[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So today, we're going to continue, once again, the roadmap. We talked about consumer theory. Then we started talking about producer theory. And we got to the point where we realized the one place producer theory is harder than consumer theory is you're not given a budget constraint.

You're not told how much you have. You get to choose how much you have. You get to choose how much to produce. And we know for any quantity the way you'll produce it. We know the efficient combination of inputs to produce any quantity. We developed that last time.

What we don't know is, how do you choose what quantity to produce? And to get that, we're going to add a third-- instead of having two equations and two unknowns, we're going to add a third element to the model. And the element we're going to add is market structure. We're going to add an element of how firms compete in the market, OK?

And essentially, that is where we're going to go for the next several lectures is to think about the structure of the market itself and how that feeds back into firms' decisions on how much to produce. And so that's kind of where we're going today is to ask, how do we take the market, impose it on the firm's decisions from last time on that long run expansion path we developed last time to determine how much they actually produce?

We're going to start with a particular assumption. We're going to start by talking about short run profit maximization. And we're going to start in the particular example of a market that features perfect competition. I'll define this more technically in a minute. But basically, stepping back, it's going to turn out the firm's decisions on how much to produce is going to depend on the kind of market they face.

And there's different types of markets. One extreme that we'll focus on for much of this class is the extreme of perfect competition-- basically, many firms competing to produce identical goods. The other extreme would be monopoly-- one firm that produces the good. And then there's a whole host of examples in between. Indeed, in between is my favorite word in economics-- oligopoly-- just a great-sounding word.

Oligopoly is really the way markets almost all look, which is a number of firms competing, but not too many. So it says we have one extreme where there's a huge numbers of firms competing. That's perfect competition. The other extreme is just one firm in the market. That's monopoly. And in between, we have oligopoly, which is most industries.

Think of the auto industry. There's more than one car company, but there aren't 1,000 car companies. Think about most of the industries that we feature in a modern developed economy are oligopolistic. And that'll be the third place we turn.

So we're going to start by thinking about perfect competition. Then we'll turn to monopoly. Oh, as always, it's easiest to trace out the extremes. Start with one extreme, which is perfect competition. We'll then shift to the other extreme, which is monopoly. You'll then take your midterm, and then we'll talk about the middle case, which is oligopoly.

So let's talk about what are the assumptions of perfect competition? What are the assumptions that underlie perfect competition? Basically, the key-- the technical definition of a perfectly competitive market-- is one where all firms are price-takers, price-takers. What do I mean by that?

What I mean by that is that no action that any given firm does can affect the price in the market. Any given firm is sufficiently small relative to the market that no action on the individual firm takes affects the market price. Well, when would this be true?

This would be true in the case where when demand for a firm's output is perfectly elastic. Demand for a firm's output-- not the market demand-- but demand for a given firm is perfectly elastic. And we show that in Figure 7-1. In Figure 7-1, here is a case-- and notice the x-axis is little q, not big Q. Little q is a firm, big Q is the market. This is a firm-level diagram of the firm's production decision, how much to produce at what price.

In a perfectly competitive market, the firm faces a perfectly elastic demand at some price, p. We'll talk about where p comes from. But just by the end of the lecture, you'll know where p comes from. But now just say there's some price, p. And at that price, demand is perfectly elastic. As a result, supply-- shifts in supply have no impact on the price charts. You always charge the same price.

So when does it make sense that demand be perfectly elastic? What are the assumptions of perfect competition? When does it make sense that firms are price-takers? Well, we make three key assumptions. The first is identical products. The first assumption is that firms are selling identical products.

The second assumption is that there's full-- transparent, full knowledge of prices. That is, every consumer knows every firm's price. Identical products, every consumer knows every firm's price, and finally, no transaction costs-- the third assumption, no transaction costs meaning you can easily shop at any firm you want. The products are identical. You know every firm's price, and you can easily shop at as many firms as you want.

Now this is obviously untrue for any market in the world. But once again, the extremes provide a useful example. A close example would be eBay. Think about eBay for a very common good, buying a certain type of silicon chip on eBay, or a certain type of kind of food, or a certain type of a shirt of a certain size and style.

The products are identical. The prices are all listed there, and it's easy to shop. There's no trick, just look at the list. There's no transaction costs. So in theory, eBay fits these conditions. In practice it doesn't. And why is that? How does eBay violate-- how does shop at eBay violate all these conditions? What are some reasons why shopping on eBay isn't perfectly competitive?

Yeah. Yeah, basically, you don't really know the price. Indeed, for many years, eBay would rank firms by the price, excluding shipping. We call that a shrouded attribute in technical terms. They would hide the cost of shipping so really, you weren't really competing on prices. What else?

Well, it's hard to find truly identical products on eBay. Every product is a little bit different. So there's a lot-- and if products are different, you have to look into it. Suddenly that takes time. That's a transaction cost. So even eBay, which is the canonical example of perfectly competitive market, isn't.

Probably my favorite example is if you go to a tourist site, particularly in developing-- in many countries, there'll be lots of people selling the same shit around the tourist site. So think of the Eiffel Tower. There are literally hundreds of guys who have spread out blankets saying the sales on the same damn little Eiffel Tower keychain.

Now there, I can easily see what people are charging. They're right next to each other. There's little transaction costs. I can see it's the same damn keychain and it's not that hard just to ask what the price is. So that is an example. In fact, I taught that example in my online class and a student actually--

I got a package in the mail once, and I opened it. It was a bunch of little Eiffel Tower keychains fell out. And someone said they'd gone and done the exploration. And it turns out all the sellers right near the Eiffel Tower did charge identical price. But once you got further away, it was a different price. Why? Because that was not identical anymore. You're in a different spot. But everyone right around the Eiffel Tower charged exactly the same price for the Eiffel Tower, for the little Eiffel Tower statuette, which is pretty cool. So that's the kind of market we have in mind.

Now, let me highlight one thing once again. This is little q, not big Q. This does not say the market demand is perfectly elastic. The demand for any given firm is very elastic. I'm not saying the demand for Eiffel Tower statuettes in general is perfectly elastic. I'm saying that across sellers, the firm-specific demand is perfectly elastic.

Armed with these assumptions, and what a competitive market is, we can now talk about how do firms maximize profits? How do firms maximize profits? Well, basically, firms want to choose the optimal quantity to maximize profits as a function of q which equals revenue as a function of q minus costs a function of q. All firms are choosing as q. They're not price-makers, they're price-takers. They don't choose the price. They just-- the price is whatever.

When they set up their blanket in front of the Eiffel Tower, they find out what everyone is charging. There's basically no scope to set the price. All they can decide is how much quantity they want to sell and how they want to make it. They have control over their costs. So basically, if we take the derivative and set equal to 0, this says that if we set-- if we take the derivative and set equal to 0 to optimize this, that says that dR/dq equals dC/dq.

And this is it. This is profit maximization in one equation. Basically, you want to set the marginal revenue equal to the marginal cost. Well, in a competitive market, we know what the marginal revenue is. What do you get from selling the next unit? Well, you get the market price, p. It's given to you. The marginal revenue is always p. Whenever you sell the next unit, you get p.

We'll talk later about markets where that's not true, where the price you get depends on how much you sell. But in a competitive market, that's not true. No matter what you sell, you get p. It's a flat line. And we know what dC/dq is, it's marginal cost. So profit maximization in a perfectly competitive market means setting price equal to marginal cost. This is the rule for profit maximization in a perfectly competitive market.

This is the general rule for any market. Any market is going to set marginal revenue equals marginal cost. That's going to translate in a perfectly competitive market into price equals marginal cost. For example, let's go back to a cost curve we were using last time.

Let's imagine cost equals 10 plus 5q squared. And we know how to derive that. We did that last time. So let's go to figure 7-2. And let's assume p equals 30, just to make life easy. We'll assume p equals 30. Assume, and I'll tell you later where that comes from. But for now, just take it as given.

The left panel shows-- the left panel of this graph shows how costs and revenues evolve as costs increase. So basically, when you-- for every-- the revenue line is a line with the slope of 30. Why? Because every unit you sell, you get 30. That's the price. The cost line is the graph we saw last time of the cost function 10 plus 5q squared.

So basically, what you see-- and so you can translate that to the right-hand graph, which is what is the profit you want on every unit sold? So basically, what that says is the first-- if you're selling 0 units, you're losing money. Because what's the revenue? 0. What's the cost? 10. 10 plus 5q squared is 10. So your profit is negative 10.

But as you sell units, your profits rise until you get to the third unit where each of the first, second, and third units, each of those first three units, the money you make by selling the next unit is higher than the cost of producing the next unit.

The money you make from selling the next unit is higher than the cost of producing the next unit. But once you've reached three units, that's no longer true. When you go to producing the fourth unit, the money you make on the fourth unit is less than the cost of the fourth unit. The money you make on the fourth unit is less than the cost of the fourth unit. That's why profit maximization-- much like utility maximization-- is a hill-climbing exercise.

Remember I said in economics, you do marginal decision-making. Here is an example, of marginal decision-making. The idea of marginal decision-making is-- you think of yourself as climbing a mountain that's in the clouds. And you're trying to get to the top of the mountain, but you can't see more than one step in front of you.

If you take a step forward and you go up, you know you're heading up the mountain. Take a step forward and you go down, you're heading down the mountain. You don't need to know where the top of the mountain is. You just need to know where you're heading up or down. And when you've reached a point where the next step takes you down, then you're at the top.

That's what we're doing here. We're not saying-- we don't look at what the profit maximizing quantity is, we solve for it. We ask, does the next unit make money? Does the revenue made off the next unit exceed the cost of producing the next unit? Is that next step taking you up the hill or down the hill? And that's how we do profit maximization.

You can see this in figure 7-3. Here, we are actually taking our cost curves and adding revenue. So the curves on figure 7-3 are the same as the ones we developed last time, same as the first figure in the last lecture. All I've done here is I've added a line at p equals 30. I've added a line at p equals 30. That dashed line, by the way, should go all the way to the y-axis. We're missing a little connection there. But there's a line at p equals 30.

How do we solve for a firm's decision? Well, what's the approximate profit maximization? You set price equal to marginal cost. Where does price equal marginal cost? Well, it happens when marginal cost equals 30. So the first step is you find the point at which marginal cost equals price. That is producing three units. That is the point at which price equals marginal cost. That is the profit-maximizing point.

The second question is, what are the profits at that point? What are the profits at that point? Well, what is the average cost at that point? The average cost when you're producing three units is 10 over 3 plus 10 over 3 plus 5 times 3, 10 over 3 plus 5 times-- 5 times 9, sorry, because q decisions squared.

The average cost is 10 over 3 plus 5 times 9. Let me see what I got wrong here. Oh, I'm sorry. No, you divide by q. Sorry, you divide by-- average cost is 5 times 3. Because average cost is 10 over 3 plus 5q. Average cost is 10 over-- is dividing this by 3, by q. So it's 10 over 3 plus 5q is the average cost. The total costs are 10 plus 5 times 9. The average costs are 10 over 3 plus 5 times 3, which equals 18.33. That's the average cost.

What-- so on average, each unit you sell costs you 18.33. What do you sell it for? Well, you sell it for 30. So your profits per unit-- profits per unit-- is 30 minus 18.33, which equals 11.67. You're making 11.67 per unit. And that is the height of the blue box. So your profit is the profit per unit times the number of units sold.

You're making 11.67 per unit. You're selling three units. So your total profit is 40, which is the area of that box. Let's go over it again. You set price equal to marginal cost. That determines-- that determines the top of the hill. That is true for producing three units. At three units, your profits are your revenues minus your costs. We can divide both by q, say it's revenue per unit minus cost per unit times q. So profits is revenue minus cost. We can also write that as revenue over q minus cost over q times q, just pulling out, multiplying-- dividing by q.

Well, revenue per q at 3 is 30. That's a flat number. It's a flat line. Costs per q at 30 are 18.33. So revenues-- profits per unit are 11.67. You sell three units, that gets you to 40. OK? Questions about that? So that's how we determine the firm's profits.

So let me stress test you on this. Someone tell me what happens to the cost function if a firm has to pay a tax of $10 per unit it sells? We start with this cost function. Now I want to add a tax of $10 per unit. What is the new cost function? Yeah.

AUDIENCE: 10 plus 5q squared plus 10q.

PROFESSOR: Exactly. The new cost function is 10 plus 5q squared plus 10q. Many people often say 20 plus 5q squared. Usually that's the typical first answer I get. That's wrong. That'd would be true if the firm only had to pay an extra $10 flat, but they pay $10 per unit. So it's 10 plus 5q squared plus 10q. What does that do to our case?

Well, what that means is now marginal cost-- if we differentiate this with respect to q-- is 10q plus 10. Marginal cost used to be 10a. Now it's 10q plus 10. So the marginal cost curve has shifted up. All the cost curves are shifted up. So now, marginal cost equals price-- equals price at two units. 10q plus 10 equals 30, which is the price at q star equals 2.

I set marginal cost equal to the price. This is the new marginal cost, 10q plus 10. I set that equal to 30, which is the price. I get the new q star is 2. So quantity has fallen from 3 to 2. The price hasn't changed. Why has quantity fallen? Because it's more expensive to produce. And you're always trading off the benefit of producing the next unit against the cost of producing the next unit.

The benefit hasn't changed. The cost has gone up, so you produce fewer units. What does that mean for profits? Well, look what's happened to profits. They've shrunken from the larger rectangle to the smaller rectangle. Two things have changed.

First of all, you're selling fewer units. So the width of the rectangles changed. I see triangle? Rectangles. The width, the rectangles changed. Second of all, you're making less money per unit. Why are you making less money per unit? Because average costs are up, too. Marginal costs are up, but average costs are up, too. So in this case, you make less profit per unit. So you make less profits per unit and you're selling fewer units. So the profit rectangle has shrunk from the larger one to the smaller one.

That's an example of how change in cost structure-- I'd have to make the tax, I can make any change in cost structure-- how changes in cost structure can impact a firm's profits. OK, questions about that? All right.

Now, there's one additional feature to make this even harder in the short run, which is we need to decide whether the firm will shut down. I've just told you how to maximize profits at p equal to marginal cost, but we're not quite done. And that's because in the short run, losing money is not necessarily a reason to shut down.

You can lose money and want to stay in business. And the reason is because staying in the business allows you to re-optimize in the long run. So, for example, imagine the price for this good suddenly dropped to $10. Let's go-- let's not have the tax. Let's go back to marginal cost equals 10q. And imagine the margin being-- the price has just fallen to $10. The price is $10 equals 10.

So that says that q star equals 1. Marginal cost equals the price. Marginal cost that says q star equals 1. What is your profits producing 1? Well profits are fixed costs-- are 10 plus 5q squared, which is 5, or the costs, I'm sorry. That's the cost. You get a revenue of 10 and you have costs of 10 plus 5. So your profits are negative 5. Your profits are negative 5. You're losing money because you sell for 10. And the cost-- remember our cost formula-- it's 10 plus 5q squared, so your costs are 15. So you lose money.

Should you stop production? And the answer is no. And why not? Why should you keep going even though you're losing money? And remember, the key is opportunity cost. Yeah?

AUDIENCE: Because you're covering the cost of the variable?

PROFESSOR: Yeah, well, you're one step ahead. What would happen if you produce 0? What would your profits be then? No. What would your profits be if you produce 0? Minus 10. Because you've got to pay the fixed cost no matter what. Just plug 0 into this formula. If you produce 0, your profits are minus 10. So you're better off producing 1 than 0. You've already given me the answer why.

But you're better off producing 1 even though you're losing money. The opportunity cost of shutting down is you actually lose 10. So you're better off losing 5. And why is that? Why would you want to produce losing money? It's because in the long run, you could change your fixed costs. So you can come back the next period and make it up.

So when you want to shut down is not about whether you lose money. It's about whether you lose so much money that it overcomes your fixed costs. When that be true? Well, this gentleman gave the answer. That will be true when ignoring fixed costs, you still lose money. That will be true when your variable costs are below revenue. It will be true when pq is less than your variable costs.

Fixed costs are, in the short run, sunk. They're irrelevant. You already paid them. You just want to ask for the next unit. Will I make money money? Or, in other words, you will keep producing as long as p is-- as long as p isn't less than VC over q, the variable costs over q. As long as the price-- which we call average variable costs-- as long as the price is above average variable costs, you will keep producing. Only once price drops below average variable costs will you stop producing. We call this the shutdown rule.

Now it turns out in our case, that's never true. Because if you look at our formula, variable costs-- so our cost function is C equals 10 plus 5q squared. So variable costs are 5q squared. Average variable costs are 5q. Average variable costs are 5q. I'm sorry, 5-- average variable costs are 5q.

Now remember for our case-- for this case, the general formula for optimization is marginal cost equals price. So that's 10q equals p. That's our optimization condition, 10q equals p, marginal cost equals price. Well, if you plug-- so if you plug that in, you get that average variable cost in terms of price is 0.5 times p.

So what I did here, I created the variable cost. I divided to make average variable costs, which is 5q. I then substituted it in for the optimization condition to get average variable cost 0.5 times p. Well, 0.5 times p can never be greater than p. Average-- this number is 0.5 times p can never be greater than price. This condition can never hold.

As a result with this cost function, you'll never shut down. It doesn't matter what profits are, you will not shut down. OK? This is a confusing concept that you need to work on, but it's important to remember that when we develop supply curves, we have to consider both what the supply curve looks like and in the short run, whether you want to shut down. So something you have to practice. It's always a confusing concept, but we'll give you some practice on that.

Now, armed with all this, we can come back to actually creating what you wanted. You are demanding. You said John, where does that supply curve come from? I need to know. And I'm going to tell you now where the supply curves comes from. Now we get to wrap it all up. Where the supply curve come from, well, that's in figure 7-5. Figure 7-5 shows you how much the firm wants to produce at different levels of prices.

Well, when the price is 30, we know it wants to produce three units. When the price is 40, it says marginal cost equals the price, that turns out to be 4 units. It turns out the firm's production decision is-- price production curve is the marginal cost curve. The supply curve is the marginal cost curve that-- so we already, two lectures ago told you, or maybe it was last lecture-- told you what supply is.

Supply curve is the marginal cost curve. And think about the intuition. The demand curve-- one second-- the demand curve was the marginal willingness to pay for a good. The demand curve was how much you're willing to pay for the next good. The supply curve is the marginal cost of producing the next good. The supply curve is what does it cost you to produce the next good.

So demand and supply meet with-- and I'll get to that in a couple of lectures-- where marginal willingness to pay for the good, how much you're willing to pay for the next good, equals the marginal cost of producing that next good. That's where we get equilibrium Yes. Question here.

AUDIENCE: So for that, p is less than VC over q. So for the right-hand side is the average--

PROFESSOR: Average variable cost.

AUDIENCE: So could the left-hand also be price instead of price times unit?

PROFESSOR: No, because I divided. It's-- let me do it. I did p times q is less than variable cost. I then divided both sides by q. So price should be less than variable-- than average variable cost.

AUDIENCE: So is that form of p would be like [INAUDIBLE].

PROFESSOR: No. Let me back up. This is a general formula. You'll shut down when price is less than variable costs. That's the general formula. Then I said in our case, what is average variable cost? It's 5q, OK? Now in equilibrium, we know what q is. We know q is p over 10. So I can just plug that in.

And then I show for our case, this formula never hold.

AUDIENCE: So then, in [INAUDIBLE]?

PROFESSOR: Giving q what?

AUDIENCE: [INAUDIBLE]

PROFESSOR: q at-- no, not at equilibrium, at profit maximization. Given q, if a firm is profit-maximizing, they'll use this rule, marginal cost equals price. If they use that rule, they'll never shut down. Does that make sense? OK. Good question. Good clarification. Other questions about this?

AUDIENCE: What is-- [INAUDIBLE] maybe happiness or something? But [INAUDIBLE] this is--

PROFESSOR: That's a great point. We are-- we're dealing in integer worlds. You can-- but you can do the whole thing fractionally. The mathematics still holds, the mathematics still holds. You just want-- really, we're always doing-- it's a derivative, ultimately. The question is for the next epsilon unit of production, is the variable cost greater than the price you can sell it for? OK? Yeah.

AUDIENCE: So at the very bottom of the board on the left, you said that ABC of equal to [INAUDIBLE]?

PROFESSOR: Yes, because it implies ABC equals 5q and 10q equals p. So I plugged in.

AUDIENCE: Thank you.

PROFESSOR: Yep. OK. So now, so the point is, the short-run supply curve is simply the marginal cost curve, but above the shutdown point. So the supply curve has two things. It's the marginal cost curve, but only above the shutdown point. In our case, there is no shutdown point, so it's irrelevant. But in general, the rule is the short-run supply curve is the marginal cost curve above the shutdown point is the technical definition of the short-run supply curve.

So it's above minimum average variable cost. So technically, the short-run supply curve is the marginal cost curve above minimum average variable costs. And that's where the supply curve comes from. Now, that is the firm's supply curve. But we started this course with markets. So let's go to markets.

What does the market supply curve look like? I showed you this firm's supply is Figure 7-6. What does the market supply curve look like? Well, let's go to Figure 7-7. Imagine initially, we have one firm. That's S1. We know what that firm's supply curve looks like. We just developed it. At a price of 10, they produce 1. At a price of 30, they produce 3, and so on. That's S1.

Now let's add a second identical firm. Well, now at a price of 10, you get 1 produced by each firm. That's 2. And at a price of 30, you get 3 produced by each firm. That's six. Now let's add a third identical firm, and so on. As you add identical firms, the market supply curve is a flatter version of the firm's supply curve.

In other words, market supply is a more elastic version of firm's supply. Market supply is a more elastic version of firm's supply. So as more identical firms produce the good, as more identical firms produce the good, the more elastic is the supply curve.

Now, this will give you a hint to where we head next time. Just want to ask a question about next time. If I said competitive markets have an infinite number of firms and I said the more firms are, the more elastic the supply curve, think about what that implies to the supply curve and we'll talk about that next time about how that works.

But for now, let's talk about one more thing. And then I want to go back and redo it all to make it clear. The last thing I want to talk about is the long-run supply curve. So far, I've been talking about the short-run only. What about the long-run supply curve?

Well, it turns out the long-run supply curve is just the marginal cost curve, but without the shutdown rule, because in the long run, if profits are negative, you do shut down. Because the long-run fixed costs aren't fixed, right? So in the long run, the supply curve is just the marginal cost curve, full stop. As long as profits are greater than 0, there's no shutdown rule. As long as you're making money, you pick a point where price equals marginal cost.

The only difference in the short run and the long run here is that in the long run, you don't have this extra annoying shutdown rule. It's just basically-- the long and supply curve is marginal cost as long as marginal-- as long as price is above-- as long as profits are greater than 0. So that is the long run.

So the idea is in the short run, you're willing to lose money because you can always re-optimize your capital. You can always go back and change your fixed costs. But in the long run, you've already optimized it. If you've done all your optimizing, you're still losing money, you're just shit out of luck. You just-- you should get out of the business because there's no more optimization you can do. All right? Questions about that.

All right, so let's now put this all together to derive the short-run equilibrium. I'm going to mathematically get you back to where we started in the first lecture. I'm going to start with the primitives. How do you find short equilibrium? The first is Step A is you get the supply curve. How do you get the supply curve? Well, we know all the steps.

You start with the production function. Step 1, start with the production function. In our case, the production function was q equals square root of l times k. That's step 1. Step 2 is you get input prices, which once again, are still given at this point. We haven't endogenized them. They're still given. And we had w equals 5, r equals 10. OK?

Step 3-- I'm sorry, you have your assumptions. And then the third thing is you have a fixed amount of capital, k bar equals 1. So you start with production function. You have your three assumptions. Once again, we know where k bar equals 1 comes from. That comes from long-run optimization where w and r come from, you haven't learned yet. But in the short run, k is fixed. So you start with production. You have those three things. We know based on that how to develop a cost function, OK?

We know how to develop a cost function. We know that with these three we develop-- we developed in a previous lecture that with this production function and these three assumptions, you get that the cost as a function of quantity is 10 plus 5q squared. We developed that in the previous lecture. So as I said, if I give you a production function, I give you an amount of capital, and I give you input prices, you can solve for the firm's cost function. You know how to do that now.

Now, if you know the firm's cost function, you know marginal cost. Well, marginal cost is just the derivative of cost respect to quantity, which in our case is 10q. We also know that in the firm's maximized profits when marginal cost equals price. In other words, the supply curve is q equals 10 equals p over 10. q equals p over 10 is our supply curve. We've just developed the supply curve. All I had to tell you to get the supply curve was a production function and those three key input assumptions. Based on that, you have all you need to develop a supply curve.

You then have to check the shutdown condition, which we know in this case is irrelevant. But you have to check it. So it's q plus p over 10 where average variable cost is greater than price, which we know is always-- which we know is always true. I'm sorry, where price is greater, less-- everywhere costs less than price. I'm sorry, which we know is always true in this case.

So that's the supply curve. We've just developed it. And I only had to give you-- I only had to give you two things, the production function and those-- and those three assumptions. Questions about that? Yeah.

AUDIENCE: [INAUDIBLE] 5q squared?

PROFESSOR: 5q squared. So basically, just quick, quick math on this. OK, let's go back and do the math on this real quickly just to review. Take this production function. q equals square root of k to l. Break out l. That's q squared over k bar. l is q squared over k bar. Then re-express-- then write a cost equation as a function of just q. You can write that as cost, k and q, I'm sorry. M cost equals 10 k bar plus 5q squared over k bar.

I just-- because that's the amount-- that's the right-- that's the amount of labor you want to use. So that's your cost function. You plug in k bar equals 1, and you get C equals 10 plus 5q squared. OK? Does that make sense?

AUDIENCE: Yeah, you bet.

PROFESSOR: OK. So now we have the firm's supply curve, the firm's supply curve, leads us to Step B, which is, we got to make a market supply curve. Well, to make the market supply curve, we need to know how many firms are in the market. Now I'm going to have to throw another assumption at you. Next time I'll tell you where to get this from.

All the assumptions I'm making, you'll eventually going to know where they come from. You know where w comes from, and r comes from, and k comes from. Another subject I'm going to throw at you is little n, the number of firms in the market. And let's just assume for now little n equals 6. And I'll tell you next lecture how you get that. But for now, it's useful to start with an assumption. Assume there's six firms in the market.

So what is the total quantity in the market? Well, the total quantity is just six little qs. And we know little q equals p over 10. We know little q equals p over-- we know little q equals p over 10. So the supply curve is q equals 3/5 times p. That's the market supply curve, which is just the firm's supply curve applied to six firms. That's the market supply curve. That's Step 2.

Step 3 is you have a demand curve. Now I'm not going to go back through all the math of that. We spent a couple of lectures on this. But you know how to derive a demand curve as long as I give you prices and income and utility. You know how to derive a demand curve. You've done that. You do it on the problem set. You'll do it a bunch of times.

Let's say, in this case, just to make life easy, the demand curve I get if I do that is 48q equals 48 minus p, the market demand curve I'm not going to go back through all those steps because that would be too much. But hopefully, you can go back and review your consumer theory to get that. And I'm just-- I mean, just making this up for this example. But how to derive a demand curve.

So now, we get equilibrium is where demand equals supply, where 48 minus p equals 3/5 p. Two equations-- we had two-- we set-- we set these equal to each other to get equilibrium. So we get that q star equals 18. The market-- the equilibrium quantity in the market is 18. Why? Because I know the market supply curve. I know the market demand curve, which I just made up for this case. But you know how to get one.

I set them equal to each other. And I get q star equals 18. But that's the market supply. What about each firm? You want to know what each firm is going to produce. Well, we know that there's six firms in this market, and they're identical. So that means that in this market, each firm for q star to be 18, little q star must be 3.

Well, how do we know if that's the case? Well, we know that if q star equals 18, what's the price? Well, we know q equals 3/5 p. So if q star equals 18, p star equals what? 30. There's two equations and two unknowns. One equation was q equals 48 minus p. One equation was q equals 3/5 p. I solved them for q star, but I can also solve for p star. p star is 30. You can get it from the demand curve too. 48 minus 18 is 30. So p star is 30.

Well, we know each firm sets price equal to marginal cost. So now we can go back to get the firm's supply. Firm supply is where price equals marginal cost. Price is 30. Marginal cost is 10q. So optimal firm production is 3, which is what it has to be. The magic of this market is that each firm is producing exactly what the aggregate market says they should produce.

So basically we solved-- we took a firm-- we took firms. We built them up to the firms that were taking a given price. Remember, we went to the firm. We didn't specify price. Prices-- they're price-takers. That's just a given from God from their perspective When you set up your blanket outside the Eiffel Tower, just look at what everybody else is charging.

We then-- so we left p to be solved for. We then used demand. We then take that firm's decision, make it a market decision. We then said we'll set that market decision equal to market demand. And that gives you an optimal market supply and market price. And that market supply implies that each firm must be producing 3. Well, it turns out at that price, each firm is producing 3. So you're at an equilibrium where everybody's satisfied.

Consumers are satisfied because the demand is being met. Consumer's willingness to pay is 48 minus p. Well, the price is 18. I'm sorry. The price is 30, so they want 18. That's what they get. Firms are satisfied because aggregate supply is 3/5 p. And each individual firm is satisfied because price equals marginal cost. We're done. That's the magic of the market, OK?

And this is the beauty of being at MIT and teaching you what is usually taught in intermediate microeconomics at most universities. You'll never find an introductory economics course that does this OK, which this is the magic of bringing mathematics to economics. Before Paul Samuelson taught this class, people would teach the graph and be done with it. And what we've done is we've said, you guys want the math? I'm giving you the math.

I have never, ever gotten a complaint in this course that I do too much math. I only get complaints I do too little math, which is unbelievably unique. I can't imagine any other university in the world where the major complaint you get is you do too little math. OK? So I wanted to explicitly write out here the math because I think this is a helpful way to understand it. But you should have the intuition, the graphics, and the math. Questions about that? Yeah.

AUDIENCE: What is the 3/5p about?

PROFESSOR: The 3/5p is the aggregate quantity. I assume there's six identical firms. Each firm, I said, was producing 10 over p. So six times 10 over p is 3/5 p. Good clarifying question. Yeah.

AUDIENCE: [INAUDIBLE]

PROFESSOR: q star equals 18. Well, I solved here. I had two equations and two unknown. One is the demand curve, Hugo's 48 minus p. The other supply curve-- q equals 3/5p. So I set 48 minus p equal to 3/5p, and I get q star equals 18 and p star equals 30. Other questions? These are great clarifying questions. Remember, if you're confused, most of the class is confused. You're representative of the class, so please ask those questions. Other questions? Yeah.

AUDIENCE: In solving for price, [INAUDIBLE]?

PROFESSOR: Well, you saw it's two equations, two unknowns. You're solving for both, right? So it's a two-equation. We have two equations-- q equals 3/5p and q equals 48 minus p. You're solving for two unknowns-- two equations and two unknowns. Yeah.

AUDIENCE: [INAUDIBLE] by supply curve?

PROFESSOR: The supply curve, all you need is the firm's production function, input prices, and the initial level of capital. And the rest is math you've done and should keep doing if you're not sure about it. Great questions. OK? Let's stop there. And I hope you guys-- this is a-- once again, it's a great month for concerts. I hope you guys are getting out to at least one this month, and I'll see you next week.