

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 MIT OpenCourseWare at ocw.mit.edu.

JON GRUBER: All right, so today we are going to continue with our discussion of producer theory. And today we're going to move beyond the unrealistic case of perfect competition to the somewhat more realistic case of monopoly.

Now, we've been discussing perfect competition thus far as a form of market organization, and that makes sense in some context like fast food and other things. But in most contexts, we think perfect composition is not the way the world works. There's some limits on competition. And we think that markets-- many markets, many of the goods we consume have only a few firms. Operating systems or cars or a lot of things we consume, typically have only a few firms in them. So the most realistic model of markets would be one which accounts for the fact that there's less than an infinite number of firms, there's only a few firms.

That turns out also to be the hardest model. So what we do is we sort of iterate there. We started with one extreme, which is a competitive market where there's an infinite number of firms, that allows us to draw some interesting conclusions. Now we're going to reverse field and talk about the other extreme, monopoly, which is only one firm. We'll talk about monopoly markets with only one firm. Then we'll talk about oligopoly, that middle case, which is multiple firms.

So let's talk about monopoly, a market where there's only one firm. The key thing to remember for monopolies is they're no longer price takers, they're now price makers. Competitive firms are price takers. They were given a price by the market and they reacted to that. And as we saw in the long run, that price settled at the minimum of average cost. So basically, they were given the price which dictated production efficiency. They reacted to that in how much they produced. You got a flat long run supply curve.

However, in a monopoly market, we don't meet the conditions for perfect competition. In particular, one condition was that consumers had perfect substitutes between your good and other goods they could buy. That's not true in a monopoly market.

So if we think about that era, sort of maybe, I don't know, I don't have my computer history isn't as good. Maybe 10 years ago when Macs were sort of at the nadir of their popularity. And really you had to be a pretty high tech guy to be using Linux and things like that. That Windows had virtually a monopoly on operating systems. Not really, but virtually. Certainly unless you were a high tech guy who could do Linux and things like that, pretty much if you wanted an operating system, you got Windows. Windows pretty much had a monopoly. And that's may be as close as we've come in the modern economy to thinking about an example of monopoly.

In that case, Bill Gates had to decide-- wasn't given a price, he had to decide how much to charge for Windows. And that decision determined in turn how much he would produce. And that's exactly what we'll talk about today, is how does a monopolist decide both what to charge and how much to produce of their good?

So to do that, we're going to turn to a new concept. Well, not a new concept, but a different way of looking at something we've talked about before, which is marginal revenue.

If you remember, the profit maximizing condition that we derived when we started our competition lectures was that marginal revenue equals marginal cost. That was our profit maximizing condition, was that marginal revenue equals marginal cost.

Now we're going to go to-- in perfect competition. And the way we thought about this marginal revenue equals marginal cost, remember the logic was this notion of climbing this profit hill. You want to produce any unit.

Think of yourself as climbing this profit hill. You're deciding, do I produce the next unit? And I produce the next unit as long as the money I make off that unit exceeds what it cost to produce that unit. So I climb that profit hill. If marginal revenue is greater than marginal cost, I keep going up. At the peak, marginal revenue will equal marginal cost. Once I go beyond that peak, marginal revenue will fall down below marginal cost and I'll stop producing. So the notion is I climb this hill, which in the peak is dictated by marginal revenue equals marginal cost.

Now for a competitive firm, we said marginal revenue was just price. So in perfect competition, the rule was set price equal to marginal cost. But that was a particular case of marginal revenue. So for example, to see that, let's look at Figure 14-1. This is another way. Probably next year when I teach this, I'll put this in the lecture on competition. Just a way to think about

how marginal revenue is priced.

Remember, a perfectly competitive firm faces a perfectly elastic demand curve. They face a perfectly elastic demand curve. So they have to think about what the implications are of the marginal unit they sell. Well if they sell little q units, their revenue's a . If they sell one more unit, their revenue is b . And the marginal revenue is the height of that rectangle B times the base. The base is 1 because it goes from q plus 1. The height is p . So the marginal revenue is price. This is a pretty basic diagram. So marginal revenue is price for the perfectly competitive firm.

Now let's look at the monopoly case. The difference with a monopolist, as you see in Figure 14-2, is they no longer face a perfectly elastic demand curve. They now face a downward-sloping demand curve. Why is that?

Well remember, this is the graph for little q . The reason the graph for little q was perfectly elastic with a perfectly competitive case was not that we said the demand for the entire good was perfectly elastic. It's just that the residual demand facing anyway one firm was perfectly elastic.

Well now, a monopolist is the only firm in the market. So their residual demand equals total demand. Their residual demand equals total demand. So as long as total demand is downward-sloping, as we typically think it is, then they'll face a downward-sloping demand curve.

So unlike a perfectly competitive firm, which faces a perfectly elastic residual demand curve, a monopoly firm will face a downward-sloping market demand curve. They face the entire market demand curve. So this is demand curve. Now we're talking about big Q 's not little q 's anymore. Or we could say little q equals big Q . There's only one firm. So little q 's and the big Q 's are the same. So now the firm reacts not to its residual firm demand curve, which is flat, but the downward-sloping market demand curve.

Now, we're going to make one assumption here that's very important. We'll come back to this at the end of the lecture. We're going to assume that the monopolist can only charge one price for their good to all consumers. So most of this lecture we're going to assume a non-price discriminating monopolist. We're going to say that Bill Gates can't look at you and say, look, you look like you really want Windows. I'm going to charge you more than her. He can't do that. He has one price. He sells it to the stores at one price. So it's a non-price discriminating monopolist we're going to work with for the first 2/3 of this lecture. That monopolist has to set

one price.

Now, for that monopolist, for Bill Gates with Windows circa 10 years ago, let's think about his decision to produce another unit. He's originally producing at Q , big Q at a price p_1 . He's originally producing a big Q at a price p_1 .

If he wants to sell one more unit, he's going to have to lower the price. Because he now faces a downward-sloping demand curve. So if he wants to sell one more unit, he's going to have to lower the price to P_2 .

I have a big cockatoo at home. This is his feather.

So he's going to have to lower the price to P_2 . What's that going to do?

Well, on the one hand, what that's going to do is that's going to mean on that next unit he's going to make p_2 . So he's going to get the rectangle B. On the other hand, on all the units he was selling at p_1 , he now gets a lower price p_2 . So he loses the rectangle C. So the marginal revenue for this monopolist, the marginal revenue is equal to the rectangle B minus the rectangle C. The marginal revenue is rectangle B minus rectangle C.

Or alternatively, you could say that if we just write that out, write out what that is, that's p_2 minus p_1 minus p_2 times Q_1 . Or rewriting, we could rewrite this then as marginal revenue equals p plus $\Delta p \Delta q$, how much the price changes when you change the quantity, times the original quantity Q_1 .

Or one more time, for those of you who prefer calculus. If revenue equals p times q , and q is a function of p , then marginal revenue, differentiating that, is p plus dp/dq times Q . So marginal revenue is the price plus the change in price from selling another unit times the initial quantity.

Once again, the graphics are in this graph. The math is here. We just differentiate the revenue equation. This term is positive. Price is always greater than 0. But this term is negative because demand curves slope down. So there's now two effects. There's a positive effect, which is if I sell another unit, I make money on that other unit. There's a negative effect, which is to sell that other unit, I have to lower the price because I face downward-sloping demand.

So there's two effects a monopolist as he thinks about wanting to sell another unit. There's the money from that unit, but the lower willingness to pay for all previous units. And that's what makes a monopolist a little more interesting. We basically think of the monopolist as basically

having to work down the demand curve.

With a perfectly competitive firm, they don't have to work down the demand curve. The demand's flat to them. They can sell as much as they want at that price because they don't affect the price. They want to sell 10 of that price or a million at that price, it doesn't matter. Their demand curve's flat. Not true for the monopolist.

If the monopolist wants to sell more, he has to face the wrath of the market. And the wrath of the market is such that if he wants to sell more, he's going to have to lower the price to do so.

Remember, once again assuming he has to charge the same price to everyone. Assuming he's going to charge the same price to everyone, if he wants to sell more, he's going to have to lower the price to do so. There's different ways of [UNINTELLIGIBLE] intuition on this. I like to call this the poisoning effect. That's how I like to think about it.

That basically, if I want to sell another unit, I'm going to poison the money I made on all previous units. Because if I want to sell another unit, I have to lower the price. That's going to take away from the money I was making on my previous units. So it's sort of a poisoning effect is how I like to think about it. But you can have your own intuition for it.

And basically, this poisoning effect did not exist for the perfectly competitive firm because they couldn't affect the price with their action. They could sell however many units they wanted at that flat price and their actions did not affect that price. The poisoning effect only exists with monopolists because to sell that next unit, they have to lower the price.

Now, basically what that means is to find equilibrium for a monopolist, it's going to be a little bit trickier. And so let's go to Figure 14-3 and slowly walk through Figure 14-3. And we'll start walking through.

What the monopolist is going to want to do is draw a marginal revenue curve. With the perfectly competitive firm, marginal revenue curve was just a price, it was given to them. There was no marginal revenue curve. For a monopolist, there is a marginal revenue curve. So here I have a demand curve. The demand curve I've drawn here in this example. The demand curve I've drawn here is $q = 24 - p$. That's a typical demand curve, downward-sloping. As the price goes up, people want less of it. And now we're in market demands. Because remember, little q equals pq . There's only one firm in the market.

Now, here's the trick with monopoly: mathematics. The first thing you're going to want to do is you're going to want to invert this. You're going to want say, OK, that takes q , but what takes the price a monopolist is going to charge? The price a monopolist is going to charge is therefore going to be $24 - q$. That's going to be the price that the monopolist is going to charge.

Therefore, revenues, pq , is $24q - q^2$. So we inverted the demand equation. We do this because we want to write down a revenue equation. So that's the trick. This is the mathematical trick here is to invert this, so then you can write down a revenue equation. We then differentiate this revenue equation to get that marginal revenues equals $24 - 2q$. That's marginal revenues.

The next unit you sell, you make $24 - 2$ times the amount you're now selling. And that's the marginal revenue curve graphed here.

Now, basically what you see is in this case the marginal revenue curve starts at the same point as the demand curve on the y-axis and lies everywhere below the demand curve. Now that first fact, that it starts at the same intercept, is not always true. That is true because in this case we assumed a linear demand curve. With a nonlinear demand curve, marginal revenue curves can start at different points on the y-axis.

The second point about the marginal revenue curve always being below the demand curve is always true regardless of the function. The marginal revenue is always below demand. Marginal revenue is always below demand. So that marginal revenue curve will always be below the demand curve because of this poisoning effect.

Now, what I want to highlight here is this means there's a very important relationship between marginal revenue and the elasticity of demand. So let's take our marginal revenue equation and put it back in change terms. $p + \frac{\Delta p}{\Delta q} \times Q$. And let's multiply and divide by p .

So marginal revenue you can rewrite as $p + p \times \frac{\Delta p}{\Delta q} \times \frac{Q}{p}$. So I just took this second term, multiplied and divided by p , the second term. I just multiplied and divided by p .

The reason I did that is because that means you can rewrite this. This now starts to look like an elasticity expression. Remember the expression for elasticity. This looks like the inverse of

an elasticity expression. Remember what elasticity of demand was, $\Delta q / \Delta p$ times p over Q . So that's the inverse to the elasticity demand. So we can rewrite this as marginal revenue equals p times $1 + 1$ over the elasticity of demand. Marginal revenue equals p times $1 + 1$ over the elasticity of demand.

Think about what this means for a second. What is the marginal revenue in a perfectly competitive firm? Well, as a perfectly competitive firm, what's the elasticity of demand facing a perfectly competitive firm? Infinity. Perfectly elastic. So marginal revenue by L'Hopital's rule equals p . So for a perfectly competitive firm where elasticity is infinity, marginal revenue equals p .

Now instead, if we took a firm where the elasticity of demand was minus 1, the elasticity demand was minus 1, the marginal revenue would be 0. Why is that? What that says is, if you're a monopolist facing an elasticity of demand of minus 1, then you make no money by selling the next unit. Because these two effects exactly cancel. It turns out with elasticity of demand of negative 1, these two effects exactly cancel. Exactly what you make by selling one more unit is offset by how much you have to lower the price on all your previous units.

So an elasticity of demand of minus 1, marginal revenue equals 0. And as you can see as the elasticity of demand gets below minus 1, as it approaches 0 from below. As the elasticity of demand approaches 0 from below-- OK, I should have said perfect competition, I'm sorry, was negative infinity, not infinity. Negative infinity. As the elasticity of demand approaches 0 from below, then you're going to see that the marginal revenue-- as you approach 0 from below, marginal revenue is going to become negative.

So for example, if the elasticity of demand equals minus 0.5, then the marginal revenue equals minus p . So if this is minus 0.5, then this becomes minus 2. So marginal revenue equals minus p . You lose money. So as that elasticity of demand approaches 0, you're going to have a negative marginal revenue from selling the next unit. And why is that?

With a very inelastic good, you have to push the price down so much to sell the next unit that you lose money. Think about a very elastic versus very inelastic good. With a very elastically demanded good, to sell another unit you don't have to change the price much. Because the demand curve's very flat. So there's not much of a poisoning effect. dp / dq is small, or dq / dp is big. OK, this is the inverse. So dq / dp is big with elasticity, so dp / dq is small.

With a very inelastically demanded good, to sell one more unit you're going to lower the price

a ton, which is going to poison the revenues you get from selling that extra unit. So that's why marginal revenue will be higher, or will be a larger fraction of p as this elasticity becomes more negative. Yeah.

AUDIENCE: [UNINTELLIGIBLE PHRASE] elastic that means [UNINTELLIGIBLE PHRASE] irrespective of the price. So couldn't you just charge a higher price and get marginal revenue [UNINTELLIGIBLE].

JON GRUBER: No, because here's the thing. You should have already been charging that high price. The point this is the margin. So you go into a market for insulin. You say, look, these guys are going to die without it. I'm going to charge \$500,000 a shot. But now the question we're asking about marginal revenue. And at \$500,000, you sell to everyone who can afford it and everyone else dies.

Now you want to ask, what's the marginal revenue as you're trying to sell that 500,000 and first unit? Well to sell that, since it's inelastically demanded, if that person could afford anything like \$500,000, they would have bought it already. They can't. They can only afford \$400,000. We have to lower the price to \$400,000. You're going to sell one more unit at \$400,000 but lose \$500,000 on all those other units you were going to sell at \$500,000. The point is it's about the margin not the level.

Yes, monopolists make a huge profit when it's inelastic. I'll talk about that in a minute. But that next unit they're going to lose money on. So that's actually a good point to segue to now, let's talk about with this in place, let's talk about how monopolists maximize profits. Let's talk about monopoly profit maximization. Let's go to Figure 14-4. Profit Maximization for a Monopolist.

Now, this is a lot more confusing than perfectly competitive firms, so let's follow along here. This is a case the cost function here is $12 + q^2$. So I'm doing the cost function, which is $12 + q^2$. That's the cost function. And the demand function, as before, is $Q = 24 - p$. So that's what's graphed here.

Now, recall the rule that profit is maximized where marginal revenue equals marginal cost. Well, we know marginal revenue. We know marginal revenue-- we derived that above-- is $24 - 2Q$. What's marginal cost with this expression?

Well, marginal cost, differentiation of the cost equation, which is $2Q$. So the optimization term for a monopolist is going to where marginal revenue, which is $24 - 2Q$, equals marginal

cost, which is $2Q$. Or Q equals 6. That's going to be the optimal production level for the monopolist.

So we can see that graphically that's where the marginal cost curve hits the marginal revenue curve. If you go downward from that point, you get that the sales are 6 units. So marginal revenue equals marginal cost at 6. You should be able to see that graphically, it's just where the curves intersect. Mathematically I just did it here. It's actually pretty straightforward.

Here's the hard part. What's the price? We might say, well, gee, marginal cost and marginal revenue intersect at 6. I'm going to draw the dashed line over. That means the price is going to be 12. Why can that not be the price? Why is that wrong? What would that violate if the price was 12? If you tried to sell 6 at a price of 12? Yeah.

AUDIENCE: [INAUDIBLE].

JON GRUBER: It's not on the the demand curve. The monopolist still has to respect the demand curve. So monopolists in setting their quantity, gets the intersection of marginal revenue and marginal cost. But then in setting the price, they still have to read off the demand curve. They can't change consumer tastes. So they charge a price of 18. That's where you sell a quantity of 6. So monopolists it's a little bit trickier in a perfectly competitive firm. You set marginal revenue equal marginal cost to derive Q . But then to get p , you've got to go back and plug that into the demand curve. So with a Q of 6, I have my Q of 6. Well, what's the p ?

Well, to get that p , I've got to go back and plug this in here. At Q equals 6, p is $24 - q$, or 18. So I've got to respect the demand curve. The monopolist has to respect the demand curve. The monopolist picks both price and quantity, but he has to pick them such that you get a point on the demand curve.

And the way we solve it, is the monopolist chooses a quantity to set marginal revenue equal to marginal cost, and then chooses the price that's consistent with demand for that quantity. Questions about that?

Now one last thing. In the short run, we still have another condition for profit maximization, which is the shutdown rule. Remember the shutdown rule we talked about perfectly competitive firms in the short run, which is even if profits are negative, you might not shut down. You only shut down if price is less than average variable cost. So there's still the shutdown rule. So you only shutd own if price is less than average variable cost.

Now in this case, what's the monopolist profits? Well, the monopolist made a profit of 60. How do we see that? Well that's graphically the box, the rectangle, that's the difference between the average cost curve and the price they get. So they're charging 18.

Now once again, marginal revenue is gone. Think about marginal revenue like an imaginary concept. Marginal revenue isn't something that actually exists in the market. Marginal revenue is just something the monopolist draws to pick what they're going to do. But then it disappears.

What the monopolist cares about then is price. They're charging 18. Their average cost for that unit is only 8. So they're making a profit of 10 per unit on 6 units. So they're making a profit of 60. And what you can see, what you should be able to demonstrate to yourself is, if a monopolist sold 5 units. I'm sorry, selling 6 units. If the monopoly sold that seventh unit, what you'll be able to see is they would lose money on the seventh unit. Because yes, if they sold that seventh unit, what happened if they sold the seventh unit? Well then their price would have to be what if they wanted to sell a seventh unit? The price would have to be 17. The price would have to be 17. So to sell a seventh unit, they'd have to have a price of 17.

So basically at a price of 17, the price was 17. Then what would happen? Well, they'd sell one more unit at 17. That'd be good. But they'd lose \$1 on the previous 6 units, which is bad. So how much revenues would they make? What would be their marginal revenue? Well, the marginal revenue would be they make 17 minus the 6 poisoning effect. So marginal revenue equals 11. What's their marginal cost?

Their marginal cost is $2Q$. Marginal cost is 14. So they lose money. So you should be able to walk through this exercise yourself. You might say, gee, the marginal cost of that next unit is only 14. They sell it for 17. Gosh, they should do it. What you're missing is by selling it for 17, they've lost the dollar extra they make on each of the previous 6 units. And that poisoning effect makes it unprofitable to do this. And that's why the monopolists stop short of what would be the perfectly competitive outcome.

What would the perfectly competitive firm do? The perfectly competitive firm would set marginal cost equal to demand. And they would end up producing where marginal cost equals demand. So demand here is $24 - p$. Marginal cost is $2Q$. So they would end up producing where marginal cost equals demand at a much higher level charging a slightly lower price. So what you see is the monopolist ends up selling fewer units at a higher price. Questions about that? Yeah.

AUDIENCE: How does this work for Microsoft where their marginal costs are very low or nonexistent?

JON GRUBER: Well, then what would happen, if their marginal costs were very low or nonexistent. Think of that marginal cost curve then as being much, much flatter. It would intersect demand at a much higher quantity. Or it'd intersect marginal revenue at a somewhat higher quantity. Not that much higher. So if marginal cost is very low, they produce more but they make even more profits. So it's a good question actually, a good comparative statics exercise.

You bring that marginal cost curve down, what's going to happen? Quantity is going to go up, but not as quickly as profits are going to go up. That's why Bill Gates is the richest man in the world. That's what happens. You get really rich.

So basically, when you're a monopoly, low marginal cost you get really rich. But that's a great thought exercise to understand how this monopoly example works. Other questions about that?

So this is a good opportunity to introduce an important concept with monopolists, the concept of market power. What monopolists have, what Bill Gates has that my local McDonald's does not is market power. Or what he had, has less of now, is market power.

Market power is the ability to charge price above marginal cost. The summary statistic of how much power a monopolist has is how much they can drive their price above marginal cost.

When Bill Gates marginal cost dwindles to 0, his market power gets bigger. Price above marginal cost.

Now to think about this, remember the condition for profit maximization. It was that marginal revenue, which we wrote as $p \times (1 + \frac{1}{\epsilon})$ equals marginal cost. So we can rewrite this as $\frac{\text{marginal cost}}{\text{price}} = 1 + \frac{1}{\epsilon}$.

Now let's define the markup. Let's define the markup as price minus marginal cost, how much money you make on the next unit. You sell for p , you get marginal cost. It's money you make the next unit.

If you define the markup, $\frac{p - MC}{p}$, that's the percentage markup. It's how much you make on the next unit, the percentage markup. Then you can see that that markup equals $-\frac{1}{\epsilon}$. So the markup for a monopoly firm equals $-\frac{1}{\epsilon}$. This

comes to the question before about the insulin example's sort of confusing.

Here we see your intuition on insulin. The lower its elasticity, the more the monopolists can mark up their price. So your intuition is shown here. Yes, the monopolist will charge an incredible price for insulin. They'll still lose a lot of money if they try to raise that price, if they try to sell one more unit. But the first initial price they'll set will be incredibly high.

Basically, what is the constraint on Bill Gates? What is the constraint on Bill Gates? It's Steve Jobs. It's substitutes. The only constraint on a monopolist is the extent to which people can sub-- actually, let me go back, that's not a good example.

Let's [UNINTELLIGIBLE] Bill Gates circa 10 years ago. The constraint on Bill Gates circa 10 years ago was a mainframe or some other form of doing a set of-- or 20 years ago it was a typewriter. It was basically the fact that there was some other way to do what Bill Gates was letting you do.

If there was no other way to do what Bill Gates was letting you do, he would charge an infinite price. Clearly if there's some other way-- and also, elasticity of course, comes from substitutes or one substitute is just not to compute. So if Bill Gates tried to charge infinity for Windows, people just wouldn't own computers.

So the reason Bill Gates can't charge infinity, and the reason he can't charge infinity for insulin is that there's some elasticity of demand. People at some point will just stop buying. Either because they'll choose to use a typewriter instead or they just won't compute. They'll write by hand or something.

So basically at some point, there is some elasticity because there's a market demand curve. And basically what's going to determine how much market power the monopolist has is going to be how elastic it is. Basically, how close the substitutes are for that good.

If there's close substitutes, the monopolist won't be able to charge a very high markup. If there's not close substitutes as of Window circa 10 years ago, the monopolist can charge a very high markup and become very, very rich. Yeah.

AUDIENCE: But if there are substitutes for the market, then it's not a monopoly anymore.

JON GRUBER: No, no, this is the key thing. Substitutes for that producer. So basically, that's why I said Steve Jobs is not a good example. Because then it's not a monopoly anymore. But the typewriter is a

good example. That's a different good, that's a different market, different good that substitutes. So my point is any given good, insulin being an exception, but any good there's always something you can do instead. Insulin there is something you can do instead, you can be sick. There's always something you can do instead. We don't have only one thing in life. So the elasticity of demand is never perfectly inelastic.

It seems silly 10 years ago, but 20 years ago it actually was a legitimate decision whether to have a PC or not. A lot of people just didn't have computers. You could always just not have one. That gives you inelasticity of demand. So basically, it's important to recognize when we talk about substitutes, I'm talking about here not substitutes within the market, but substitutable activities, other things you could do with your money. And the more other things are you could do with your money, the less markup that Bill Gates can make on his Windows operating system. Questions about that?

OK, now we can ask, OK, gee, John, this is all good and interesting, but why did you just waste the last lecture and a half teaching us about welfare if you're just going to go back to producer theory? Well, the reason is because now we come to what the welfare effects of monopoly. And ask, what effects do monopoly have on society? And in fact, we can show you that there's a deadweight loss on society imposed by monopoly. To see that, let's go to Figure 14-5.

And here we can show the deadweight loss of monopoly. And here's the same example we were using. Demand is $Q = 24 - p$. Marginal cost is $2Q$. The cost function is $12 + Q^2$. So marginal cost is $2Q$. As we saw before, the monopolist chose to sell 6 units at a price of 18. 6 units at a price of 18.

The perfectly competitive firm sets demand, which is $24 - Q$, sets price, I'm sorry, equal to marginal cost. Well, price comes to demand curve as $24 - Q$. Marginal cost is $2Q$. So the perfectly competitive firm sets Q equal to 8. The perfectly competitive firm sets q equal to 8. They choose to sell 8 units at a price of 16. So you get the competitive quantity $Q_{sub c}$ is eight and the competitive price $p_{sub c}$ is 16. That's where graphically demand equals marginal cost. Or price equals marginal cost.

The monopoly firm sells 6 units at a price of 18. So what is the welfare effects of monopoly? What we see is we know that the competitive firm maximizes welfare. We learned that last time. We know that the best you can do is to sell 8 units at a price of 16.

What happens when you sell 6 units at a price of 18? What happens is consumer surplus falls

from A plus B plus C. So with perfect competition, consumer surplus is A plus B plus C. With a monopoly, consumer surplus falls to the area A. So you lose B plus C with monopoly.

Producer surplus under perfect competition was the area D plus E. Now under a monopolist, the producer surplus is equal to D plus E plus B. The monopolist, in this case, gained the rectangle B, but gave up the rectangle E. The consumer lost the rectangle B, that was a transfer to the monopolist. So there was a transfer of the rectangle B from the consumer to the monopolist. But C plus E have disappeared. They're a deadweight loss. They're deadweight loss because in the perfectly competitive equilibrium these are trades that would have made both parties better off.

That is, these are trades which socially should happen. They are trades where the value to the consumer exceeds the cost of producing that unit. Those seventh and eighth units are units-- so take the seventh unit. What's that worth to someone?

Well, it's worth 17. We can read that off the demand curve. That's a willingness to pay curve. People are willing to pay 17 for that seventh unit. What's it cost to produce? It cost 14.

So you have a unit which people want more than it costs to produce, yet it's not getting sold. That's deadweight loss. So monopolists induce deadweight loss because units that people value above their marginal cost doesn't get sold. Units people value above their marginal cost don't get sold. And that's because this poisoning effect. Because while it's socially optimal to sell those units, while society is better off, it's privately sub-optimal. From the monopolist's perspective, it's bad to sell that unit because of this poisoning effect. So basically, the monopolist is underselling, underproducing. In general, monopolists will underproduce goods. They'll sell too few goods because to sell the right amount would not be profit maximizing.

Because remember, what's the profits for the perfectly competitive firm? Profits for the perfectly competitive firm? Well, we know the profits of perfectly competitive firm. We know cost if they sell 8 units. We know the cost function is, the cost here is 12 plus Q squared. So if they sell 8 units, their costs are 12 plus 64, which equals 76.

Their revenues if they sell 8 units are 8 units times the price of 16. 8 units time the price of 16, which is 128. So what are their profits? Their profits are 52. So their profits are 52. The monopolist's profits are 60. So the monopolist is better off than the competitive firm would be.

The competitive firm would only make profits of 52. Obviously the short run. The long run they

make no profits. But in the short run they make profits of 52. The monopolist makes profits of 60. So the monopolist is better off than the competitive firm. The difference of course, is to do so they cause a social deadweight loss. Questions about that? Yeah.

AUDIENCE: Is that what the OPEC is doing right now?

JON GRUBER: I'm going to come to that actually. Time out on that. Because OPEC is more of an oligopoly. And we'll come to that when we talk about that in a couple of lectures.

But I want to talk about one more thing before we stop, which is I want to talk about the key assumption we made here, which was the monopolist could only charge one price to everyone. In fact, we know that's not true. In fact, we know in the world, there's a large amount of what we call price discrimination. There's a large amount of price discrimination. We know that for many goods, different prices get charged to different consumers.

If you ever tried to book an airline ticket the last minute, you know exactly what I mean. Basically, different prices in many, many contexts get charged different consumers. Everything from discounts for senior citizens, to higher price last-minute flights, to specials, two for one specials. People who buy two get a special price on the third, et cetera. There's all sorts of price discrimination just out there in the world. And in fact, there's very few goods that are sold at just one price. McDonald's hamburger is typically sold at just one price. They don't say like, fat people got to pay more for McDonald's hamburgers or something.

But many, many goods we buy in the real world are sold at many prices. And that's an example of a price-discriminating firm. And here's the crazy part. Here's the crazy part.

It turns out that a price-discriminating monopolist maximizes social welfare. A price-discriminating monopolist is as good as a competitive outcome. How can that be? Let's go to Figure 14-6.

Here's the price-discriminating monopolist. Now, the price-discriminating monopolist, what does he do? This is a perfectly price-discriminating monopolist, someone who can charge a different price to every single consumer.

Well, if you were a price-discriminating monopolist, perfectly price-discriminating monopolist and you could charge a different price to every consumer, what do you charge the first consumer? 24. What do you charge the second consumer? 23. Third consumer, 22. You literally charge them their willingness to pay.

If you're perfectly price-discriminating, then what you do is literally charge every consumer exactly their willingness to pay. You say look, I know your willingness to pay function. Your willingness pay function is $p = 24 - Q$. That's your willingness to pay function. So I'm going to literally charge you that. I know that about you, it's stamped on your head. So I'm going to say, ah, you're willing to pay 24 for the first unit, 23 for the second, et cetera.

In that case, what will the perfectly price-discriminating monopolist do? Will they stop at 6 units? No, they won't. Because for that guy, there's no poisoning effect. There's no reason to stop at 6 units.

That seventh unit, as we just did the math, there's money to be made on that seventh unit. Because that seventh unit is worth 17, but it only costs 14 to produce. So the perfectly discriminating monopolist will sell it at 17. Likewise the eighth unit, people willing to pay 16 and it cost 16 to produce. So they'll sell it or not. They're basically indifferent. So we typically say they'll sell it.

The point is, the perfectly price-discriminating monopolist will work all the way down the demand curve to the competitive outcome. They will move to the competitive market outcome because there's no poisoning effect. There's no reason not to. No reason not to sell as many units. No reason not to climb the same hill the competitive firm climbs and sell any unit where the price exceeds the marginal cost.

Well, what's interesting is let's ask what's happened to social welfare with this perfectly price-discriminating monopolist. Well, consumer surplus is what? What's consumer surplus with the perfectly price-discriminating monopolist? Somebody raised their hand. Yeah.

AUDIENCE: Zero.

JON GRUBER: Zero. Why is it zero?

AUDIENCE: Because they're charged exactly how the value is.

JON GRUBER: Exactly. Consumer surplus is defined as willingness to pay minus price. But your price is set equal to your willingness to pay. So by definition, consumer surplus is 0. With a perfectly price-discriminating monopolist, there's no consumer surplus. But what's producer surplus? Same person, what's producer surplus?

AUDIENCE: Everything else.

JON GRUBER: Everything else. $A + B + C + D + E$. There's no deadweight loss. You get exactly the same social welfare as you got with perfect competition. It's just divided differently. With perfect competition, consumers got $A + B + C$. Producers got $D + E$.

With a perfectly price-discriminating monopolist, the monopolist gets everything. But the total shaded area is the same. So really fascinating because here we have the ultimate screw on consumers. We think about competition as being the best thing for consumers. Lots of firms selling goods at a competitive market where you can shop and do what's best for you. It's not surprising intuitively that that's the best thing for society.

What's very surprising intuitively is having a producer who can screw every single consumer out of every penny they value something is equally good for society. And why is that? That's because we've made a particular assumption, which is social welfare is the sum of producer surplus and consumer surplus. The linear sum.

Since it's a sum, we don't care in that function who gets the dollars. We just care about the total amount of dollars, the total size of the pie. We don't care who gets what slice of the pie, we just care about the total size of the pie. And the total size of the pie is the same with a perfectly price-discriminating monopolist and a competitive firm.

What this highlights is that that's a pretty stupid way to think about social welfare. Clearly, we don't feel the same way about a market where people get everything they want and a market where people-- all they're willing to pay is sucked out of them by a greedy monopolist. Clearly we don't. And that's why we're going to need to think more richly about equity and think more richly about the division of resources in society. Because it turns out that you can have equally good outcomes from an efficiency perspective that are very, very different from an equity perspective. And this is the first example we'll see of that.

What we'll do when we get towards the end of the course lectures, like 23, 24, lectures like that, we're going to start talking about equity. And what are different rules we can think of for dividing this pie that might give us a different answer? OK, questions about that?

All right. OK, so anyway, we'll stop here then. Let's remember that perfectly price-discriminating monopolist is obviously also a silly concept just like a perfectly competitive firm's a silly concept. What we're going to do next time is come back and talk about price

discrimination in reality and what firms do to try to approximate this golden outcome.