They're just given now. We'll come back to that later. Like I said, this course is sort of like peeling an onion. We raise things, then we come back go to the next layer. Where'd that come from?

Right now-- we'll tell you where w and r come from. Right now we're just going to take them fixed. And we'll assume they're always \$5 and \$10, regardless of the amount produced. OK.

Now, here's the question. You're a firm that wants to release a certain amount of units. You have a production function and a cost function. How do you graphically figure out the right combination of capital labor to use to produce a certain amount of units? Yeah.

AUDIENCE: Could it be the tension between the isoquant and the isocost curves?

JONATHAN GRUBER: That's exactly right. Just as I asked you, what is the right combination of pizza and cookies, and you told me that it was the tangency of the indifference curve and the budget constraint, it's the exact same logic here. The optimal mix of capital and labor comes from the tangency of the isoquant with the isocost, as we see in figure 6-3.

And ignore the, like-- somehow that curve is sort of connected at the top. It's sort of a glitch a PowerPoint. Just ignore that. It's not actually like a square or a trapezoid. It's just a curve. What would you call that-- the curve and the two sides. Is that a-- that's not a trapezoid. That's not-- it's not a polygon. It's just a line.

All right. So basically-- is there a name for that? A curve and two lines? I don't think so. It's just a polygon, right? OK. So it's not a polygon, it's just a curve.

So the curve is the isoquant for the square root of 12.5. What do I mean by that? I mean that is the combination of capital and labor that delivers square root of 12.5 units of production. So what that curve is all possible combinations of capital and

labor that deliver square root of 12.5 units. Just like the indifference curve is all possible combinations of pizza and cookies that leaves you equally happy, this is all possible combinations of capital and labor that leads you to a given production level.

And as we said, the further out the isoquant, the more you can produce. So you want to produce as much as you can given the prices you face in the market. Well, those prices you face in the market are delivered by the isocost curve. So the tangency is the best-- is the cost minimizing point. That's when you're producing the most you can given the costs you face in the market-- the most you can, given the costs you face in the market.

And that tangency condition-- once again, considering our parallels to consumer theory, the tangency condition is going to deliver that the marginal product-- is going to deliver that the marginal product of labor over the marginal product of capital, which we remember called last time the marginal rate of technical substitution, is going to be equal to w/r -- actually, the negative of these is going to be equal to each other.

But we'll just cross out the negatives. So the negative of MPL over MPK, which we called the marginal rate of technical substitution, is equal to the negative of w/r . The slope, the optimal point, is where the marginal rate of technical subsection equals the slope, which is the wage to rental rate ratio.

Alternatively-- I don't know if anyone besides me likes this intuition-- we can rewrite this as MPL/w equals MPK/r , my bang-for-the-buck formulation that I like-- that the next dollar of wages, if you ask, should I spend next dollar on wages or machines, you should do it until the next dollar of wages delivers you the same return as the next dollar of machines. This is what you get for the next dollar of wages, MPL/w . This it what you get for the next dollar of machines, MPK/r . You want to continue to trade off machines and workers until that condition is true.

So let's actually now solve for this for our example. Let's actually solve for that. If we solve for this, we know that the marginal product of labor, which is dq/dL , is 0.5 times K over square root of K times L . And the marginal product of capital, which is dq/dK , equals 0.5 times L over the square root-- 0.5 times L over the square root of

K times L. Just taking the derivative of the-- all I did was take the derivative of the production function.

So therefore, putting these together, we're going to-- we know that marginal rate of technical substitution in this example is equal to minus K over L. That's not a general formula. That's just this example. The marginal rate of technical substitution is equal to the negative of the ratio of capital to labor.

We also know that the wage rate-- we also know we want to set this equal to the negative of the wage rate-- I'm sorry, the wage rental rate ratio. And we know that's negative 1/2. We know that's a 1/2, because that's 5 and that's 10.

So we want to set the marginal rate of technical substitution to the wage rental rate ratio, which means we set minus K over L equal to minus 1/2. Or at the optimum, that means that your labor, your capital, at the optimum-- that means the amount of capital-- should be half as much as the amount of labor. In this example, we just solved for the efficient combination of inputs, the efficient combination is you should use capital to labor in a ratio of 1/2. You should use half as much capital as use labor.

So let me pause there. And let's talk about where this is coming from. Yeah.

AUDIENCE:

So does that mean that at any given price of cost [INAUDIBLE] line, that is the optimal point where it will be tangent [INAUDIBLE]

JONATHAN

Exactly. Exactly. That's the graphic intuition. Let's come to the economics intuition.

GRUBER:

The economics intuition is the following. The production function delivers this relationship-- that the marginal rate of technical substitution was minus K over L. In other words, when you're producing goods, given this production function you're indifferent between the next machine and the next worker.

That's just the way this production function worked out-- that one more machine delivers you the same amount as one more worker. Now, I've just told you that one machine costs half of one more worker. So which you want more of? No? One machine delivers the same return as one worker.

You want more workers. Workers cost half machines, they're equally productive, so you want more workers. So the optimal amount of machines is going to be half as

many as the number of workers.

You want more workers, because you're indifferent-- look at that production function. You're indifferent. You don't give a shit about L versus K. They're the same to you. You're a hard capitalist, man. Machine or worker, you don't care.

But the market's telling you you can get a worker for half the price of a machine. So you, as a good cost minimizing capitalist, take twice as many workers as machines. And that's the outcome that you get here. Questions about that?

OK. So that's basically what we do to derive this. Now, what I want to do is take this and then derive our ultimate goal, which is, what is the long run cost function? That's sort of what we-- why we started this lecture. What is the long run cost function?

Let's do the math. We'll do the math in five steps. Step one, q equals square root of K times L . Step two, we know from up there that K/L equals w/r . We derived that, leading us to the conclusion that K -- lead us to the conclusion that K equals $1/2 L$. We just derived that.

Therefore, we can rewrite q as the square root of $1/2$ times L squared, just substituting it, because K equals $1/2 L$. Therefore, we can solve for L is going to be square root of 2 over q . And K is going to be square root of 2 over 2 over q .

L is square root of 2 over q , K is square root of 2 over 2 over q . I'm sorry-- no, not over q . I'm sorry. That's my bad. Error. Error. Go back. Should always look at my notes.

It's square root of 2 times q . My bad. L is square root of 2 times q . And K is square root of 2 over 2 times q . OK. Therefore, armed with this L and K , we can rewrite our cost function.

So step five is that the cost function-- given this stuff, the cost function equals r times square root of 2 times q -- I'm sorry, r times square root of 2 over 2 times q -- plus w times square root of 2 times q . I just plugged in the optimal L and K into my cost function.

Now I can plug in the 10 and the 5 to get C equals 10 times square root of 2 times q .

And I'm done. I just derived the cost function. That's what we came here for. This is what you got up this morning and wanted to see.

You got up this morning, you said, I want to know how does the cost of a firm vary with the quantity it produces? And I've just told you. This tells you how the costs that you pay vary the quantity you produce.

And I did that by deriving the optimal mix of L and K you want to use, and then simply imposing the prices of those two, and I get a cost function. Yeah.

AUDIENCE: Wouldn't you be adding the two terms?

JONATHAN I'm sorry?

GRUBER:

AUDIENCE: Wouldn't you be adding the two terms?

JONATHAN Which two terms? Oh, I see. Yeah, plus. I'm sorry. You're right. My bad. That's a plus.

GRUBER: Thanks. This is the most math I'll do in a lecture all year, you'll be pleased to know. It's why it's my least favorite lecture. Yeah?

AUDIENCE: So is five generally true?

JONATHAN You mean this particular functional form?

GRUBER:

AUDIENCE: Yeah.

JONATHAN No. This is all dependent on that production function I wrote down. What's generally true is this-- or actually, K wouldn't be-- what's generally true is just C equals-- in the long run, what's generally true is C equals wL plus rK . That's what's generally true. I just made--

But what I've showed you is, given three things-- a production-- all I gave you was a production function, a wage rate, and a rental rate. Given those three things, you can then derive the cost function. Given those three things, you can derive the cost function.

In fact, given two things you can derive the cost function. You could actually derive

the cost function given one thing-- given just the production function, derive the cost function as a function of these input prices. OK. So that's a lot of results from one function, from one production function.

Now, other question about this? Other math I got wrong? Sorry about that. Yeah.

AUDIENCE: Sorry, could you just repeat the three inputs that you were using?

JONATHAN GRUBER: [INAUDIBLE] all I used. Somebody tell me. What do you need to get the magical cost function? What three things do you need? Yeah.

AUDIENCE: w and r and then q.

JONATHAN GRUBER: No, w, r, and-- what about q? You need q, but what about-- what do you need? w, r, and the production function. So armed with the production function, that mathematical equation, this mathematical equation, w and r, I'm done. Everything I've done in this lecture comes from those three things.

The math is hard and annoying. We will have you practice it. You will not like it. OK. It's just kind of what you've got to do. All right? I don't like it, you're not going to like it. It's just what we've got to do to get to the more interesting stuff.

OK? Yeah.

AUDIENCE: Is the bar on the K?

JONATHAN GRUBER: It's fixed in the short run. The bar on the-- shouldn't see a bar on a K over here. That's all over here when I was doing the short run. Yeah?

AUDIENCE: Can you use r and w in order to get to the fifth step?

JONATHAN GRUBER: Yeah, we needed r and w.

AUDIENCE: Well, in the sense that, to simplify, K equals $1/2 L$?

JONATHAN GRUBER: To K-- oh, you're right. That's a good point. I needed r and w back here. You're right. That's a good point. Good point. But I still could have done this whole thing as a function of r and w if I wanted to-- if I wanted to really screw up my math. All right?

OK. So now, armed with this, let's talk about what happens when input prices change. We talked about with consumer theory, what happens when the price of pizza and cookies change. What happens when the price of labor and capital changes? What does that do?

So let's talk about changes in input prices. OK. Let's go to figure 6-4. And let's look at, with the same production function, square root of L times K-- we're not changing our production function-- we're going to change the wage rental ratio.

So line-- we have our initial line, our initial wage rental ratio, which is that basically you have a wage rate, but the budget constraint, essentially, that's flatter is our original budget constraint. The flatter budget constraint is our original budget constraint. That's the budget constraint with the price of capital of \$10 and a wage of \$5. And that intersects our isoquant at point x. So we chose five units of labor.

Now we have a new-- we chose five units of labor and two and a half machines. That was our original. This is sort of a messed up graph. But our original intersection was at point x. The cost minimizing combination, the square root of 12.5 production, was to have five workers and two and a half machines.

Now let's say the price of workers rises to \$10. The wage rate rises to \$10 an hour. So now, workers and machines cost the same. What is now the optimal mix of workers and machines? Well, graphically we know we still want to produce the square root of 12.5. So we want to stay tangent to the same isoquant.

So based on-- what we're saying is, this is as if we said to consumers, keep your utility changed, change the price. What do we call that? Remember what we called that? Keeping utility constant, changing the price? Anyone remember what we call that?

Who said that? All right. Raise your hand next. Be proud. Substitution effect. That's the substitution effect.

It's the same idea here. We want to know, for a given level of production, what happens as the price of the inputs change? And so we shift along the isoquant for point x to point y. And you'll see we choose a mix where we use fewer workers and more machines.

And just as the substitution effect is always nonpositive, this shift as the wage rate rises, the price of the good, the x-axis rises. You will unambiguously use no more, and almost certainly less, workers. OK.

And you can see that graphically-- think about graphically, you're looking for the tangency between this curve and the line. The slope of the line just got steeper. Therefore, you must move to the left on the curve. It's the same proof as we used the substitution effect, where substitution effect was always nonpositive.

OK. It's the same intuition here we use for y . A rise in the wage rate will lead you to hire fewer workers and more machines. Guess what? You just entered the debate on the minimum wage.

And if you follow the debate in minimum wage, what do people say? Well, if you raise the wage you have to pay workers, they'll be replaced by machines. That's this. This is the math-- this is the mathematical and graphical intuition behind the debate on minimum wage, which we'll get into later in the semester.

But the basic idea of that debate is, gee, if you force firms to pay more to workers, they're going to substitute towards machines. That's exactly right, in theory. And practice, there's a lot of complications. But this gives you the theory of why people make that argument. Yeah.

AUDIENCE: So in this example, only the wage for workers [INAUDIBLE] not the machines.

JONATHAN Not the machine.

GRUBER:

AUDIENCE: So why does the isocost not have the same y-intercept?

JONATHAN Ah, great point. Because here I'm drawing a new-- I am drawing the isocost that I would use while still producing square root of 12.5. So that's it's just the substitution effect.

I'm not drawing the full set of isocosts at the new price. I'm just saying, to produce the same amount what's my new-- if I want to produce the same amount, what combination do I now have to use? OK?

Yeah.

AUDIENCE: The total cost of production [INAUDIBLE].

JONATHAN GRUBER: The total cost of production-- let's see. Yeah, it has to be-- no. Let's see. No, total cost doesn't have to be the same. The total cost used to be five workers at \$5 an hour, that's \$25. So it used to be \$50.

Now what is it? Now it's \$70. The total cost has gone up.

AUDIENCE: So it's not like the budget constraint or the--

JONATHAN GRUBER: Exactly. It's not like the budget constraint where your income is fixed. That's what's hard about producer theory. Because basically, the budget constraint was sort of asking, keeping your budget fixed. This is like asking, keeping your total production fixed.

And so now you have to pay more to get that level of production. OK? Good questions. OK.

Now, ultimately what does this lead to? So that gives us our change in input prices. Other questions about that?

Now, remember I said at the beginning of the lecture, we are first going to solve for what is the cost minimizing combination of inputs for a given quantity? We derived that up there. It's half as much-- going back to our old prices, it's half as much capital as labor.

Now we want to ask, how does your cost change as the quantity changes? And we call that the long run expansion path. The long run expansion path, which is, how do your costs expand as you produce more?

And we see that in figure 6-5. In figure 6-5, we show the particular case of a linear long run expansion path. That's what you get in this example. It's a particular case.

What this case says is, at any given level of production, the optimal mix of labor and capital is the same. In other words, you always want to have-- essentially, given the price of labor is half the price of capital, you always want to have twice as many workers as machines.

So if you want to produce square root of 12.5, you want five workers and two and a half machines. If you want to produce square root of 50, you want 10 workers and 5 machines. If you want to produce square root of 112.5, you want 15 workers and 7 and 1/2 machines.

So the long run-- given this production function, the long run expansion path is linear. You always want the same ratio of workers to machines. Yeah.

AUDIENCE: [INAUDIBLE] consumer and firms, is the reason why we don't necessarily have a strict budget, per se, and then isn't the idea that if we really want increase production, we can take a loan out?

JONATHAN GRUBER: This is what I tried to say. It's sort of-- I always say it over here, but it's hard, and we have to come back to it. The reason producer theory is harder is because we're not given a fixed constant we are with consumers.

Consumers, we're saying, look-- you've got a resource, you've got to constrain maximization. We haven't constrained the maximization yet. There's another constraint we need. They have an extra degree of freedom relative to consumers.

Now, in fact, consumers have degree of freedom, too. When you grow up, your parents don't give you money. You decide how much to make. So in reality, consumers-- you can do this-- will have the same degrees of freedom.

But we started with the easy consumer theory case, where you constrict-- we took away a degree of freedom. Now we're writing it back, which is, you can choose how much to produce. Like, you being able choose your income as a consumer. That leads to long run expansion path.

Let me go on, because I want to make sure I get through this stuff. OK? Now, the long run expansion path does not have to be linear. So think about-- look at figure 6-5b and 6-5c. So 6-5b is a long run expansion path for a production function such that capital becomes less productive the more you produce. I don't have the example of production function.

But when you write down production functions which have the feature that the more you produce the less productive capital becomes, the less each additional unit

capital helps [INAUDIBLE] additional unit of workers. So you know, we could think of this roughly as sort of like a fast food restaurant. That kind of-- you know, each addition-- that basically, there's so much stuff to do where workers can efficiently share tasks and things. Each additional worker-- that the marginal product of labor essentially diminishes less quickly than the marginal product of capital.

On the other hand, in figure 6-5c we can have a long run expansion path where labor becomes less productive relative to capital. Think of it as like heavy machinery, where basically all workers can do is run the machine. So that second worker-- workers don't really do much but sit there and flip a switch.

You need the worker to flip the switch. That's all they do. So the second worker, you're already flipping the switch. So really, adding more machines is a more productive way to expand the output.

None of these is right or wrong. We're just saying that the shape of this expansion path can basically vary with how much-- with different production functions. But they're all the same idea.

Now, so basically, that tells you-- but here's the bottom line that we wanted to come to. That long run expansion path is a long run cost curve. So ultimately, if you want to ask, how do my costs vary with how much I produce, this curve tells you.

Because what it does, it says, for every level of production I'll tell you the optimal combination of L and K. Given the price of L and K, that will tell you the costs. And so you trace out the costs with every level of production. This is your cost curve.

This long run expansion path tells you what the costs are for every level of production. And it tells you that, because you've made-- you're doing the efficient level of production. That's what the long run expansion path is telling you. OK. This is hard. I'm about to make it harder.

Which is, we're now going to talk about the relationship between short run costs and long run costs. And the key insight is that long run costs are everywhere lower than short run costs. Without looking at the figure-- because the figure doesn't help with this-- why does that make sense? Why our long run costs-- why, if you can optimize over the long run, will you always have costs that are no higher, and in general

lower, than optimizing the short run? Yeah.

AUDIENCE: [INAUDIBLE] already had the right capital.

JONATHAN GRUBER: Because you have an extra degree of freedom. I think that's LeChatelier's Principle. Is that right, for the chemists among us? That basically, like-- essentially, an extra degree of freedom means, the more you can optimize over, the better you can do in optimizing.

In the short run you're constrained by the size of the building. In the long run, you could choose. So let's-- to see that, let's go to figure 6-6. This is a confusing figure. So bear with me as I walk you through it. OK?

Consider a firm with three possible sizes of plants. They're going to build a plant. So the capital here is the building. And there's three possible sizes-- small, medium, and large. The small plant has the curve SRAC1. What does that curve mean?

That means that the small plant-- I'm sorry, the small plant is SRAC2, the medium plant is SRAC2, and the large plant has SRAC3. Compare SRAC1 to SRAC3. What this is saying is, for small quantities of production, SRAC1 lies below SRAC3. For small quantities production, if you extend SRAC3 out, you see at levels of production like q_1 or even q_2 , SRAC3 is way above SRAC1.

When you go to a level of production like q_3 , SRAC1, if you extend that dashed line out, is going to be much, much higher than SRAC3. So the right-- and SRAC2 is in between. So essentially, for different levels of production these give the different optimal short run cost curves.

In the long run, you get to choose. So the long run average cost curve is the lower envelope of the short run average cost curves. Because in the long run you say, well, here's my production level. I know in the long run-- so if I know I'm going to build a lot of things, I choose SRAC3. I choose the biggest plant. If I know my production is going to be low, I choose the smallest plant.

But I can optimize in the long run by choosing the right sized plant for my production level. This is hard. And I'm almost out of time, so let me end with an example that perfectly illustrates this. Tesla. Elon Musk. Everybody's favorite guy these days.

Tesla, when they came out, had to decide how big a plant to build-- how many batteries to make. Batteries are the key [INAUDIBLE] Teslas. And they expected to make-- to have demand for 20,000 cars by the year 2017.

So they built a plant like SRAC1. They built a plant that was the efficient plant to produce 20,000 cars. The problem is, demand was for 200,000 cars. And as a result, there's a three-year waiting list to get Teslas. It turned out that was not the right size to produce. They lost money-- relative to the optimum. They made money. Musk is incredibly rich. But they didn't do what was most efficient given the underlying demand.

But now, Musk can re-optimize. Now he's saying, wait a second. People want way more cars. Well, producing them at the tiny plant was exorbitantly expensive. I had to run it over and overtime-- pay workers overtime. To produce 200k cars in that tiny plant just was exorbitantly expensive. That's if you take that dashed line and extend it way the hell up. The SRAC1 extended way the hell up, incredibly expensive.

So what is Musk doing now? Building the largest battery plant in the world. In Nevada, he is building a battery plant that can produce batteries for 500,000 cars. So he shifted from SRAC1 to SRAC3. He's now saying in the long run, I can more-- if I'm going to produce 200,000 cars, I can do that more efficiently with a giant battery plant.

And that's what he's doing. So he's re-optimizing. Now, what if Musk is wrong? What if it turns out Teslas suck and people are like, I don't want them anymore? Someone else-- or, you know, Chevy finally figures it out and makes a good electric car.

Then what's going to happen is he's going to have made a mistake in the long run. Then the third period, he'll go back to a smaller plant again. But he always can do what's efficient in the long run, given the underlying demand. So Tesla is an example of this sort of long run, short run dichotomy.

Anyway, it's a lot of stuff for one lecture. We'll come back next time, talk more about costs. And then we'll start getting into competition.