

[SQUEAKING]

[RUSTLING]

[CLICKING]

**JONATHAN
GRUBER:**

OK, why don't we get started? Since I had some problems with the end of last lecture, I'm going to pick up right where things got a little dicey in the last lecture, and we're going to start over. So we're looking back at figure 7-3, which, if you remember, was the cost curves for our cost function $10 + 5q^2$. And you remember where it came from. This cost function we derived ourselves from the production function and wages and rental rates. We derived this cost function.

We're now graphing the cost curves that come out of this cost function, and we're talking about profit maximization. And we're talking about measuring profit. So we're talking about perfect competition. And remember, we said that profits are revenues minus costs. That means the profits per unit are revenues per unit minus costs per unit. Revenues per unit are price, and cost per unit is average cost. So profits per unit is price minus the average cost.

So we just go to the diagram. You see the height of the profit rectangle is price minus the average cost. So what do you do? You start by finding the point. So what are your steps here? Step one is you find the point where price equals marginal cost. That gives you your production level. That gives your optimal q^* -- we derived that last time-- of 3. We derived that last time that you want to optimize the price equals marginal cost. That gives our q^* of 3.

Then at that level, we compute the average cost. Average cost at q of 3 is simply going to be this cost function divided by 3-- so $10/3 + 5q^2/3$, or just $5q$, plus 15, OK? So the average cost at three is going to be \$18.33. That's going to be the average cost at production of three.

That means that the profits we're making for that third unit, the profits on our third unit, is the price-- profits per unit are the price, which is a fixed level of 30, minus the average cost, which is \$18.33. So profits per unit equals \$11.67. So that's the

profits. That's the height of the rectangle. The profits per unit is \$11.67. The price minus the average cost is \$11.67.

We're selling three units, so that means our total profit rectangle is 35, OK? So our total profits are 35, which is three units at a profit of \$11.67 each. And that's how we get that rectangle. Questions about that? OK.

So now, let's return to what we did last time. Let's imagine there's a tax of \$10 per unit. That would shift the cost curve to C equals 10 plus $5q$ squared plus $10q$. Remember, it's a tax per unit, plus $10q$. That would mean-- so that is illustrated in figure 7-4.

That means that the marginal cost curve and the average cost curve both shift up. Marginal cost is now equal to $10q$ plus 10 . We want to set that equal to the price to get the optimal q , and you solve this and you get a new q star equals 2 .

So now your optimal production level is two. You set your new marginal cost equal to the price. Marginal cost equals price. Marginal cost equals price at a new optimal quantity of two.

What's the profits there? Well, once again, profits per unit are just price, which is \$30, minus average cost. Well, what's the average cost at two? The average cost at two is going to be 10 over q -- so 5 . 10 over 2 is 5 , OK? Plus $5q$ because we're dividing this by q , $5q$, which is 10 , OK? Plus 10 , OK? 30 minus 25 , which equals 5 .

So your profits per unit, the height of that rectangle, has fallen from \$11.67 to \$5. So you're now making less profit per unit and you're selling fewer units. The height of the rectangle has shrunk. The width of the rectangle has shrunk. So the total profits have fallen from \$30 to \$10. You used to make profits of \$11.67 on each of three units. Now you make profits of \$5 on each of two units. So that tax has lowered your profits from \$30 to \$10, OK? Question about that? Yeah.

AUDIENCE: Is that by changing the cost [INAUDIBLE]?

JONATHAN GRUBER: Yes, exactly. Because you have a higher cost-- now, it's not just that your costs change. Your costs changing fed through to your production as well. So because your costs change, you produce at a different level and you made different profit

per unit, OK? All right.

Now, let's go to the other point I tried to cover last time that I want to get straight on, which is the shutdown decision. Now remember, we are, in the short run, shut down. We're in the short run. In the short run, there's no entry and exit. The firm can't literally leave, but it can just produce zero. That we call the short-run shutdown decision.

A short-run shutdown decision, you're still in the market. You still paid your fixed costs. You just produce zero. That's what we mean by shutdown, as opposed to exit, which is literally you take your toys and leave. This is you're still in the market. You've paid your fixed costs. You can't go anywhere-- remember, those costs are fixed in the short run-- but you just produce zero.

So now let's ask, for example, what happens if the price suddenly dropped to \$10. Let's say the price dropped from \$30 to \$10. Well, the fundamental profit-maximization rule never changes. We're not having the tax anymore. We're back to the original cost function. So original cost function is C equals 10 plus $5q$ squared, OK?

So the marginal cost equals $10q$, and our profit-maximization rule does not change. Our profit-maximization rule is that price equals marginal cost. So we set $10q$ to 10 , marginal cost equal to price. \$10 is the price. And we get that the optimal quantity is now one. We now want to produce one unit. That's the optimal quantity.

Well, what is our profits if we produce one unit? Well, profits equals revenues, \$10, minus costs. Well, what's our cost if we produce one unit? 10 plus 5 , 15 . So profits are negative 5 . I'm going fast, so jump in if I get the math wrong here, OK? Profits are negative 5 .

So you might think that's terrible, using money. You should get out of there, shut down. But the answer is, as we discussed last time, you should not shut down because shutting down still means paying your fixed costs. No matter what you do, you have to pay those fixed costs. So shutting down means production of zero.

What are your profits at a production of zero? Your profits at a production of zero are zero revenues minus the fixed cost of \$10. If you produce zero, what's your

cost? Plug 0 in here. Your costs are still 10. So if you produce zero, your profits are minus 10. So you should continue to sell even though you're losing money. Even though you're losing money, you continue to sell that one unit. You don't go to zero. Because at zero, you're even worse off.

This is the key thing about short-run shutdown decision. Yeah?

AUDIENCE:

What happens if you're in a position where, like, you are afraid you're going to keep selling one unit the next couple of months?

JONATHAN

Ah, but that's the key thing. Once you move to the long run, you can reoptimize your capital and you can exit. But in the short run, you've already paid that. You've already laid down your blanket, which cost you \$10. You might as well sell at a loss. If you just sell zero, then you've laid down your blanket for \$10 and gotten nothing out of it, OK? That's the transition to the long run.

GRUBER:

And so basically, that's how you think about the short run versus the long run. So what do we think about how this works? So basically, the shutdown rule, the way you think of the math of the shutdown rule is that basically you only want to shut down if your revenues are less than your variable costs because you got to pay your fixed cost no matter what. That's done. You're in the short run. That's gone. Forget it. You laid your blanket down in the bazaar. That's done.

So you only want to shut down if your revenues are less than your variable costs. You don't care about your fixed costs. That's done. You only shut down if your revenues are less than your variable costs.

Well, that's the same as saying you only want to shut down if your price-- variable costs are pq -- I'm sorry, your variable cost. You only shut down if your price are less than your average variable cost-- so divide by q . You only shut down if your price is less than your average variable costs. That's the shutdown rule. Shut down if the price you get is less than, on average, what you'll get unit for the units you sell.

So in our example, you never shut down. And why is that? Let's look at the math of our example. Well, in our example, what are variable costs? The variable cost is $5q$ squared. I'm going to put these down here so I don't have to-- the variable costs in our example, variable costs are $5q$ squared. So what are average variable costs?

$5q$. Those are average variable costs.

Now, remember that you want to shut down if variable costs are greater than the price. We can express the price in terms of q . Because remember, marginal cost equals price. So $10q$ equals p . So at the optimum, it is always true that p equals q over 10 .

That's a profit-maximizing condition, that p equals q over 10 . We can just substitute that in, and we have that average variable costs equals 5 times q . So average variable costs equals 0.5 times p at the optimum. Because we know that the optimum marginal cost equals price, we can plug that in and say, at the optimum, average cost is 0.5 times p . Yeah?

AUDIENCE: Do you mean q equals p over 10 ?

JONATHAN GRUBER: q -- let me see if I got that wrong. Yes, I'm sorry. q equals p over 10 . You're right. q equals p over 10 , my bad. Sorry. So the average variable cost is 0.5 times p . What's the shutdown rule? The shutdown rule is that price is less than average variable cost or price is less than 0.5 price, which can never be true. So you'd never shut down with this cost function.

To say it again, you only shut down if price is less than average variable cost. We've computed average variable cost in terms of price, and it's 0.5 times price.

Therefore, you never shut down. And you could see this when we actually go to-- so you can actually see this in figure 7-5 when we look at the firm's supply decision.

So what figure 7-5 does is show you, at every price, what the firm wants to produce in the short run. At a price of $\$10$, it wants to produce one. We showed you that. At a price of $\$20$, it wants to produce two. At a price of $\$30$, it wants to produce three, and so on. That dash-- that line that runs from $0, 0$ all the way up to $4, 40$, that's the marginal cost line.

At each point, price equals marginal cost, the optimal production decision. And you never shut down. As long as price is positive, you produce. Yeah?

AUDIENCE: So in this case, we don't set price to 10 minus p ? Because if it's 10 , then [INAUDIBLE] will be higher than 10 .

JONATHAN I'm sorry. I don't understand.

GRUBER:

AUDIENCE: So--

JONATHAN Price is always above average variable cost, so you'd never shut down, right? That's
GRUBER: the shutdown rule. This can never be true with this function. Price is forever above average variable cost. You can see that in this graph, right? The price equals the marginal cost, and that is always above the average variable cost, OK?

So the bottom line is you'd never shut down. But here's the other cool thing. After all this confusing math, guess what this line is, this marginal cost line. What else do we call that? That's the supply curve, people. We just derived the supply curve.

What's the supply curve? The supply curve is the relationship between the price in the market and the amount the producer desires to produce. Well, that's what this marginal cost curve is. So we've just shown you that the firm's supply curve is simply its marginal cost curve. So all you need to know to derive what a firm's supply curve looks like is what its marginal cost curve looks like, and you're done.

So literally, you take production plus input prices, and that gives you the supply curve. Why? Because production plus input prices gives you the cost function. Take the derivative of the cost function with respect to quantity. That gives you the marginal cost function. That's the supply curve.

So just like I could give you a utility function and a budget constraint and you got the demand curve, here, I can give you a production function and input prices and it gets the supply curve-- however, only under perfect competition, OK? This is only for the case of perfect competition. That's different from consumer theory. That was an everywhere rule.

Because I've given you an extra constraint, which is a constraint on the market, I've allowed you to derive the supply curve just like we easily derived the demand curve. So we have now derived a firm's supply curve. Questions about that?

So let's go back one more time. Just like if I give you on, say, a problem set or an exam a utility function and a budget constraint, you should be able to draw a demand curve. If I give you a production function and input prices and I tell you

you're in perfect competition, then you should be able to draw a supply curve because it's just the marginal cost curve, OK?

Now, that's the firm's supply curve. What we care about, actually, in the end, is the market supply curve, right? I've been doing little q . We care about big Q . That's what I drew in the lecture, was a big Q diagram.

So how do we get there? Well, all you do to get the market supply is to horizontally sum each firm's supply. So we can see that in figure 7-7, OK? What figure 7-7 does is take multiple identical firms and put them in the market. So for example, if there's one firm in the market, the supply curve is what we just drew. At a price of \$10, you get 10. At a price of \$10, you get one. At a price of \$30, you get three.

Now let's add a second identical firm. Well, that firm behaves the same way. At a price of \$10, it produces one. And at a price of \$30, it produces three. So the market now suddenly has twice as much. At a price of \$10, it has two. At a price of \$30, it has six. Now let's add a third firm. It behaves the same way. At a price of \$10, it produces one. At a price of \$3, it produces three. So each additional firm you add just literally shifts the supply curve down.

So in other words, the point is that a market supply curve, as you add more and more firms, the market supply curve is always more elastic than the firm's supply curve because any given increase in price calls for, from each firm, an increase in quantity. As long as there's multiple firms, that means a flatter market supply curve than a firm's supply curve.

So that's how we get the market supply curve. We solve for the firm's supply curve. We then just horizontally sum it over the number of firms in the market. Get the firm supply curve, horizontally sum, OK? Questions about that?

Now, now we have everything we need to go back to the first lecture, back to the future, back to the first lecture. And we can actually get to market equilibrium. In the first lecture, we started with a supply curve and a demand curve and got equilibrium. Well, we derived the demand curve a few lectures ago. We've just derived the supply curve. So let's do short-run equilibrium. Let's do short-run equilibrium, OK?

Now, the key thing is that-- so how do we do short-run equilibrium? Let's go through the steps. Step one, each firm picks a fixed amount of capital it's going to have in the short run. So each firm has a K bar. And based on that, it has a production function. Each firm has a production function, q equals f of k , K bar, L . And we have some input prices. We have some w and some r . Taken together, we can use those to create a cost function, which, in our example, was C equals 10 plus $5q$ squared, step one.

Step two. Step two, based on this, we can get optimal production levels from the fundamental profit-maximization rule, MC equals p . So this yields-- so here, we say $10q$ equals p . And this yields a supply function which is q equals p over 10 , and that's our supply curve. That's what I drew here. q equals p over 10 is the firm's supply curve. That's step two.

Step three is to create a market supply curve. Well, let's say there's six firms in the market. I'm just pulling this out of thin air. In a minute, we'll get to how many firms there are. But for now, in the short run, there's no entry and exit. So whatever number of firms I tell you, that's what's there. It's just given.

So let's say n equals 6 . Let's say there's six firms. And I'll come in a minute to where six comes from. But for now, let's just assume it. That means that the total market supply, big Q , is just equal to 6 little q , 6 little q , which equals $3/5 p$ or $6/10 p$. That is our market supply curve. Yeah?

AUDIENCE: Would firms that change [INAUDIBLE] who are out there? Would that change the overall--

JONATHAN GRUBER: The number of firms that are out there? Yeah, but once again, once they're in, they're done. They know there's six firms. So they're done. I'm going to come to this. You're talking about the long run, OK? Other questions? Yeah-- OK.

So that gives you your market supply. Finally, we go back to lecture one and the first recitation where you solved mathematically for equilibrium. We have a demand curve. I'm just going to make this up. Let's say the demand curve is Q equals 40 minus p . I just made that up to make the math easy.

Where does this come from? Well, you know where that comes from. You solved

where that comes from. That comes from consumer maximization, $48 - p$. So to get equilibrium, we just set $48 - p$ equal to $\frac{3}{5}p$, demand equal to supply. And we end up with $p = 30$, conveniently familiar number, which means that the total demand, big Q , is $48 - 30 = 18$, also convenient number. And I'll show you why.

The fifth step, each firm says, given this price, what do I want to produce? Well, we know that given a price of 30, each firm's little q^* is 3, right? That's what we solved. Each firm's little q^* is 3 given a price of 30. So each firm produces three units. How many firms are there?

AUDIENCE: Six.

AUDIENCE: Six.

JONATHAN GRUBER: What's 6 times 3? Supply equals demand. Six firms produce three each at a price of 30. That's 18. At a price of 30, people want 18, equilibrium. That's how it all works.

So leaving aside where the six came from, everything else you see here is just taking what we did and working out the math. And we get six firms in the market. They each want to produce three. That's the 18 that people want.

So to get equilibrium, you just need-- to get equilibrium, you need a demand curve, you need a cost function, and you need the number of firms in the market. You need a demand curve. You need the cost function and the number of firms in the market. Given that, you can solve for the equilibrium. Questions about that? Yeah.

AUDIENCE: So thinking about this in terms of the intersection of the graphs, would intersecting the demand curve with the whole market's supply curve--

JONATHAN GRUBER: Yeah, because it's the whole market demand curve.

GRUBER:

AUDIENCE: OK.

JONATHAN GRUBER: Yeah?

GRUBER:

AUDIENCE: $Q = 48 - p$ is the demand curve.

**JONATHAN
GRUBER:**

Q equals $48 - p$ is the demand curve. Other questions? Good questions. Those are good, clarifying questions. OK, so now let's ask where the hell six comes from. And where six comes from is the fact-- where six comes from is from the long run. So now we get to long-run competition, which is, in the short run, there's a certain number of firms in the market. But where those firms come from? Well, they come from the fact that each short run is a repeated exercise that makes up the long run.

So in the long run, perfect competition in the long run looks just like perfect competition in the short run-- full information, no transaction cost, lots of firms, with one difference. We're now going to allow entry and exit. Nothing else is going to change in the short run except we're now going to allow entry and exit.

So now, what-- and one other thing. I'm sorry. One other thing changes. Since all costs are variable, there's no shutdown decision. There's no issue of this shutdown rule. There's no more fixed costs. So if you lose money, in the long run, you shut down.

So all you have to worry about is profits. You don't have to worry about this extra shut-- so in the short run, we had an extra shutdown condition. You don't worry about that in the long run. You just worry about if you're making money or losing money. In the long run, if you're making money, you're in. If you're losing money, you're out.

So now let's ask how that pins down the number of firms in the market. This is complicated, so the best way to do it, I think, is an example. And think about our-- the fundamental rule I want you to keep in mind-- and the example gives you this rule. The simple rule is if there's profits to be made, you enter. If there's losses, you exit.

And these two things together imply the fundamental rule of competitive equilibrium-- that in the long run, profits are zero. Because if there's ever profit in the market, firms will come in. And if there's ever loss in the market, firms will exit. And that process will continue until profits are zero. So our fundamental conclusion is the long-run competitive equilibrium features zero profits. Now let's see why.

Let's go to figure 8-1. I want to talk about the market for personal computers. Let's

see. You guys were born what year? So you're now, like, 18. So you're born, like, 1990. This is about when you guys are born, OK? 1990 was a very interesting time for PC market. You all grew up in a PC world, a personal computer world.

But basically, in 1990, we're in a very different world. We're in a world that was dominated by giant mainframe computers. My graduation speaker from MIT in 1987 was Ken Olsen, the chairman of DEC, which was one of the premier manufacturer of giant computers. He gave the worst fucking graduation speech you've ever seen.

If you ever become famous and give a graduation speech, the number one rule is don't spend the whole speech talking about yourself. He spent the whole speech talking about how wonderful he was, how DEC was everything. And five years later, he was bankrupt. So take that, Ken Olsen. That's why you don't give a speech talking about how wonderful you are.

What happened? Well, let's talk about what happened. Let's talk about the PC market, OK? I hope Ken Olsen doesn't watch this.

[LAUGHTER]

Let's talk about the PC market. Let's talk about Dell, who was an early PC manufacturer. And what we're going to do in figure 8-1, there's going to be two side-by-side graphs. I'm going to go back and forth between them. If you're ever not clear about which graph I'm talking about, stop me. But I'll try to be very clear because you have to think about both these graphs in tandem.

On the right-hand side is the market for PCs, the market graph. On the left-hand side is Dell's cost curves. So the left-hand side graph is firm-specific Dell information. The right-hand side graph is the market.

In the short run, we were at a position where people wanted PCs, and there weren't many people making them. You guys, PCs were super cool. I mean, I guess, you're on laptops and everything now. But a desktop PC was an amazing thing at this point. Everyone wanted them, and there weren't that many people making them.

So at that point, Dell could-- the short-run market cost curve was SR1, and the demand was D. That was the short-- so we're starting on the right. We're going to go

from the right to the left. Start in the right. The initial market in 1990 was the short-run cost curve SR_1 and the demand curve of D .

At that intersection, the price was p_1 . Meanwhile, Dell, in the short run, had a marginal cost curve such-- now shift to the left-hand diagram. At a price of p_1 , Dell wants to produce little q_1 PCs-- marginal cost equals price, right? So find the intersection of that price with its marginal cost curve. Once again, this is very important. So stop me if I'm going too fast. Find where the price intersects the marginal cost curve on the left. Dell wants to produce little q_1 .

Well, what's its profits? Profits are price minus average cost. Well, its average cost was way below that. So Dell made this huge profit of the lightly shaded rectangle. That was its profit because it produced little q_1 . At little q_1 , average cost was way below price. So it made this lighted dot rectangle on profits.

Now what happens? Other companies see this and say, hey, we want in. What happens as more companies enter a market? The supply curve flattens. You're horizontally summing firm supply curve. Supply curve flattens. As it flattens, you move from SR_1 to SR_2 . So now we're back on the right.

In the market, as more firms enter-- and let's assume there are more firms identical to Dell. That's our perfect competition assumption, that each firm that enters is sort of identical. So more firms just like Dell enter. Gateway and all these guys start entering. And you shift to SR_2 , a flatter curve. That intersects the demand curve at the new price p_2 .

Now shift to the left-hand diagram. The price facing Dell falls to p_2 , but their cost function hasn't changed. They're still Dell. They're still the same underlying technology, paying the same wages and rental rates. So now, their marginal cost curve is the same. So at this new lower price, their production drops to little q_2 . Their production drops to little q_2 and their profits fall to the darkly shaded rectangle, π_2 . They make less money as firms enter.

And this process will continue until profits go to zero. Even at π_2 , another firm-- I forget what the third PC firm was. Some other firm will come in and say, wait a second. There's still money to be made. I want to come in. And it'll keep going till profits go to zero.

These are repeated short runs. So SR1 was a short run. Dell made a ton of money. Then we get to the next short run. Gateway enters. SR2's the new short run. Gateway and Dell still make money. So the third period, another firm's going to enter. And it's going to go till profits equal zero. Yeah?

AUDIENCE: Does this assume that the manufacturers or all the different companies produce the same amount of quantity and that they're the same quality, right?

JONATHAN GRUBER: Yes, absolutely. I'm going to come to those-- there's a large set of assumptions under this. But once again, think of Dell and Gateway selling their computers on rugs in front of the Eiffel Tower. It's all the same. You just go. You can compare equally. They're all the same thing. So it's that kind of market, OK?

Now let's think about poor, old IBM. Flip the page. IBM dominated-- despite Ken Olsen's claims, IBM dominated the big computer market. And so they were initially, on the right-hand side, in equilibrium with supply curve SR1. That intersected demand-- now we're on the mainframe market.

So in figure 8-1, we're in the PC market. Figure 8-1 is the PC market. Figure 8-2 is the mainframe market. The mainframe market, they're initially at supply curve one, which is very flat because there's lots of firms making mainframes. It intersects demand at a price p_1 , OK?

Now we go to the left. Well, remember, IBM has to produce where its marginal cost equals price. That occurs at production level q_1 . At production level of q_1 , it is losing money. IBM is losing money. It is losing that entire rectangle, the entire large rectangle.

Now, does IBM exit? It can't in the short run, but in the long run, the next period, some firms will exit. Some firms will say, yeah, we lost money in the short run. We're stuck. But in the long run, we just don't think this is a winning game. In particular, firms with high fixed costs will exit. They don't want to pay those fixed costs again in the second period. They know they have to build a new factory. They're not going to do it. They exit. Yeah?

AUDIENCE: What constitutes as a period? I know we said the long run is when--

JONATHAN
GRUBER:

I told you I can't tell you that. It's just some period of time more than a month, less than 10 years, OK? It's the period over which capital is variable. Think of it as a period of time over which you can build a new plant to make PCs. Think of it that way. So it's years, OK?

So what happens, then, somebody exits, says, I'm out of this. I'm shutting down the plant and moving somewhere else. That's steepens the market supply curve because now fewer firms are in the market. As that steepens the supply curve, the new intersection of supply and demand is at the higher price p_2 .

IBM stayed in the market. They just built a new plant. At p_2 , they now produce an amount q_2 . And at that combination, they are literally zero profit. That's the point where marginal cost equals average cost, which is a zero-profit point. Why is that zero profits? It's zero profits because, remember, profits are price minus average cost. What's price? Marginal cost. So when marginal cost equals average cost, profits are zero. You should be able to see that from the math we did before.

So what that means is that basically, when firms are losing money, they will leave and drive profits from below zero toward zero. So when they're making money, they enter and drive profits from above zero towards zero. When they're losing money, they leave and drive profits from below zero towards zero.

What does that mean? That means that our long-run perfectly competitive supply curve is in figure 8-3. It's flat. Long-run perfectly competitive market supply is flat at the price level. The market long-run supply is the point where marginal cost equals average cost or where supply equals average cost, OK?

And why is this? This is true because at any price above that point, if any firm tries to charge more than \$10, they'll be driven out of business. If any firm tries to charge below \$10, they won't make any money. Remember before, I said earlier-- you said, well, gee why don't firms just-- somebody asked, why don't firms just come in and take the whole market? This is why. Because cost is upward sloping. You don't want to come and take the whole market. You'll lose money. That's why firms don't want the whole market.

If marginal cost was flat, then you would then be undefined. You would want the whole market. But you don't. Marginal cost is rising. So you never want to produce

at a price above \$10 because no one will buy from you. You never want to set a price below \$10 because you'll lose money. Therefore, supply is perfectly elastic at a price of \$10. Yeah?

AUDIENCE: Thinking about this long-run logic then, wouldn't [INAUDIBLE]?

JONATHAN GRUBER: No, no, no. Once again, there's a lot of assumptions under this. But under these assumptions, you never make money in the long run because they shut down. There are short-run periods. So let's say what happened was-- let's go back to figure 8-2. Let's say that another firm shut down in the next period and suddenly IBM started making money. What would happen? Someone would enter and drive profits back to zero.

So under the assumptions we've laid out here, profits are zero. Because in the long run, we've achieved cost minimization. Firms minimize their costs where marginal cost equals average cost. Look at the average cost curve. Firms are producing at the minimum. The minimum of average cost is where average cost equals marginal cost. That is, competition has forced cost minimization. I'm just doing all sorts of mind-blowing stuff here for you guys. It can take hours to recover from this.

Competition forces cost minimization. Why? Because competition forces each firm to produce where price equals marginal cost, and it forces entry and exit until marginal cost equals average cost. Therefore, it forces each firm to produce at the most efficient point, which is where marginal cost equals average cost, because that is the point at where average cost is minimized.

So under competition, in the long run, every firm is cost-minimizing. They're producing the minimum of their average costs. And therefore, the supply curve is elastic, and it's defined purely by the minimum of average costs. That is, if you know-- here's a cheat. If you know you're in a perfectly competitive market and I give you a cost function, then you know-- I guess you still need the demand function too.

If I give a cost function and a demand function, you know you don't need to know how many firms are going to be in the market. You don't need to know little n . You know, in the long run, profits are going to be zero. So long run, firms are just going

to produce at the minimum of average cost. You find that, use the demand to find the price, and you're done, OK? So competition leads to cost minimization.

Questions about that? OK.

Now, you're all thinking, well, wait a second. We've got rich parents, many of us who make money in these businesses. I don't see zero profits. There's a stock market that's been booming. Where's the zero profits, buddy? I don't see zero profits. Well, the answer is twofold. First of all, remember, firms can make money in the short run in this model.

But that doesn't explain the stock market. The stock market's supposed to be forward-looking-- not months, but years and decades. So firms-- really, long run profit's zero. Why would stocks be expensive? Why would people want to invest in these companies?

And the answer is because these assumptions are unrealistic-- that this is an extreme version of the model that delivers some nice intuition, but doesn't apply to the real world. So I want to take the last few minutes to talk about the assumptions that we've made that don't really work in the real world to make this model work. It doesn't mean the model's invalid. We learn a huge amount.

And the key lesson from this model is competition pushes you towards cost minimization. Always think about these models as not delivering a level truth, but a directional truth. The directional lesson is this is why competition forces firms towards cost minimization, but firms won't actually necessarily get to zero profits.

And there's at least three complications. The first one is limited entry. I assumed that firms could costlessly enter and exit in this market. But in fact, that might be hard because, in reality, we have sunk costs, which I talked about last time. We have sunk costs-- or two times ago-- costs which, once paid, can never be recovered.

And therefore, firms might say, look, I don't want to get into this market because it's not like I can get my fixed cost back out next period. So if my fixed costs are building a building, next period, I can just sell that building to someone else. But my fixed costs are going to med school. I can't sell that med school degree. Therefore, if doctors aren't going to be profitable in the long run-- there'll be zero profit in the long run-- I'm going to go be a lawyer instead.

There are sunk costs, and those sunk costs create what we call barriers to entry. There are barriers to entry that come from costs that are sunk in the long run-- med schools. But there are other sorts of barriers to entry. Take our vendor market. Even in vendor markets, they're not perfectly elastic. One barrier to entry could be they could come in the middle of the night and steal your stuff and beat you up if you tried to enter the market. There's lots of reasons why entry and exit is not costless and easy. There's lots of barriers to entry.

Once there's a barrier to entry, this graph goes away. Because go back to figure 8-1. In the second short-run equilibrium, we've got these small profits, this small, dark gray rectangle, right? I then told you another firm would enter and squeeze those away, but another firm will enter only if what? Only under what-- somebody raise their hand and tell me. Under what condition would that third firm enter? Yeah?

AUDIENCE: Profits are greater than zero.

JONATHAN GRUBER: No, not profits are greater than zero. In the model, that was true. But in reality, what has to be true? Profits have to be greater than what?

AUDIENCE: Sunk cost.

JONATHAN GRUBER: Sunk costs. Raise your hand, people. I'm going to give you credit because you raised your hand regardless of who yelled out the answer. I'll assume you were right. Sunk costs. Profits have to be greater than the barriers to entry.

So in the short run, you enter if profits are greater than zero. In the long run, you only enter if profits are greater than the barriers to entry, which might not be true. There's always some cost to starting a firm. So profits will never really go to zero. They'll only go down to the barriers to entry. So that's problem one.

Problem two with this model, problem two is firms may differ. I've assumed-- and this was raised in one question. I've assumed identical firms here. But in fact, firms differ. And in particular, different firms have different cost functions. And with different cost functions, you can get some firms making long-run profits.

So for example, let's consider, in figure 8-4, the long-run market supply for cotton. This is from a textbook example. And this is from estimates that people have made

of the minimum average cost of producing cotton by country. So in other words, this is old news, but it doesn't matter. The country names don't matter. The example's what matters.

In the period of time this study was done, the cheapest place to produce cotton was Pakistan. You produce cotton in Pakistan for \$0.71 per-- I don't know-- kilogram, \$0.71 per kilogram, OK?

However, that was only for a certain amount. At some point, Pakistan ran out of cheap cotton and they had to start producing cotton using more expensive methods. Remember, marginal cost, at some point, has to slope up. So the point is marginal cost was flat for a while in Pakistan, then has to slope up. And suddenly, it becomes cheaper to produce cotton in Argentina at \$1.08.

And then it becomes cheaper to produce cotton in Australia at \$1.15. And eventually, if the price gets to \$1.56, it's finally cheap enough to produce cotton in the United States. The point is these flat segments represent the minimum average cost in each country. It's just, instead of making it a point, there's an amount of production they can do with that minimum average cost.

Now, let's say world demand for cotton was 1 billion kilograms per year. Then what would happen would be what I taught. You would have Pakistan-- competitors in Pakistan would compete, driving profits down to zero. And price would be \$0.71 per kilogram.

Now let's say, however, demand for cotton is 5 billion kilograms per year. Well, now, that intersects this supply curve. This is the world supply curve. There's the supply curve at \$1.71. So now, US producers are making zero profits because their marginal cost's \$1.71. The price is \$1.71.

But what about producers in Pakistan? They still make cotton. They make the first almost 2 billion kilos. But they're selling at \$1.71, and it's costing them \$0.71. So they just made profits. The point is that if there are firms which have rising costs and demand is high enough that the high-cost firms are actually producing, that higher price means profits for the low-cost firm. Yeah?

AUDIENCE:

Would this increase the price variable in Pakistan?

**JONATHAN
GRUBER:**

You could imagine those profits could then have a feedback effect in asset markets. And you could imagine the long run, in the very, very, very long run, as people buy land, that could dissipate the profits. That's a good point. We'll come back to that. But no, that's very long run, OK?

But for now, people have their land in Pakistan, and they're making their money on it. And so that becomes long-run profits. So long-run profits can come from heterogeneous costs. If some firms are particularly efficient in a multimarket firm, those firms can make money. That's a second feature.

Now, the third feature is, in some sense, the most interesting, which is input prices may not be fixed. And in fact, input prices, input prices-- ah. Input may have an upward-sloping supply. There could be an upward-sloping supply for inputs. We've assumed input prices are, everywhere, fixed, but that's not true. And in a few lectures, we'll come and teach you about that. But for now, let's recognize that inputs may have an upward-sloping supply.

So let's go through that for a couple minutes before we stop. Let's take a market in figure 8-5. This is the market for labor. Now, we've only been doing markets for goods so far in this course, and I'm sort of shortcutting by jumping to here. We'll spend a lot more time on this graph. But the bottom line is this is a graph of the amount of labor supply to the market.

In this graph, people are now the suppliers because they're supplying labor. Firms are the demanders. They're demanding labor. So what happens is you have supply curve of labor, and let's assume it's upward sloping. What I mean by that is as the wage goes up, people want to work harder. You want to work harder as the wage goes up. So it's an upward-sloping supply. It makes sense, right? Upward-sloping supply of labor.

Now let's imagine a firm suddenly wants to produce more. They used to produce little q_1 . Now they want to produce little q_2 . To do so, they need more workers. That represents a shift out in the demand for workers.

With upward-sloping supply, what does that do? Raises the wage. If you want more workers, you've got to pay them more. We didn't do that before. We assumed W was

a constant. But imagine if, to produce more, you have to pay more. What does that do? Well, you see that in figure 8-6.

Now, I used to produce-- the market equilibrium used to be at point E_1 with n_1 firms producing little q_1 . So that point big E_1 -- biggie, big comma E_1 -- is little n_1 . There was n firms producing little q_1 units per firm. And they were making these profits. Their profits were where price equaled marginal cost 1, now shifting to the left. p equals marginal cost one at little e_1 . So they were producing where marginal cost equaled average cost. So little q_1 at price p_1 meant that firms producing at little e_1 , which was the zero-profit point, OK?

Now what happens? Now the firm wants to produce more. Demand goes up. Firm wants to produce more. It wants to produce q_2 . You have a new long-run equilibrium with n_2 firms producing q_2 . Well, in that case, now if you want to produce q_2 , you're going to have to pay a higher wage. A higher wage means a higher marginal average cost. Higher marginal average cost means that you're now producing at a higher price and, therefore, the supply curve slopes upward.

Now, this is very different. One thing to think about-- I'll let you go. One last thing to think, think about the difference between this third case and the other two cases. In the other two cases, firms made profits. In this case, firms still don't make profits. So notice that profit's still zero, but the supply curve's upward sloping. So you don't need positive profits have an upward-sloping supply curve. So let's stop there. I've given you a lot to think about. And we will come back and talk more about this stuff on Wednesday.