[SQUEAKING]

[RUSTLING]

[CLICKING]

JONATHAN GRUBER: So today, we are going to move on to a new topic. Once again, to remind you of the path of how we got here, within producer theory, we said that it's harder than consumer theory because you don't only do the standard optimization. You do the consumer theory, but you need a third condition. You need a market structure condition to pin down the equilibrium.

For the last few lectures, we relied on a particular market structure, which was perfect competition. That allowed us to show that in perfect competition, price equals marginal cost. And that allowed us to solve the system for how firms choose their quantity and prices in perfect competition. But as we discussed, that's an incredibly extreme assumption and pretty much unrealistic for almost any market.

There are some markets that are more competitive than others, but there's really no such thing as a perfectly competitive market. So what we're going to do now is move on in the next few lectures and talk about alternative market structures and how they change the conclusions we draw. This is going to be important for two reasons.

First of all, it's going to be important because it's more realistic. Second of all, it's going to be important because having established the first fundamental theorem of welfare economics, which is the perfect competitive outcome is the best outcome, we're now going to talk to you about why that doesn't work in reality and why, in fact, governments aren't just agents of chaos but can actually do good things in improving potential market failures.

So by introducing these alternative market structures, we'll also introduce the concept of a market failure, which is markets not working to deliver the efficient outcome. So we're going to start today by talking about monopoly. So we're going to go from one extreme to another.

The one extreme, we started with was a perfectly competitive market, with a large number of firms competing to deliver identical products. We're going to move all the way to the other extreme, which is monopoly, which is not the game, but rather the name for a market with one firm that delivers the goods for the entire market. So moving from many firms competing to one firm dominating the market.

And the key thing about monopoly, the key thing that's going to change is that firms are going to be price makers, not just price takers. Remember, in perfect competition, the price was given to you from god. It came from some market equilibrium. You don't know where it came from. You just know there's a price and you couldn't affect it.

Well, that's kind of silly. Firms can obviously choose the prices they set. So now, when we move to monopoly, we'll actually let firms be price makers as well as price takers. Now, in fact, once again, monopoly is very extreme. It's not an extreme perfect competition. Monopolies do exist, even if perfectly competitive markets don't.

But it's still a relatively small share of reality. After a couple of lectures of monopoly, we'll then turn to the more realistic case-- and my favorite word in economics-- oligopoly, which is the case of several firms competing, where there's competition, but it's less than perfect competition. And we'll come back to that.

But for the next two lectures, we're going to move to the other extreme, which is monopoly. And we're going to talk today about how monopolists maximize profits. How do monopolists maximize profits? Well, to do that, we're going to turn to a concept that we've already talked about, but it's been of meaningless so far, which is the concept of marginal revenue.

The marginal revenue is the revenue earned on the next unit you sell, which in perfect competition wasn't interesting. It was just the price. But it's going to get more interesting today. So we're going to talk about what happens. So let's go back and think about a perfectly competitive firm and think about what their marginal revenue is.

So in figure 11-1, you have a perfectly competitive firm. The firm faces a perfectly elastic demand curve. A perfectly competitive firm. This doesn't mean the market demand is elastic, but each firm in perfect competition faces a perfectly elastic demand curve. So if they sell little q units, they make revenue A.

If they sell q plus 1 units, they make revenue A plus B. They add B. Well, we know what B is. The width is 1 and the height is the price. So the marginal revenue from selling the q plus 1 unit is the price. It's just an illustration of what we learned before.

Monopoly is going to be different. And in particular, monopoly is going to be different because we're going to make a critical assumption for how monopolists behave. We're going to assume that monopolists cannot price discriminate.

That is, when a monopolist sets a price, they set that price for all units they sell. So a monopolist sets the price. They're a price maker. But it's not quite free price setting because we're not going to let them set a different price to me and you. When they set a price, they set the price for everyone.

So we're going to assume no price discrimination. We'll come back to that. But for now, we're going to say that-- we're going to let the monopolist set its price. But critically, it's one price for everyone. And you'll see in a minute why this is important. So to understand that, let's go to figure 11-2.

11-2 is for a monopoly market. All that's changed here is that, now, firm demand is market demand With perfect competition, firm demand was perfectly elastic because there are many firms. But now firm demand's market demand. And unless the market itself features perfectly elastic demand, the firm will now face a downward-sloping demand curve.

So now, little q is big Q. You notice on the x-axis, I moved to big Q, which we use to refer to the market. Well, the monopoly market's little q is big Q. There's one firm, so the firm is the market. And that firm faces a downward-sloping demand curve.

So let's say that firm, initially, wants to sell Q, Q1. I'm sorry-- wants to sell Q. If it wants to sell Q, it has to respect demand. So if it wants to sell Q, it's going to set the price at P1. A price of P1 allows them to sell Q units. Remember, they get to set the price now.

So a price of P1 allows them to sell Q units. Now imagine you want to sell one more unit. What happens? Well, two things happen. And in particular, if they want to sell one more unit, they have to lower the price. Why? Because demand is downward sloping and there's just one price.

So if the monopolist wants to move from selling Q to selling Q plus 1, it has to lower the price from P1 to P2. There's no choice. It has to be on the demand curve. So when the monopolist wants to sell two units, the marginal revenue they get once they sell one more unit, they get the amount B, which is sell one more unit at the price P2, but they lose the amount C.

They get the amount B, but they lose the amount C. Or in other words, what does the monopolist get from selling one more unit? Well, the monopolist sells one more unit, the monopolist gets P2 minus P1 times-- I'm sorry-- P2 minus P1 minus P2 times Q1. The monopolist gets some money from selling an additional unit but loses some money because they had to lower the price on all its previous units.

Once again, stop and think for me for a second. You might think that doesn't make sense. Why does it have to? Well, it has to because it can only set one price. So if it wants to sell more, it's got to set a lower price, which means it's going to make less on all its previous units.

So basically, when the monopolist sets the price, there's now a trade off. The trade off is, if he wants to sell more, he has to lower the price. But in doing so, the monopolist makes less money on the previous units they sold. And that's because they set the same price for all consumers.

This is the key intuition that drives our analysis of monopoly. I like to refer to this as what I call the poisoning effect. I call it the poisoning effect-- it's just the term I like to use-- because the idea is, if I want to sell one more unit, I have to poison all my previous sales by lowering their price. Because I can't sell another unit without a lower price.

And if I lower price, I make less money on my previous sales. Don't think of this as changing your price day by day. Think of this I have to set one price. Yeah?

AUDIENCE: [INAUDIBLE] [? on the board ?] [? next to ?] [INAUDIBLE] marginal revenue?

JONATHAN GRUBER: Yes, this is the marginal revenue.

AUDIENCE: Was P2 just multiplied by [INAUDIBLE] or--

JONATHAN GRUBER: Yeah, P1 minus P2 times Q1. It's in parentheses. So basically, the marginal revenue is, if I want to sell one more unit, I get P2 for that unit. That's what I get for that unit. But I give up what I was making on all the old units, the difference in the price that I was making on all the old units.

And that is basically the definition of marginal revenue in the monopolist's case. So in other words, more generally, if you write revenue as P times Q, then marginal revenue, you can write-- if you write P times Q, then marginal revenue is dR/dQ. OK, if you take marginal with respect to Q, it's dR/dQ.

Well, that equals P plus dP/dQ times Q. So when you want to sell one more unit-- the marginal revenue, there's a positive term, which is I sell one more unit. But there's a negative term, which is I had to lower my price to sell that unit. So when I'm setting my price, I have to trade these things off.

I have to trade these things off, which is basically that I have to trade off the fact that I'm going to sell more units, but I'm going to make less money on each unit. Why didn't the competitive firm face this trade off? Why was this trade off not an issue for the competitive firm mathematically? What's different for the competitive firm? Yeah?

AUDIENCE: dP/dQ is 0.

JONATHAN GRUBER: dP/dQ is 0 for the competitive firm because it's a price taker. It can sell as much as it wants at the market price. So for the competitive firm, marginal revenue is just price. But for the monopolist, that's not true. To sell more, it has to lower the price.

So we can graph a monopolist's marginal revenue curve. And we'll do that in figure 11-3. So let's go through the math here. Let's do an example. Imagine we have an example where demand is described by the equation Q equals 24 minus P. Just a standard, well-behaved demand equation. You know how to derive that from utility.

Imagine that's the demand equation. Well, the first step to solving the monopoly problem is let's rewrite this as P equals 24 minus Q. Makes the math a little easier. And now let's multiply both sides by Q to make it a revenue equation. So revenue, which is P times Q, equals 24Q minus Q squared.

That's the revenue that the firm earns. It's P, which comes from demand times the amount of units sold. Differentiating this, marginal revenue is 24 minus 2Q. That's the marginal revenue curve for that firm. And you can see this in figure 11-3.

Here we graph the demand curve. That's the blue line. The demand curve is just a standard demand curve, where you can have 24 units per day at a price of 0. And you want 0 units per day at a price of 24. And then the marginal revenue curve is everywhere below the demand curve.

It has the same intersection on the y-axis, but it's everywhere below the demand curve, is the marginal revenue curve. Now, the fact that the marginal revenue curve starts at the same y-intercept is not general. It's only because this is linear. In a nonlinear case, the marginal revenue curve can start at a different y-intercept.

But the fact that marginal revenue is everywhere below demand is general. Marginal revenue is always below demand because of this poisoning effect. Marginal revenue is always below demand. Questions about that? This is not a simple intuition, so questions about how this works.

OK. Well, armed with that, let's talk about how monopolists maximize profits. Let's go to figure 11-4. Figure 11-4, I'm going to now add a cost function. Let me say that C equals-- I'm just going to assume C equals 12 plus q squared. Once again, you know how to derive that.

If I gave you production function input prices, you could derive that But I'm just going to assume C equals 12 plus Q squared. So in that case, what do we know? We can graph the marginal cost. Marginal cost is this 2Q. That's the marginal cost line in figure 11-4. Marginal cost is 2Q.

We can graph average total cost, and average variable cost, and average fixed costs. Those are all things we learned to graph before. We can put those all on the graph. Now we can solve for the monopoly-profit-maximizing decision. Remember, profit maximization always happens.

Regardless of market structure, profit maximization happens when marginal revenue equals marginal cost. Remember, just to remind you, profits equals revenue minus cost. So to maximize this, we say marginal revenue minus marginal cost equals 0. That's how we maximize profits.

So if profit is revenue minus cost, then profits are maximized when marginal revenue minus marginal cost equals 0. So the profit-maximizing condition is marginal revenue equals marginal cost. That's true with any market structure. That's always profit maximizing.

Marginal revenue here is 24 minus 2Q. Marginal cost is 2Q. Oh, I should have made this a big Q. Well, it doesn't matter because little q is big Q. But to make life easier, let's make that a big Q. So what this says is that Q star equals 6, that the optimal monopoly quantity is 6.

And how did I do this? Once again, not too hard from the perspective of your awesome math skills. I simply took the demand curve, created a revenue equation, differentiated that to get marginal revenue, have a cost equation, differentiate that to get marginal cost, set them equal, and you get Q star. That's all you got to do.

Now, what is the price? I solved for Q, but it's not obvious here how you get back. So how do you get price? So what price does the monopolist set? Yeah?

AUDIENCE: 18.

JONATHAN GRUBER: Why 18?

AUDIENCE: Because they have to obey the demand curve.

JONATHAN GRUBER: Exactly. So you skipped the usual thing where someone gives the wrong answer. People often say, well, look, marginal cost equals marginal revenue. The quantity's 6 and the price is 12. That's the typical answer. You say, well, look, if cost equals marginal revenue at 6, then you go over to the y-axis and it's 12.

And why is that wrong? That's wrong because the monopolist still has to respect the demand curve. A monopolist cannot charge a price that people aren't willing to pay. At a price of 12, they're not on the demand curve. People want way more at a price of 12.

To be on the demand curve to get to price, we go back to the demand equation. P equals 24 minus Q. Well, I just solved for Q star. So P star equals 24 minus Q star equals 18. And this is where people get screwed up in the math here. Remember, the monopolist has to respect the demand curve.

They're a price maker, but a price maker within the constraints the market gives them. And the constraint the market gives them is the demand curve. So the profit-maximizing point is a price of 18 and selling 6 units. And you could see, at that point, what are their profits?

Well, the profits are price minus average cost. Well, at 6 units, average cost is 8. So they are making 10 per unit. Times 6 units is 60. And that's their profits. Now, you might say, gee, consumers are willing to buy a lot more units.

Think about the seventh unit. Say the seventh unit-- well, this is what's weird to me. If you look at the seventh unit, demand is above cost. Consumers' willingness to pay at the seventh unit is 17. If you had a price of 7-- I'm sorry-- if you just wanted to sell 7 units, consumers are willing to pay 17.

And what's the marginal cost at 7 units? Well, it's 14. So this is weird. Price exceeds marginal cost. So why don't they produce that seventh unit? Because of the poisoning effect. To produce that seventh unit, they'd have to lower the price to 17 to sell it-- not to produce it, to sell it.

They lower the price to 17, they make less money on the first six units. Indeed, what are their profits if they sell 7? Well, if they sell 7 units, what are their profits? Well, they sell 7 units, profits equals 7 units at a price of 17 minus costs. What are their costs if they're producing 7 units?

Well, it's 12 plus 7 squared. So it's 12 plus 7 squared. So that's 12 plus 49. And if you do the math, you get profits are 58. They get lower profits selling 7 units. This is crazy. Think about this. We have units where people are willing to pay more than it cost to produce, and yet to sell them, the monopolists would actually lower their profits.

That's the power of this poisoning effect. The reason they lower the profits is because they can make more money by selling less at a higher price. Yeah?

AUDIENCE: [INAUDIBLE]. [? Basically, ?] let's say they sold at 18 for the first 6 [INAUDIBLE] 17 [INAUDIBLE]

JONATHAN GRUBER: What I said, they can't do that. I just ruled that out. We're going to come back to that. But that's the key, is don't-- the way I described it isn't the way you think of the decision making. Don't think of them saying, I'm going to set this for this. Just think, I want to sell x units. What do I have to charge?

So they're sitting back, saying, how many units should I produce? Well, if I produce 6, I can sell for 18. I make 60. If I produce 7, I sell for 17. I make 58. I'm going to produce 6. That's the way they think about it. So that is the power here. Now, what does this mean?

This means that monopolists, because they can set price, have something we call market power. Monopolists have market power. What do I mean by that? Market power means that, basically, monopolists can make money even in the long run. To see this, let's think about the marginal revenue expression.

So marginal revenue equals P plus dP/dQ times Q. Now, take that equation and take the second term and multiply and divide by price. So I want to take the second term, I want to rewrite it as P plus dP/dQ times Q over P times P.

I just multiplied and divided by price. That's all I did. Can anyone tell me what this term looks like? Yeah? The elasticity of demand but inverted. This is the inverse of the price elasticity of demand. So I can rewrite marginal revenue.

I can rewrite marginal revenue as P times 1 plus 1 over the elasticity of demand. All I did-- no trick here-- well, there's a little math trick-- wrote marginal revenue, multiplied and divided the second term by P, kept the first term, then realized this is the inverse of the elasticity of demand. So I can write marginal revenue as P times 1 plus the inverse of the elasticity of demand.

Why is this important? Because what this says is the ability of monopolists to make money depends on how elastic is the demand they face. So what would happen here if the elasticity of demand was negative infinity? What would marginal revenue be? If the elasticity of demand is negative infinity, what would marginal revenue be?

Come on, you guys. You're MIT students. Don't be shy. Raise your hand. Yeah?

AUDIENCE: P.

JONATHAN GRUBER: P. What does that sound like? Marginal revenue equals P. We've heard that before. That's the perfectly competitive case. Well, in a perfectly competitive case, firms face perfectly elastic demand. So if demand in the market was perfectly elastic, monopolists would behave just like perfectly competitive firms.

On the other hand, as the elasticity falls or rises, or falls in absolute value but rises from negative infinity towards 0, monopolists start to be able to charge more than-- monopolists start to have marginal revenue above the price Monopolists start to have marginal revenue above the price. So we can rewrite this.

Let's rewrite this condition for profit maximization. Profit maximization is that marginal revenue, which equals P times 1 plus 1 over the elasticity of demand, equals marginal cost. That's our profit-maximizing condition. So in other words, we can write that marginal cost over price equals 1 plus 1 over epsilon.

Marginal cost over price equals 1 plus 1 over epsilon. Now let's define something I'm going to call the markup. The markup is how much more monopolists can charge than perfectly competitive firms. I'm going to find the markup as P minus MC over P. That's what we call the markup.

That is how much more can monopolists charge than competitive firms. It's the markup. Well, the answer is equal to-- if we do the math, the answer is equal to minus 1 over the elasticity of demand. If you do the math, the markup is minus 1 over the elasticity of demand.

In other words, the profits monopolists can make are inversely proportional to the elasticity of demand. Why does that make sense? Why does it make sense that, the elasticity of demand, as markets get more inelastic, monopolists make more money? As markets get more elastic, monopolists make less money. Why does that make sense? Yeah?

AUDIENCE: [INAUDIBLE] the insulin example for [INAUDIBLE].

JONATHAN GRUBER: Yeah, insulin, the perfectly inelastic good.

AUDIENCE: Yeah. No matter what the [INAUDIBLE] charge, you have to buy it.

JONATHAN GRUBER: Exactly. You have to buy it. So given that, they can charge whatever the hell they want. This is the answer to the question you might have first thought or someone might think is, well, gee, if a firm is a monopoly, why don't they just charge whatever they want? Why do monopolists have to-- if they're price makers, they're the only game in town.

Why don't they just charge whatever the hell they want? And the answer is because while people can't substitute to other firms, they can substitute to other goods. So if you're monopolizing a good where there's very good substitutes and you try to charge a very high price, they'll just switch to the substitute.

So if the elasticity of demand is very large, a large negative number, they can't make much profit. Sure, they're the only firm. But if they're the only firm making spearmint gum and people are pretty indifferent to spearmint gum and peppermint gum, they can't charge outrageously high prices because people will just switch to peppermint gum.

But if they're the only firm making insulin, then they can charge an outrageously high price because there's nowhere to go. So the money monopolists make-- in a perfectly competitive market, firms are disciplined by competition from other firms. In a monopoly market, firms are disciplined only by the elasticity of demand. That's the only disciplining effect on the firm.

So let's do a simple example. Two books you guys probably had to read in high school, *Huckleberry Finn* and *Fahrenheit 451.* If you go on Amazon right now, you can get *Huckleberry Finn* for about $4 in paperback or $0 on Kindle. *Fahrenheit 451* is about $9 and about the same on Kindle. Why?

They're both two old books. People have to basically read them at about the same. They're all high-school-required books. The marginal cost of production can't be that different. They're slightly different lengths, but not that different. Certainly, on Kindle, the marginal cost of production is comparable.

So why? Why is *Fahrenheit 451* so much more expensive than *Huckleberry Finn?* Does anyone know? Go ahead. You've got an idea. Any thoughts? OK, well, the reason is because in the US-- oh, yeah, go ahead.

AUDIENCE: Is it something about copyright law [INAUDIBLE]?

JONATHAN GRUBER: Exactly. In the US, we have something called copyright laws, which say that for 75 years after the author's death, only the author or his estate or her estate can approve production of the book, which gives them a monopoly. Mark Twain died more than 75 years ago. Ray Bradbury did not.

So as a result, there's no more monopoly. Anybody with a printing press or whatever it is to make an ebook-- I don't know how you do that-- can make a copy of *Huckleberry Finn.* So it's perfectly competitive. And so profits get driven to $0. But *Fahrenheit 451,* they still have a monopoly right, so they're making profits.

Now, why is *Fahrenheit 451*-- everyone has to read *Fahrenheit 451.* Why don't they charge $100? Because then, high schools, some of them will say, you know what, *Fahrenheit 451* is fine, but there's other ones that are OK too. So there's some elasticity demand disciplining the market for *Fahrenheit 451,* but it's not as much as disciplining the market for *Huckleberry Finn.*

Questions about that? OK. Let's go on, then, and talk about the welfare effects of monopoly. Because now we have the tools in this class to move beyond positive economics to talk about normative economics. And we have the tools to show you that monopoly causes our first example, what we call a market failure.

Let's look at figure 11-5. Back to the example we used. The demand curve is 24 minus Q. The marginal cost is 2Q. Same example we were just using. The monopolist sells 6 units at a price of 18. The competitive equilibrium would be 8 units at a price of 16.

How do we know this? Well, we just set price equal to marginal cost. Marginal cost equals 2Q. So that would be true at Q equals 8, P equals 16. We all could see it from the intersection of the demand-supply curves. So the competitive case, they sell 8 units at a price of 16. The monopoly case, they sell 6 units at a price of 18.

What is the consumer surplus in the competitive case? What letters represent the consumer surplus in the competitive case at point EC? Don't raise your hand. Tell me. Come on, guys. Help me out here. Yeah? Loud, loud.

AUDIENCE: A plus B plus C.

JONATHAN GRUBER: A plus B plus C. Same person, what's the producer surplus in the competitive case?

AUDIENCE: D plus E.

JONATHAN GRUBER: D plus E. Now, in the monopoly case, what has happened to consumer surplus? It has fallen from A plus B plus C to just A. So consumer surplus, from the competitive case, the monopoly case has gone from A plus B plus C to just A.

What about producer surplus? Producer surplus in the competitive case was just D plus E. Well, now, what is it? Well, they lose E, but they gain B. So what's happened is two things. We have transferred area B from consumers to producers.

We don't care for that in the total-social-welfare framework. We just care about the sum producer-consumer surplus. We don't care who has it. So that transfer is not important, not consequential. What's consequential is we've created a deadweight loss of C plus E.

We've created a deadweight loss of C plus E. There's an inefficiency. Why? Because there are trades that would make society better off that are not happening. Remember, if we sold that seventh unit, the benefit of that unit to consumers exceeds the cost of producing it. Therefore, society is better off if it's sold.

But it's not being sold because of this poisoning effect. So monopoly has created a deadweight loss. And this first example, the deadweight loss of monopoly, is the first example of what we call in this class a market failure. When I did welfare the last couple of lectures, the only thing that caused deadweight loss was the government screwing things up.

As long as you let the market rip, you got the perfectly efficient outcome. Well, here we're letting the market rip and not getting the perfectly efficient outcome. Now, I'll stop right there. This does not imply government can make things better. We'll talk about that next time.

All I'm saying here is a simpler statement, which is, unlike the perfectly competitive case, if you let the market rip, you get an inefficient outcome. That's a fundamental change. It's a fundamental challenge to everything we've learned so far. It doesn't come from any magic or any government intervention. It's just the way monopolists make decisions.

So because monopolists don't want to lower the price because that's not profit maximizing, we end up not selling units that would make society better off. So the private incentives-- put another way, the positive and normative outcomes don't align. With a competitive market, the positive outcome was no profits.

That was also the welfare-maximizing outcome. In the monopoly case, the private outcome is different than the social-welfare-maximizing outcome, and that leads to a market failure. Questions about that? OK. Now, this raises the natural question, which was raised here and which has been on your mind the whole time, which is, why do monopolists have to charge just one price?

What law says that? No firm charges just one price. Firms charge different prices for different sorts of things, and we'll come to some examples. So what if we moved to the more realistic case where monopolists can price discriminate? What if monopolists could charge different prices for different quantities?

Well, in fact, what would happen-- let's go to the extreme of perfect price discrimination. Imagine a monopolist could go to every individual, figure out what their willingness to pay was, and charge them that amount. That would be a perfectly price discriminating monopolist. They would charge every person their own individual price, and that price will correspond to their willingness to pay.

What would that lead to? That would lead to figure 11-6. The first person, Person One, they are willing to pay that entire trapezoid that would be to the left of unit 1. They're willing to pay a lot, so you charge them a lot. All the way down to the eighth person.

You charge them what they're willing to pay, which is 16. You charge the eighth person 16, the seventh person 17, the sixth person 18, and so on. Why would you want to do that? Well, now, what has happened to consumer surplus? Someone raise their hand and tell me. Yeah?

AUDIENCE: 0.

JONATHAN GRUBER: 0. Why is consumer surplus 0?

AUDIENCE: They have to pay what they want to pay.

JONATHAN GRUBER: Exactly. Consumer surplus is the benefit you get from valuing something above its price. That's gone. They're charging you what you're willing to pay. So by definition, there's no consumer surplus. So all the rest is producer surplus. So the firm has converted the entire surplus to producer surplus by perfect price discrimination.

What else is notable? In this case with perfect price discrimination, firms will sell 8 units. So what else is notable? No deadweight loss. A perfectly price-discriminating firm creates no deadweight loss. Why? Because there's no poisoning effect.

The poisoning effect comes from the fact that if I want to sell more units, I got to lower my price. Well, here, I don't have to lower my price to sell more units. Given that, I might as well sell as many units as people are willing to pay for that are above my cost. So I might as well do the welfare-maximizing thing, which is sell 8 units.

This is the first inclination that we have to care about fairness in economics. Because what this says, what the existing course would teach you so far is I'm completely indifferent between the perfectly competitive case, as illustrated by EC above, where consumers get A plus B plus C-- I'm indifferent between that case and the perfect-price-discrimination case. That is, I'm indifferent between whether consumers [? get ?] [INAUDIBLE].

I'm completely indifferent. All I care about is the total. It's the sum. That can't make sense. Can't make sense that I'm completely indifferent to a world where producers have everything and where consumers have some, producers have some. There's no way that anyone would be truly indifferent to those worlds.

And that's the first hint that, in economics, we have to care about fairness. We have to not just care about the total sum of social surplus, but who gets what. And we're going to come back to that later. For now, we'll stick with our definition. We just care about the total. But this is the hint of why we need to come back to that.

So that is what would happen with perfect price discrimination. Now, in fact, firms can't perfectly price discriminate. In fact, firms can't go and get your willingness to pay. So in reality, what do firms do? Well, they price discriminate by trying to guess who is more price sensitive.

Firms engage in price discrimination, and the way they do it is they look for correlates of epsilon. Basically, if you're a monopolist, you want to charge a higher price to a low-epsilon guy and-- remember, guys, gender-neutral term-- and a lower price to a high-epsilon guy. So basically, if I knew everyone's epsilon, I'd be done.

But I don't. So what I do is I look for signals that are correlated with your epsilon. So for example, airlines charge more if you buy at the last minute than if you buy four months in advance. Why? It's the same seat, same marginal cost of flying on the plane. Yeah?

AUDIENCE: There's less substitutes if you buy at the last minute.

JONATHAN GRUBER: Yeah. If you buy at the last minute, you clearly got to go. OK, you're a business person who's got a last-minute meeting. You've got a loved one who's sick you got to go see, et cetera. If you buy four months in advance, you're clearly shopping. You have lots of options.

But if you're buying at the last minute, you don't have a lot of options. Now, I can't charge you infinity because there's other airlines you could fly. But clearly, you are less price sensitive if you're buying at the last minute than if you're buying four months in advance. That's one example.

Why do movie theaters charge less for matinees than movies at night? Why do movie theaters charge less when you go during the day? It's the same marginal cost, again, same popcorn cost, same projector cost. Why do movie theaters charge less when you go during the day? Yeah?

AUDIENCE: Lower demand.

JONATHAN GRUBER: Well, lower demand by itself is not necessarily the answer. It's really about elasticity of demand. Why are people more price sensitive during the day than at night? Think about who goes to the movies during the day.

AUDIENCE: Is it you have other options?

JONATHAN GRUBER: Yeah. I mean, people like retirees and things, which have plenty of time in their hands. If you go to a movie at night, it's like, you're done with work, you got to go at night. Maybe you can decide between a 7 o'clock showing and an 11 o'clock showing. But the bottom line is, there's one time you can go.

If you're going during the day, you've shown you have lots of time to go see movies. You can go at night, you can go during the day. So anybody who goes to a movie during the day, unless they work a weird shift, has basically shown that they have more flexibility on when they go to the movie.

As a result, you charge them a lower price. This is everywhere you look. All of a sudden, I've opened up to you a world of understanding why prices are differentiated. Remember, we discussed this a long time ago. We talked about why the incremental large drink at McDonald's, the incremental cost of a large drink is very small compared to the cost of a small drink because there's different elasticity of demand for the small drink and the incremental to the large drink.

So everywhere you look, when you go out in the real world now and look at prices, you'll see this-- firms pricing to try to target the lower prices to the most elastic consumers. Now, perhaps, the best example of this was an example from a few years ago with Tesla. Tesla sells electric cars, and the original Model S could go for 250 miles on a battery charge.

What they said is we are going to offer a cheaper Model S that can only go about half as far, about 125 or 150 miles per battery charge. So you can buy the more expensive one because it's got this fancier battery-- 250 miles you can go, it costs more money-- or the less expensive one with a less fancy battery and only go 150 miles.

Well, in September of 2017, there was a huge hurricane in Florida, and people had to evacuate. And Tesla understood that the limits on how far you could drive-- Florida is a big state-- could get in the way of people evacuating. So Tesla said, you know what? All cars can now go 250 miles on their battery.

People were like, what the [MUTED]? What do you mean we can go 250 miles on the battery? Well, it turned out it's the same battery. They just had a piece of code that limited how far you could go in the cheaper car. Why would they do that? Why would they literally have the same battery and literally add money to add code to limit how far you could go? Why did they do that? Yeah?

AUDIENCE: [INAUDIBLE] prices [INAUDIBLE], it's better to get a little bit more money from what the [? consumer ?] [INAUDIBLE] [? pay ?] less than [INAUDIBLE] at all.

JONATHAN GRUBER: Exactly. There's two types of people-- people who have a lot of money and are willing to pay for a better car and people who don't have a lot of money and are not willing to pay for a better car. If they just simply said, we're going to sell the Model S with 250 miles, a bunch of lower-income people wouldn't have bought it.

If they said, we're going to let everyone go 250 miles on a cheaper price, then they wouldn't have made as much money. So they set a high price for a premium product to grab all the low-elasticity consumers and the lower price for a less premium product to grab the high-elasticity consumers, even though the products were the same. Actually, Tesla is not alone in this.

When IBM first created the digital printer, it had a home version which was much less efficient than the office version. And it turned out-- someone took them apart, and it was the same printer. They just added crap to the motor in the home version to make it go slower because they wanted to take that same advantage. They wanted to be able to charge a lot for the office one because people are less price elastic in the office.

Turns out you can take this too far. Amazon tried to go further with this. Amazon actually tried to set prices by IP address. So Amazon, literally an exercise where, literally, they would differentiate the price you'd see online by your IP address. Why would they do that?

Once again, regardless of your feelings about Jeff Bezos, why would they do that? Why would they differentiate the price? Yeah?

AUDIENCE: [? If ?] they're willing to spend more money, [INAUDIBLE].

JONATHAN GRUBER: Yeah, but how would they know by your IP address?

AUDIENCE: Past sales.

JONATHAN GRUBER: What's that?

AUDIENCE: Past sales.

JONATHAN GRUBER: Past sales is one reason. What else? Yeah?

AUDIENCE: Zip code.

JONATHAN GRUBER: Your zip code. That's another reason. What else? Who's willing to pay a lot? Yeah?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: What's that?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Yeah. Basically, if you're ordering from Goldman Sachs' IP address, they're like, you're rich. We're going to charge you a lot. If you're ordering from some homeless shelter or some place where people don't have any money, they'll charge you less. So basically, they said, IP address is a great signal of willingness to pay.

We'll suck up all that consumer surplus by varying the price by willingness to pay. They got busted and couldn't do it anymore. They were told they couldn't do that. Now, it's not clear why not. In fact, the reason why not-- you might say, well, look, why shouldn't they do that?

They're increasing efficiency. If they don't do that, we're losing C plus E. By doing that, we're gaining efficiency. They're selling more units they wouldn't have sold to those low-income consumers. Tesla wouldn't have sold those cheaper model Xs or whatever they were. So what's wrong with that?

And the answer is, in the traditional economics world, nothing. Traditional economics world, we should have let Tesla do that. We should have let IBM make the crappy home printer. We should have let Amazon differentiate by IP address. The reason we didn't speaks to the fact that people don't think consumer and producer surplus are really valued equally, that we care about both separately. And that's what we'll start talking about more in a few lectures.

OK? Questions about that? All right. We'll come back on Wednesday, do a second monopoly lecture. You're responsible for the lectures this week on the midterm. Monday's lecture will go on to a new topic, which is not in the midterm. But it's on the final, so you don't want to miss it. Plus, it involves oligopoly, which is fun.