

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit mitopencourseware@ocw.mit.edu.

PROFESSOR: All right. So today we're going to continue our discussion of the perfectly competitive market outcome. And remember, once again, where we're coming from here. We're trying to figure out how firms decide how much to produce. We talked about the firm's production decision in costs. Then we said how the firm decides how much to produce is going to get dictated by the market. We're going to talk about different market structures. We'll start with our benchmark of perfect competition and then move on to some more interesting and realistic cases next.

Now, I just want to finish up where we were last time. Remember last time we were dealing with a firm that had a cost function of the form $10 + 0.5q^2$. And if you remember, the key condition we derived last time for profit maximization with a perfectly competitive firm is that price equals marginal cost. If you differentiate this with respect to q , you get that that means that $p = q$ is the profit maximizing condition for this firm. It sets the price equal to quantity it's going to sell. That's the profit maximizing condition with this particular functional form of the cost function.

Now, before we left last time, I said, this was not enough. There's one other thing you have to consider which is the firm's shutdown decision. In the short run, a firm might not shut down even if it's losing money. And the reason is because the firm has already paid its fixed costs. It's already paid 10. So even if it's losing money, it might still not shut down.

So, for example, imagine that the price, as I said last time, imagine the price fell from 6 to 3. The price equals 3. If the price equals 3, the firm will choose to produce 3 units, because it will still follow the profit maximizing condition. It will still choose to produce three units. If it produces 3 units, its profits are its revenues which is $9 - 3$ units at a price of 3-- minus its cost which is 14.5. So that equals negative 5.5. So its profits are negative.

So now I say, well, if profits are negative, maybe I should shut down and stop doing business. Well, what are its profits if it shuts down? Negative 10. If it shuts down, its profits are $0 - 10$ equals negative 10. So it actually makes more money by staying in business than shutting down, because it has these fixed costs.

So because it's going to pay the 10 anyway, as long as it's going to lose less than 10, it might as well stay in business. More generally, what we say is a firm will stay in business in the short run as long as its price covers its variable costs. So a firm will stay in business so long as the price is greater or equal to its variable costs. Then it will stay in business.

If we go further-- let me just derive it for a second-- as long as they cover its fixed costs, that means a firm will stay in business as long as its revenues are greater than its variable costs. As long as its revenues are greater than or equal to its variable costs, it will stay in business. And that means that it will stay in business as long as its price is greater than or equal-- this should be a little q , sorry, this is just a firm-- as long as its price is greater than or equal to variable costs over quantity, or as long as price is greater than or equal to average variable cost.

As long as its price is greater than or equal to its average variable cost, it will stay in business. As long as its revenues cover its variable costs, it will stay in business. And that's saying the same as long as its price is greater than or equal to average variable cost, it will stay in business.

Now, what are the average variable costs for our firm? Well, the variable costs for our firm are 0.5 times q squared. The variable costs for our firm are 0.5 times q squared. We know that in equilibrium, if it's profit maximizing, it will produce where q equals p . So we can replace the q with the p . The variable costs are 0.5 times p squared, because we know we'll produce where q equals p . So its average variable costs are 0.5 times p . Its average variable costs are 0.5 times p .

Well, by definition, p is always greater than 0.5 times p . So our firm will never go out of business in the short run. In the short run, our firm will never go out of business, because at the profit maximizing price, at p equals q , it will always be producing a point where the price is greater than its average variable cost. So it will never shut down.

So, more generally, when we think about a short run supply decision, a short run supply decision for a firm, there's two steps. The first step is set price equal to marginal cost to figure out what the firm is going to produce. So step one is set price equal to marginal cost. And that will give you the firm's q^* . That will give you what the firm is going to produce.

The second step is check that price is greater than or equal to average variable costs.

Because you may solve for an optimal quantity that turns out to be a money loser for the firm. So they'd rather shut down. So it's a two-step process. You've got to first solve for the optimal quantity that the firm is going to produce. But then you've got to make sure that the firm actually makes money on that quantity, or it won't produce at all.

And that's how we do the profit maximization decision in the short run for the firm. You've got to produce at the efficient point and make sure the firm actually makes some money. All right, questions about that?

Now, armed with these rules, we can now, finally, derive the supply curve. Remember we derived the demand curve a number of lectures ago by getting the tangency at different price ratios with the indifference curves. Well, to derive the firm's supply function, what we now need to do is say, OK, at different prices, how much will the firm produce?

Well, we can now get that if we go to Figure 11-1. We can now see the supply curve for this firm. What we see is that at a price of 3, it will produce 3 units. At a price of 4, it will produce 4 units, et cetera.

The supply curve is the marginal cost curve. So now we know where supply curves come from. Supply curves are marginal cost curves above the point where price equals average variable cost. So the definition of a firm's supply curve is the marginal cost curve above p is greater than or equal to average variable cost. That is the firm's short run supply curve.

Now, in our case, p is always greater than average variable cost. So the second condition is irrelevant. The firm's supply curve is just literally that marginal cost curve. With different functions, which you may someday see in a problem set or an exam, that won't be true. So, in that case, you'll need to check that shutdown condition.

But the supply curve is the marginal cost curve above that 0 profit point. And that's where supply curves come from. So where supply curves come from is the same kind of maximization we did with consumers. But instead of their parents giving them their income, the market conditions firms face are dictated by the competitive nature of the market. And that's the firm's supply curve.

Now, this is the firm's supply curve. Now, of course, what we talked about in the first lecture was not firm supply curves but market supply curves. So now let's take the next step and say, well, where do market supply curves come from? We now know where firm supply curves

come from. The marginal cost stork brings them. Now, where do market supply curves come from?

Well, to do that, we need to now imagine that there's not one firm in the market but many firms in the market. And we need to recognize that the market demand may not be perfectly elastic. But, as we talked about last time, the firm's own demand will be close to perfectly elastic. Or, in this case, a perfect competition will be perfectly elastic.

So, basically, the way to get market demand is to say, look, we're going to take each firm. It's going to take a market price as given. I'm sorry, we get market supply. I'm sorry. Each firm is going to take a market price as given. Based on that market price, it's going to decide how much to produce.

We're going to add up that production. That will make a market supply curve. And that market supply curve will then interact with market demand to give you a price. If that price is the same one the firms were using, then the whole thing is in equilibrium.

Let me explain that in less steps just to make it clear. Let's talk about the steps involved in getting to short run equilibrium in the market, the steps involved in getting to short run market equilibrium. The first step is each firm chooses an amount of capital.

So the first step of the short run is you're going to enter this market. And to enter this market, you're going to have an amount of capital you're going to pick. So each firm is going to have some cost function which involves picking some amount of capital or fixed costs. It's going to say, I want to build a building this big. Having built that building, we're going to get the firm's supply curve which is $p = MC$. That's step one. That's a step we've derived.

The second step is we're going to add up the firm's supply curves to get a market supply curve. We're going to add up the firm's supply curve to get a market supply curve. So, for example, suppose that there's five firms in the market. Suppose that there's five firms in the market. Those five firms are going to produce.

Now to see that, let's go to Figure 11-2. This is the second step. It's how we get to that short run market supply curve. Each firm has a marginal cost curve. Here we're using our same cost function we've used. That same cost function up there where price equals marginal cost, where $p = q$, is the supply curve. So each firm has that supply curve you see in the first panel.

Then what you see is as you add more firms, the second panel gives you the market supply curve. So if there's only one firm in the market, the market supply curve would be S_1 . Now, if there were two firms in the market, the market supply curve is S_2 . That is, at a price of 2, you're now producing 4 units in the market. If there's three firms, the curve is S_3 , four firms, S_4 , and so on.

As you add more firms, that market supply curve shifts out and becomes flatter. Remember firms are identical here. We're adding identical firms. That was an assumption of perfect competition. We're adding more and more identical firms producing the same good. You can see that market supply curve is shifting out and becoming flatter.

That is the supply of goods is becoming more elastic as there are more firms. The more firms in the market the more elastic the supply. And that comes to what we talked about last time when I derived residual demand. It's sort of the flip side of that. Basically, the more firms you have with a given supply curve in the market, the more elastic it's going to become.

Why is that? Well, just think about it. Think about what elasticity of supply is. It's saying, if I increase the price by \$1, how much more production do I call forth? Well, the more identical firms I have, every time I increase the price by \$1, I call forth production from all these firms. So the more firms I have, the more production I call forth.

So for every increment in price, the more firms in the market, the more production I call forth. Therefore, the more elastic is the supply. So as there are more firms, that market supply curve becomes more and more elastic. And that's the market supply curve. That's the second step.

The third step is we intersect market supply with market demand to get the equilibrium price. So, in other words, we say, look, there's some market supply, which we've derived. Now let's imagine there's some market demand. And that will give us the equilibrium price.

So, for example, in our case, market supply is what? Let's say, for example, there's five firms in the market. Just to make an example, let's say five firms have entered. There's five firms in the market. Well, the total market supply Q is 5 of the little q , because there's five identical firms in the market. Five identical firms are in the market.

Well, we know, from the marginal cost condition, that's the same as saying Q equals 5 times p . So our market supply curve, which is actually S_5 on Figure 11-2, is Q equals $5p$. You can see that. Because when the price is 2, Q equals 10. When the price is 5, Q equals 25. So you can

see that $S_{\text{super } 5}$ is the market supply curve. Q equals $5p$.

Let's just make this up. So this is the quantity supplied. Let's say the demand function is that the quantity demanded is $30 - p$. I just made this up. I'm making all this up. But this is just an example demand curve. We have a downward sloping demand curve with a slope of negative 1. The quantity demanded is $30 - p$.

So to get equilibrium, we set these equal. And we get that $30 - p = 5p$ or $p = 5$. $30 - p = 5p$ or $p = 5$, that's the equilibrium price. Given the market supply curve, given the demand curve, I've derived the equilibrium price of 5.

Now, at a price of 5, what's the quantity demanded? At a price of 5, quantity demanded equals 25. So, at a price of 5, the market wants 25 of these things, whatever the heck it's producing. At a price of 5, the market wants 25 of them. That's the quantity demanded.

Then the final step in solving for equilibrium is that each firm then decides how much to produce. Well, what is each firm going to decide to produce? How much is each firm going to decide to produce? Somebody raise their hand and tell me. Yeah.

AUDIENCE: p .

PROFESSOR: p which is?

AUDIENCE: 5

PROFESSOR: 5. So each firm is going to produce 5. How many firms are there?

AUDIENCE: 5.

PROFESSOR: How much gets produced?

AUDIENCE: 25.

PROFESSOR: Which is exactly what people want. We're done. That's the magic of the market. Through these four steps, we've gotten equilibrium which is defined as the quantity supplied equals the quantity demanded, just by following these steps.

The firms didn't come at it saying, hey, let's figure this out beforehand. Let's all get together and coordinate and figure out how much we're going to produce at what price. They didn't do that at all. The firms just entered the market. They said, this is my supply curve. We added it

up. We interact it with demand. We find an equilibrium price. At that equilibrium price, the quantity demanded equals the quantity supplied. We're done. We've now finally gotten to where we started this course.

We started this course with a supply and demand graph. I told you how to derive demand. You took a test on that already. Now I've just told you where supply comes from. Now we interact, and we're done. And this tells us now, given this price, how much each firm is going to produce.

So, basically, what do you need to find equilibrium? What do you need to find the short run equilibrium? To find the short run equilibrium, you need a demand function, a cost function, and a number of firms. You have to be given a number of firms, because there's no entry and exit in the short run, remember. So, basically, the firms are magically put on the earth in the short run.

So if you're asked about the short run equilibrium, you have to be given the number of firms. You can't derive that yet. That comes in the long run. But given a number of firms, given a cost function, and given the demand function, you can find the short run equilibrium. Questions about that?

Now, with that in place, now let's get to where it gets really interesting. Let's talk now about the long run. Now what makes the long run interesting is you no longer have to be given the number of firms. Now, in the long run, we're going to actually figure out how many firms are in the market. So at the end of the day, in the long run, all we'll need is two things, a demand function and a cost function. And then we'll be done.

Now, here's the key thing for the long run. The key point in the long run is that in the long run, no one can lose money. So we don't have to worry about the shutdown condition anymore. The shutdown condition goes away. In the long run, nothing is fixed. So in the long run, no one loses money. There's no reason to be in the market if you're losing money in the long run.

So, in long run, the first thing is now we only have one condition to worry about, which is price equals marginal cost. Price equals marginal cost is what we need to worry about. And you're only going to be in a situation, in the long run, where you're only going to be making either 0 or positive profits. If you're making negative profits, you'll be gone.

Now, the key difference in the long run, is now we can't take the number of firms as given.

Now we need to derive the number of firms. And the way we do that is by thinking about entry and exit. Now, what's going to determine entry and exit? Well, it's quite simple. If in the market, as it stands today with some number of firms, there's profit to be made, new firms will enter. If in the market, as it stands today with some number of firms, there's losses being made, some firms will leave. Remember, no one stays in making losses anymore.

And that continues until you reach a situation where all firms make zero profit. And here's the key lesson. In a perfectly competitive long run equilibrium, all firms make zero profit. It's the fundamental lesson about perfect competition.

Obviously, there's no place that works like this in the world. This is an extreme. But, nonetheless, you should understand this extreme. In a perfectly competitive long run equilibrium, all firms make zero profit. Why? Because if there's any profit to be made, a new firm will enter and take it away. And if there's any unprofitable industry, a firm will exit until the profits go back to zero. So profits will always be zero in the long run equilibrium.

Now, to understand how this works, let's think about a realistic example. Let's think about the PC market circa 1990 when all you youngins were being born. It's circa about 1990, the PC market, cast your mind way back. It's a history lesson for you guys.

In 1990, not that many folks had PCs. There was still a vibrant use of mainframe computers. And, basically, you had big firms like IBM who were producing these mainframe computers, which is what I did my computing on. A lot of people did their computing on it in 1990. And there was starting to be a market, however, for personal computers. The chip strength had gotten large enough that it was actually viable to have desktop computing that was powerful enough. Firms like Dell were starting out-- not starting out, necessarily-- but were starting to make money making desktop computers, making PCs.

Let's actually think about how that market might look. So let's actually talk about, in Figure 11-3, this is the market for PCs circa 1990. In 1990, if you were making PCs, that was a great business to be in. Because people wanted them, there weren't that many firms making them, and you were making a killing. So if you were Dell in 1990, that was a great place to be.

So let's say Dell in 1990 was facing a demand curve D and the supply curve SR_1 . The supply curve was pretty steep because there weren't many firms making PCs. So the market price was P_1 .

So on the right, you have the market. We should label this actually. On the right, you have the market. On the left, you have Dell. On the right, you have the market. In the market, in initial equilibrium, there's a price of P_1 with big Q_1 being sold. So Q_1 PCs are being sold at a high price of P_1 . It's the novel technology. People want it, but not many firms are doing it.

What happens with Dell? Now let's go to the left-hand side diagram. Well, Dell is producing where that price equals their marginal cost. That price equals their marginal cost at little q_1 . So Dell's producing little q_1 .

But its average costs at that point are all the way down. It's not really labeled. But you can see it's where that vertical line for Q_1 intersects average total cost. That's where their costs are. So, in each unit, they're making the height between marginal cost and average total cost at that Q_1 . They're making that vertical bar. So they make that entire rectangle of profit.

So Dell makes a big profit, because not many firms are in this business. And yet demand for PCs are high. And that's where things are circa 1990.

So what happens? Well, Gateway arrives. Does Gateway still exist? Do people still buy Gateway? Gateway is gone, right? So, while Gateway was big, they were sort of a big upstart. They had the boxes with cow colors on them and stuff. You guys didn't even know Gateway?

Well, they were big. They were the first cheap knockoff computer makers that competed with the big folks. And they came and said look, we can do this. OK, we can produce. This is a profitable business. We can make PCs. It's not that expensive. It's largely a variable cost business. You just have to build your plant first.

So they built their plant, and then they come in. Well, what happens when a new firm comes in? The market supply curve flattens. Because now, at any price, you're producing more. So the market supply curve flattens to the point SR_2 . In fact, maybe it's not just Gateway. Maybe you get a bunch of entrants until you get the market supply curve SR_2 .

Well, SR_2 intersects demand at a new higher market quantity big Q_2 . Well, that higher market quantity going to the left, there's now no longer profits to be made. Because at that market quantity, Dell is going to produce little q_2 . Little q_2 is exactly at the minimum of the average total cost curve. It's where the marginal cost curve intersects the average total cost curve, the minimum of that average total cost curve. So Dell no longer makes profits.

Dell shrinks its production. The price has fallen. It doesn't lower its price. Dell doesn't set the

price. Remember, this a price taker in a perfectly competitive market. The price is given by that diagram on the right. Dell gets a price of P_2 . It says look, at P_2 , I have to produce along my marginal cost curve. I have no choice. That's what's profit maximizing. I can't choose a point not on that marginal cost curve. That will not be profit maximizing.

Well, where does that price intersect that marginal cost curve? At little q_2 . So I'm producing at little q_2 . And at little q_2 , I make no profit. So the entry of Gateway and other firms into the PC business has removed the profit from the PC business.

Now note what's interesting here. Market quantity has gone up. Big Q_2 is bigger than big Q_1 . But Dell's quantity has gone down. Little q_2 is smaller than little q_1 . That's because more firms are in the market producing. So as more firms come in, total market quantity goes up. But any given firm is going to produce less. And that will continue until profits go to zero. That is how firm entry wipes out profits. That is how firm entry wipes out profits, by driving firms to the point where price equals average cost.

So, in the long run, firms make zero profit because, first of all, entry drives price down to average cost. Entry drives price down to average cost. And when price equals average cost, profits are zero. Profits are zero when price equals average cost. Because profits are pq minus C . So if you divide by q profits are p minus average costs. So if price equals average cost, profits are zero.

So entry drives profits to zero. It drives price to equal average cost. Since price equals marginal cost, it's the point where marginal cost equals average cost. That's the technological outcome in a perfectly competitive long run equilibrium. You'll end up producing where marginal cost equals average cost. That's what will end up happening naturally through the forces of entry.

Likewise, you see this through the force of exit. Now let's go to Figure 11-4. And now let's look at IBM. I guess we're calling this the broader computer market now. It's not just a PC market. It's the broader computer market.

So IBM, they're producing these mainframes. This is the mainframe market. This isn't the PC market. This is the mainframe market. In the mainframe market, now people don't want mainframes much anymore. But there were a lot of firms producing. There was IBM and tons of other firms producing these mainframes. Because that's what everybody wanted.

So the original supply curve was very flat at SR_1 . So, initially, in the mainframe market, we're in equilibrium with quantity Q_1 -- big Q_1 -- and a price P_1 . And what's happening there? Where does that price intersect marginal cost? It intersects marginal cost for IBM at little q_1 which is below average total cost. So IBM is losing money. In the initial short run equilibrium, IBM is producing at little q_1 , and it's losing money.

Now, why is it still in business? Because it's the short run. And as long as those losses are less than its fixed cost, it's staying in business. So, in the short run, you can lose money. So IBM is losing money. Because it's built this big plant. It's cranking out these mainframes. People don't want them anymore.

The price has fallen so low that they can still make enough money than it costs to produce the next mainframe. Price is still greater than marginal costs, but price is below average costs. I'm sorry. Prices are greater than average variable cost. But it's lower the total average cost. They're losing money. So what happens? They leave.

There was a company called [? Deck ?] that went out of business. And what happened was that then raised the supply curve. It steepened the supply curve in the mainframe market. It steepened the supply curve, because now you have fewer firms producing mainframes. That supply curve gets steeper. That raises the price. And, indeed, exit will continue until you raise the price to the point where marginal cost equals average total cost.

And what you'll see is the market will shrink from big Q_1 to big Q_2 . The remaining market participants will increase from little q_1 to little q_2 . And profits go to 0 with price going to the minimum average cost. So through both entry and exit, we get this condition that's illustrated in Figure 11-5. In 11-5, we see that in the long run, firms always supply not on a single curve but at a single point.

In the long run, with a perfectly competitive market, for a given firm, there is no longer even meaningfully a supply curve to a firm. There's just literally a supply point. Every firm produces at exactly the point where marginal costs equal average costs.

So in some sense, once again, for a given firm-- this is not the market-- but for a given firm, there's not even meaningfully a supply curve anymore. For a given firm, in the long run, they literally choose one production point which is technologically given.

So this is interesting. For a given firm, the market doesn't matter. For a given firm in a

perfectly competitive market, we don't need to know anything about demand. All we need to know is the firm's production function. That's all we need to know.

We don't even need to know anything about costs. Well, we need to know cost. We need to know their cost function. All we need to know is their cost function. And then all we need to do is derive where marginal costs equals average costs, and we're done.

This is the power of the perfectly competitive equilibrium. This is why economists love it so much. Because we don't need to go through all this. This is all short run stuff. In the long run, it's easy. You just say, give me a cost function, I'll tell you what the firm will produce. And I'll tell you what the price is.

The firm will produce where marginal cost equals average cost. And the price will be where marginal cost equals average costs. I can tell you the p and the q in equilibrium just if you give me a cost function. And that's the beauty of the long run perfectly competitive equilibrium. That's why it's so attractive to economists for modeling purposes and other things. It's incredibly easy to work with. Because all you need is a cost function, and you're done.

The key lesson is what is true at the point where marginal costs equals average costs? Well, look at our graph. That is the point of cost minimization. Note that that is the very minimum point of the long run average cost curve. So where marginal cost equals average cost is the point of cost minimization.

So we're saying further-- this is even more powerful-- we're saying that in the long run of perfectly competitive equilibrium firms will, by definition, minimize their costs. They will produce as efficiently as possible not because God told them to, but through the power of the market.

Because what happens if you start a firm and you aren't cost minimizing? What happens? What happens is, in the short run, you might make money even if you aren't cost minimizing. But, in long run, you'll get driven out of business. Because if there's someone else who can produce more cheaply than you, they'll be able to charge a lower price and drive you out of business. Your price will end up above the long run equilibrium price if you're not cost minimizing.

So any firm that is not cost minimizing will get driven out of business. And the equilibrium will be a market where all firms are producing at the cost minimizing level. And that's why we get the high tech Figure 11-6, which is that the long run market supply curve is perfectly elastic.

Now this comes all the way back to what I talked about at the beginning of the last lecture. Remember I said, what determines perfect competition? Two things, the demand curve to the firm was perfectly elastic, and the supply curve to the market is perfectly elastic. And we talked last time about why the demand curve to the firm is perfectly elastic. Because with lots of firms, any given firm has a perfectly elastic demand. Now we've just derived why the market supply curve is perfectly elastic.

It's perfectly elastic at the cost minimizing point. If the price ever rises above that cost minimizing point, what happens? What happens if the price should suddenly rise above it? What happens? Firms enter and drive the price back down. If the price ever drops below that cost minimizing point, firms exit, and the price goes back up. So through the power of firm entry and exit, in the long run, you end up with a horizontal or perfectly elastic supply curve. And that's perfect competition. Questions about that? Yeah?

AUDIENCE: You said that if it's not profit maximizing, then it cannot [INAUDIBLE PHRASE]. But [INAUDIBLE PHRASE].

PROFESSOR: Yes, they are. So that's a great point. So what will happen is it's not that a firm will charge a lower price. I shouldn't have said it that way. It's that another firm will enter. And, by definition, the price will then fall, because the supply curve will flatten, and the price will then fall.

So if you're in there producing inefficiently and I say, hey, your firm sucks. I can come in and produce much more efficiently than you. I'll hop in, that will flatten the supply curve. The price will fall. At that price, if I'm not cost minimizing, I'll be losing money. So I'll leave.

That's a good clarifying question. So if we have a market with a bunch of guys in it, and one of them is not cost minimizing, well, that means someone else can come in. They'll, in the short run, expand the market. That will flatten the supply curve, drive the price down. At that lower price, the non-cost maximizing firm says, I'm losing money. So I leave.

The price goes back up, and it goes up and down and up and down until it settles at this point where costs are minimized. So perfect competition leads to a perfectly elastic supply and cost minimization. So that is our extreme. That is the theoretical point of no return. Of course, in reality, we never get there. In reality, there's no such thing as a purely perfectly competitive equilibrium.

Why not? Well, in the long run, supply is actually upward sloping. In the long run, market

supply will be upward sloping. And that's going to be for at least three reasons. So why is long run supply upward sloping in reality?

Well, in reality, it would be for at least three reasons. The first reason is that entry and exit are not free. There could be barriers to entry or exit. Even in the long run, there could be barriers to entry and exit. There could be features of the market which make it hard to leave or hard to join.

A classic example is something that we introduced a couple lectures ago, the notion of sunk costs, costs which even in the long run might be fixed. If they're large sunk costs, if one firm's incurred them, another firm is going to say, look, it's not worth it for me to get in.

If you have to build such a massive plant to produce your good, then it takes an unbelievable amount of capital, it's a huge investment, and once you're in, it's going to be really hard to drive you out, because you made that big investment, other firms might say, forget it. Dell made this huge investment in this huge plant. They're never going to leave having built that plant. That's virtually a sunk cost. So I'm not even going to bother entering.

So Dell can exist making some profit, maybe not too much profit. If they make too much profit, then another firm will build a plant. Well, as long as they're not making too much profit, they can make some profit. Because another firm says, you know what, I can't fight that battle. I can't build a plant that big. Or there could be other things like that.

That's sort of a natural barrier to entry. There are some artificial barriers to entry we see. For example, take med school. The number of slots to be doctors is limited by the physician profession. So even if docs make lots of money, which they do-- especially specialists-- you can't just compete and have new doctors enter. Because the number of slots are actually limited. You have to be licensed by an organization which is run by the people making all the money.

So if you're a doc, and you're in, this is a pretty good deal. You say, hey, let's have a system where new docs can't be licensed. They'll have to come to me, and I can charge whatever I want. That's a barrier to entry. We call that occupational licensing. We see that in lots of professions. Plumbing, taxi drivers, we see it everywhere. It's occupational licensing.

A second example, of course, is patents. And we'll talk about patents more in a few lectures. If I invented a new drug, and it's patented, nobody can sell that same chemical compound for 17

years. That's a barrier to entry.

There could be more informal barriers to entry. Let's say you're around Port Authority setting up these little shlocky stands, the ones we said were perfect competition. You've got yours, and the guy comes in next to you, you just beat the crap out of him. That's an informal barrier to entry.

You say, you come in, and you're going to get beaten up. The guy says, well look, if it's a big profit, it's worth getting beaten up. Or I'll hire protection if it's a big profit. But it's not that big of a profit, I'm not going to bother getting beaten up over it. So I'll let you make your profit. I won't come in.

So barriers to entry exist all over the place. And they're a big reason why we don't get a perfectly elastic supply curve in most markets because of things like patents or thuggery. So that's one example of why you don't get it.

Another example of why you might not get a perfectly elastic supply curve is that firms might differ. In particular, we have assumed critically, through the last lecture and this lecture, that firms are identical. We have assumed that firms are identical. But, in fact, of course, firms aren't. And one firm's cost minimizing production level might be different than another firm's cost minimizing production level.

Not all firms will have exactly the same cost minimizing production level. In particular, some firms may have a lower minimum average cost than others for a while. So it may be that as long as I'm producing less than x units, I have a lower minimum average cost than you do. But once I produce more than x units, my minimum average cost rises to above yours. Well, in that case, I might be able to make money for a while staying in. But then once I produce too much, I'm going to have to raise the price.

So to see this, there's a great example in Perloff which you see in Figure 11-7, where he talks about the international long run market supply curve for cotton. And he says, look, in Pakistan, you can produce cotton incredibly cheaply. This is a dated example, but in Pakistan, you can produce cotton incredibly cheaply. You can produce it at \$0.71 per kilogram.

So if the world demand for cotton is less than \$2 billion kilograms per year, or less than \$1.8 billion kilograms per year, then Pakistan would provide it all, and the price would be \$0.71. But let's say the demand is more than that. Well, Pakistan just runs out of cotton. They can't do

that.

Well, then you have to go to the next cheapest country. Well, the next cheapest country is Argentina. It costs a lot more to produce cotton there. And then comes Australia, Brazil, Nicaragua, Turkey, then finally the US. And then Iran is the most expensive.

So this is, effectively, an upward sloping supply curve. It's stepwise, but it's an upward sloping supply in the sense that as you want more quantity, the price goes up.

So if the market wants \$5 billion kilograms of cotton a year, then that means that the marginal producer is the US even though they're much less efficient than Pakistan. Because Pakistan hit a constraint. You have to go to that next less efficient producer. So that's taking an upward sloping supply curve. Basically you're getting an upward sloping supply curve, because you have constraints on how much any given firm can produce. Those constraints can make you move onto less efficient firms as you go on.

And so in a market, you'll end up, in reality, with a distribution of firms ranging from most efficient to least efficient. The most efficient would produce as much as it can at that efficiency level. But then some less efficient ones will get in the game as well just depending on where demand is. So you can see if you put in demand curves at different points in the supply curve, you get different prices. That's an upward sloping supply curve. That's a second reason.

And then a third reason-- these aren't a comprehensive list, but types of reasons why supply curves will slope up in reality-- is that input prices might rise as the market expands. We've assumed fixed input prices. I gave you an r and w , and I assume they were fixed. But, in fact, that might not be true. It might be that, in reality, as you want to produce more, you need to buy more of the input. Well, if you need to buy more of the input, and the input has an upward sloping supply curve, then you'll have to pay more to get more of that input.

So to see that, let's run through an example. Let's imagine that you want to produce something in the long run, and you need more labor to produce it. So as you produce more of it, you need more labor.

So now let's go to Figure 11-8. You're initially, in your firm, demanding L_1 units of labor at a the wage of W_1 . And let's say that at that point, at that wage, you're cost minimizing. That's the cost minimizing point, and you've got this flat supply curve. You're at that point.

Now let's say you want to produce more. Well, to produce more, you've got to go to the market

and hire more labor. If the supply curve for labor is upward sloping, if labor is not a perfectly competitive market and, therefore, is an upward sloping supply, they'll say, fine. If you want more workers, you've got to pay more. You've got to pay W_2 . You've got to pay more for your workers.

Well, think about what that does. Now, go to Figure 11-9. What that means is that as I produce more units, I have to pay more for the labor. Now let's start on the left-hand side figure. That says that if I'm producing little q_1 as a firm, my marginal cost is MC^1 , and my average cost is AC^1 . So I'm at P_1 .

Now if I want to produce more, if I want to produce q_2 , my average cost is going to be higher, and my marginal cost is going to be higher, because I have to pay a higher wage to workers. So that's going to shift me up to have to charge a higher price.

Now, I'm still cost minimizing. Given the wage the market gives me, I'm still cost minimizing. There's nothing non-cost minimizing about this. I'm still cost minimizing. But to cost minimize, I have to charge more. Because the market's charging me a higher price.

If you go back to solving our initial production decision, you'll see that because this is a higher wage, that's going to shift my cost function up. A high wage is going to make my costs higher. That's going to make my cost minimizing price be higher. So I'm going to shift.

I'm going to need to charge a higher price. That, itself, will also yield an upward sloping supply curve. So an upward sloping supply curve comes from the fact that as I produce more, I've got to pay higher prices for my inputs. That means I've got to charge higher prices for my outputs. So these are three examples of reasons why, in reality, we don't see a perfectly flat long run supply curve.

So once again, to review, because this is the end of this particular topic, to review where we are, the way it works is firms are given a cost function. Well, they choose a technology. That gives them a cost function. They enter the market. In the short run, they're stuck with that technology. So they decide to produce where price equals marginal cost as long as they're not losing more money than they've paid in fixed costs.

In the long run, firms come in and out until the point where every firm is producing efficiently. As long as entry is free, as long as there are not barriers to entry, every firm is producing efficiently. Then every firm is producing at a single point, which is where marginal cost equals

average cost. That yields a flat long run supply curve at the technological minimum.

In reality, supply curves slope up because there might be barriers to entry which leads to non-cost minimization. So this leads to non-cost minimization. And then there's two reasons, even if you're cost minimizing, you still could have capacity constraints, which is like our cotton example, or you could have upward sloping input supply. So that's why, in reality, we draw the upward sloping supply curves that we started with at the beginning of this course even with a perfectly competitive market.

So let's stop there. That's a lot of stuff to digest. We're going to come back next time and talk about why all this is crap, and firms don't really cost minimize or maximize profits or any of that.