

[SQUEAKING]

[RUSTLING]

[CLICKING]

**JONATHAN
GRUBER:**

Why don't we get started. We're going to start by finishing up our lecture on costs, with a concept I didn't get to cover last time. Then we'll move on to talking about competition. So I want to start by talking about one concept on costs we didn't get the cover last time, which is fixed versus sunk costs.

And sunk costs are an important term in economics, so I want to make sure people understand what they are. Sunk costs are essentially costs that cannot be changed, no matter what action you take from this date forward. So in some sense, some costs are long-run fixed costs.

I know we said in the long run, nothing's fixed. But a sunk cost is the idea of an investment that once made, can never, ever be changed. So the thing is, if we think about it in the short run, let's imagine you're a doctor.

In the short run, your variable costs are how many hours you work, or the nurses you employ, or the physician assistants you employ. In the long run, your fixed costs or how big your office is. You can get a bigger office or smaller office.

But your sunk cost is having gone to med school. You can ever undo having gone to med school. Having paid that cost, it is sunk. It is gone.

You're never going to-- you could go to more med school. You could take supplementary classes, but you can never undo having gone to med school. So that's the sense in which some things are sunk costs. They're investments that are made that can never, essentially, be undone.

And essentially, we can think of these as long-run fixed costs, which I know is really confusing, because the long run is when costs aren't fixed. But that's essentially what they amount to.

Now, sunk costs turn out to have a very important place in economic lore, because

the thing about sunk costs is it's hard to think about them. It's hard to remember the rule that sunk costs are always sunk. So let me give you one example-- it's in one of my videos for the 1401X videos that supplement this course. But if you've seen it, I apologize.

We'll go through it again, which is literally a case where I got this wrong two years ago. So about two and a half years ago, I made the decision that I take my wife to go see the band Journey. So I went and I bought tickets, and the tickets were \$240 for the pair-- pretty good tickets at a big stadium, \$240 for the pair.

Then about a month before the show, I was looking, and realized I didn't actually like Journey that much. I like-- I love a couple of their songs. I looked at the set list. I'm like, I don't really like that many of their songs.

There's a couple I love, but I don't want to sit there. It's a couple hours to get there. I don't want to sit there all night for two songs that I like. I want to sell them.

So you can sell tickets through StubHub or other secondary mechanisms, but you have to decide what price you're going to ask. And I said, I gotta at least make back my \$240. That's my goal.

And then I realized I was thinking about it completely wrong. I didn't realize immediately, but I was [INAUDIBLE] about it completely wrong-- that the \$240 was already gone. The fact I spent \$240, given that I'd already spent it, is irrelevant.

How should I think about it? I should only ask how much am I still willing to pay to see Journey? And as long as someone else is willing to pay more than that, I should sell them my tickets. And if they're not, I should go.

So let's say, for example, I decided, I'd still pay \$100 for my wife and I to go see Journey. Then I should say, as long as the tickets sell for more than \$100, I'll sell them, but if it's less than \$100, I won't. And the fact I paid \$240 is irrelevant-- that's gone. That is a sunk cost.

So a common [INAUDIBLE] is the sunk cost fallacy-- the notion we pay attention to things that don't matter anymore. Having spent that money was irrelevant. All that mattered was the decision looking forward, which was, do I want to go or do I want to sell the tickets? And that decision simply depends on how much I was willing to

pay now.

Now that's confusing. I've explained this. If you guys get it, you're a rare breed. I've tried this story on many people, and they're, like, no, of course it matters how much you paid. It doesn't.

AUDIENCE: Does this kind of [INAUDIBLE] apply to the housing market crash, and homeowners are trying to sell their houses for way more than they're worth, because they're trying to get some of their money back?

JONATHAN GRUBER: The sunk cost fallacy finds its name very much in the housing market. Many homeowners set as a benchmark what they paid for their house. They'll say, I don't want to sell less than what I paid.

That's silly. Doesn't matter what you paid. What matters is how much you can sell it for and whether or not you want to sell at that price. And what you paid is irrelevant.

Now, it might matter because there are borrowing constraints we'll get into. You might have a mortgage you have to pay back, and you can't afford to pay it back unless you sell for a certain price. But aside from that, let's imagine you just bought your house outright without a mortgage for \$1 million.

Then that's irrelevant. Whether you sell your house for \$900,000 just depends on whether it's worth \$900,000 for you to stay there or not. The million you paid is irrelevant, except if you have a mortgage and you've borrowed it. That can make it complicated.

So basically, sunk costs are sunk. Now, I ended up selling the tickets for \$220, so I did OK. But if I thought about it incorrectly, I would've said that's a bummer. I lost \$20.

And in the big picture, it is a bummer. I never should have bought them in the first place. But having made that decision to buy them, the fact I was only willing to pay about \$100 at that point to go and sold for \$220 meant I did well, not badly. Questions about that?

It's a confusing example. None of your friends understand this. You could totally

mess with your friends' minds by explaining this to them. Yeah.

AUDIENCE: [INAUDIBLE] more I might make [INAUDIBLE]

JONATHAN GRUBER: What I mean is, think of them-- I have a set of Journey tickets. I have a choice-- I can sell them or go. So the only question is, how much is it worth to me to go to Journey? I decided it was worth \$100 to go to Journey.

So I simply put them on StubHub, and said if they sell for more than \$100, I won't go. If they sell for less than \$100, I'll go. The fact that I paid \$240 is irrelevant. Yeah.

AUDIENCE: So, basically you've lost [INAUDIBLE]. You've lost overall you've--

[INTERPOSING VOICES]

JONATHAN GRUBER: I lost \$20.

AUDIENCE: --by making the mistake of buying them.

JONATHAN GRUBER: Yeah, ex ante, there was a mistake made. But having made the mistake, that's irrelevant. So that's an important concept to keep in mind, is sunk cost.

Let's go on now and move on to the next topic, which is the focus of the next two lectures, which is perfect competition. Perfect competition. Now, stepping back, as we highlighted, a lot of what we did for producer theory is the same as what we did for consumer theory. We have isoquants and isocosts instead of indifference curves and budget constraints, but the same damn exercise, same math, same graphics.

The difference was we have a series of tangencies between the isocosts and isoquants, and we didn't know which one was the right one to choose. With consumer theory, we pin that down by the fact our parents gave us a certain amount of money. For producer theory, there was nothing pinning down q . There's nothing pinning down our total costs we can afford to spend.

What pins that down is an additional constraint we bring into the system, which is the market. So now, we're going to take that next step of actually deciding what a firm produces. What a firm produces partly comes from all the math we do before developing the cost curve.

But it also comes from the fact that firm exists in a market. And we focused on the cost side. Now, we turn to the revenue side. The markets can determine what the firm can make from selling that good, and that's and that's going to pin down how much they produce.

Now, we're getting set of three different market settings. Today, we'll talk about perfect competition, which is the market with many, many firms competing to sell a homogeneous good. I'll talk about it more precisely later.

Then, we're going to talk-- that's one extreme. That's the extreme of markets that economists sort of dream about is a perfectly competitive market. The other extreme is monopoly, where there's one firm, only one firm that sells the good.

And then in between, we have my favorite word in economics-- oligopoly, which is when there's several firms competing to sell a good, not as many as perfect competition, but more than one. That turns out to be super complicated. So what we do is we start with the two extremes, develop a set of intuitions and rules, and then we sort of hand-wave a bit around oligopoly. And that's where I introduce the notion of game theory, which we'll talk about some in a couple of lectures.

So today, we're going to start with one extreme-- perfect competition. What is perfect competition? Basically, the technical definition of a perfectly competitive market is where producers are price takers. A perfectly competitive market is one where a producer doesn't have any influence over the price that they sell their good at. They are price takers.

So I can't actually, as one-producer market, can't actually affect the market price. The market price is given to me. When will this be true? This will be true when the demand for a firm's output is perfectly elastic, when the demand for a given firm's output-- little q , not the market elasticity, but the firm elasticity-- when you have an infinitely elastic little q .

So let's think about that case. Let's turn to the first figure. So we have little q^2 and-- so we basically have little q , because it's the firm's output on the x-axis. The y-axis is the market price. The firm faces a perfectly elastic demand curve.

That means that they can't change the market price that's charged. So what does that mean? That means that no matter what their cost function is, whether their supply curve 1 or supply curve 2, they always sell at price p .

So shifts in the cost function-- that is shifts in the supply curve-- only affect how much you sell, not the price you get. That would be true in a perfectly competitive market. So what conditions make a perfectly competitive market?

There's basically three conditions. The first is identical products. A market will only be perfectly competitive if all the firms are selling at least what consumers perceive to be identical products. They don't have to be technologically identical. But from a consumer's perspective, they have to be viewed as identical.

The second condition is there's full information about prices. That is, consumers know what every firm is charging. So I go to market, and I know what every firm is charging. I have full information.

And the third condition is that there's low transaction costs, or what you might call "search costs," that it's very costless for me to search across opportunities. It's very costly for me to care-- everybody's price is perfectly posted, and it's very costless for me to search across them. Sort of two and three are kind of related.

So this is obviously never true, much like many of our assumptions are never true. But it's a useful benchmark for thinking about a lot of what we're going to think about. So you want to think about markets like this-- so for example, the classic case you might think of is eBay.

If you go on eBay and you search-- I just bought a pair of 72-inch red shoelaces. There's 72-inch red shoelace or [INAUDIBLE] It turns out not quite-- some are flat, some are oval. But within oval 72-inch red shoelaces, there's literally no variation. They're 72-inch, red shoelaces.

Now, they could be of different quality, and that might be unobservable. That's why it's not perfect, but it's pretty close to identical. I go on eBay, and all the prices are listed there. So I have full information on my prices, and they're easy to compare.

Now, it's not perfect because, a, there could be unobserved quality differences. Some shoelaces may be made by cheaper manufacturers. Others are easier to

break or fray.

And the second reason is that I might not have full information about prices, because at least in the old days-- eBay's changed this-- they can price compare on the non-shipping cost price. Now eBay's fixed this. You get the price including shipping costs. So in the old days, you could shop, and think you had the cheapest deal, but it turned out when you add shipping costs, it wasn't. So even eBay, which is sort of the economist's dream platform, doesn't quite meet these conditions, but it's about as close as you can come.

Now, the other example I like to point to is like buying little knickknacks in a tourist area. That basically, if you've ever gone to a tourist area and tried to buy a little knickknack like a replica of the Eiffel Tower around the Eiffel Tower, there's a bunch of guys with blankets, out selling them. And it's pretty easy to get see they're all the same replicas. And it's pretty easy to ask the guy, what do you want for your Eiffel tower replica? So that's another market.

Now, I taught this as a casual example. Someone watched this video from when I did this four years ago, and went to the Eiffel Tower-- some guy in France-- and actually did the exercise of marching from blanket to blanket. And he found that when you got further from the Eiffel Tower, they're more expensive, because there weren't many people selling them.

But once you got close to the Eiffel Tower, all these guys selling them, he found that everyone charged the identical price. And he sent me a little package of 24 little Eiffel towers. It was really cute. So that was sort of this exercise in reality.

Now, the important thing to remember, though, is this is a perfectly elastic, firm demand. So I want to talk for a second about demand for a firm's good versus demand for a market's good, so firm versus market demand. These are two different things.

And the way to think about this, think about the concept of residual demand. So you could think about the demand for a market, demand function for a market, as being some Q of p . That's a demand function of the market. That's this demand curve we've been looking at since lecture one. It's a downward-sloping function of the

price.

And we can then think of any given firm, little q 's demand, q of p , as equal to big Q of p minus S_0 of p , where S_0 of p we call "residual demand." S_0 of p is residual demand. It's what everyone else is supplying.

So the demand to my firm-- once again, under these conditions, perfectly competitive market-- the demand to my firm is simply the total market demand minus what everyone else is selling. So if I'm going to set up my little tchotchke blanket by the Eiffel Tower, my demand is going to be the total demand for little Eiffel towers minus what everyone else is selling at on their blanket. The key point is, even with fairly inelastic, big Q , you can get really elastic little q .

So let's do an example. Let's first differentiate this. If you do dq/dp , what do you get? You get dQ/dp minus dS_0/dp .

Now, the first term is negative, because demand is downward sloping. This term is positive, dS_0/dp . The higher the price, the more suppliers are going to be in the market. So it's a positive minus a negative.

So already, we see dq/dp is going to be bigger in absolute value than dQ/dp , because dS_0/dp is positive, and you're subtracting it. But we can actually go further if we assume firms are identical. Let's assume firms are identical, so that little q equals Q/N . Assume there's N identical firms.

And what that means is that S_0 equals $N - 1$ times little q . The supply of everyone else is the number of firms in the market-- that's an N , big N -- number of firms in the market minus 1, because that's your firm, minus little q , because you're all identical. Then we can rewrite this-- you can do the math at home-- we can rewrite this equation as-- under these conditions, we can rewrite this as the elasticity of demand facing firm i is equal to the number of firms times the market elasticity minus $N - 1$ times the market supply elasticity.

η is the market supply elasticity. η is the market demand elasticity. η_i is the firm demand elasticity. So the firm demand elasticity is equal to N -- I keep mixing my little n 's and my big N 's. It's all the same thing.

N times e -- so firm demand elasticity is equal to number of firms times market

demand elasticity minus number of firms minus 1 times the market supply elasticity. We know this is positive. Supply curves slope up. We know this is negative. So we know it's a negative number.

What's good about this formula is it gives us examples showing how much bigger firm elasticity can be than the market elasticity. So for example, suppose, as a simple example, suppose N equals 100 suppose it's 100 people selling Eiffel towers on their blankets around the Eiffel Tower. And suppose that the market elasticity is minus 1-- that is, it's an elastically demanded good. It's not super elastically demanded.

Basically, we call that an elastically demanded good, minus 1. So it's sort of a 45-degree line demand curve, a 45-degree line demand curve sloping down. And let's also imagine that the elasticity of supply is 1.

That's a 45-degree line, upward-sloping supply curve. These are pretty standard assumptions. It's what I drew in lecture 1, a 45-degree demand curve, a 45-degree supply curve, with 100 firms.

That gives you that the firm specific elasticity is minus 199. That is virtually flat. That is approximately negative infinity, as far as these things go.

So basically, even with just a regular-looking, downward-sloping demand curve, if it's 100 firms of independent marketing, you get a virtually flat firm specific demand curve. So it's not crazy that firms in a really competitive market would face essentially perfectly elastic demand. And that's the key thing that's going on in a perfectly competitive market, is firms themselves-- not the market demand, the market demand can be sensible-- but firms themselves face virtually perfectly elastic demand.

If you're all selling little Eiffel towers next to each other, then if you try to raise your price by one euro, you're gone. No one buys from you. If you lower it by one euro, everyone buys from you. That sort of makes sense. Questions about that? Yeah.

AUDIENCE: What was the difference between big Q and little q ?

JONATHAN Little q is the firm and big Q is the market-- very important to remember. Always

GRUBER:

remember. That's a middle of the night thing.

When I wake you up in the middle of the night, you have to know that one. Little q is the firm. Big Q is the market. Now, I'm not guaranteeing I always get that right on the board. But I am guaranteeing if I get it wrong, one of you guys will correct me.

So armed with this-- any other question about this? Armed with this, we now turn to how firms maximize profits, short-run profit maximization. This is what we've been heading for. With consumers, our goal was to model how they maximize their utility. With firms, our goal is to model how they maximize profit.

Now, what is the key thing in the short run? The fact is that short run means we're going to make-- we talk about short run being capital is fixed, but there's one other assumption we're going to make about the short run. We're going to assume no firm entry or exit.

It's sort of a complement of capital being fixed. Firms are you're in. You make your capital investment and you're in.

You could stop producing, but you still already made your capital investment. Or in other words, in the short run, capital is a sunk cost in the short run. You're in. And once the market starts, no one new is coming in.

So you sort of roll in this market. People have announced to begin. They've set up their blankets. No one else has come to set up a blanket and no one's rolling up their blanket and going home.

So now, let's ask first question-- what is profit? Well, that seems sort of easy. Profit, as I wrote down earlier, is just revenue minus costs. And if you're taking one of those goddamn boring counting courses in course 15, that's where you'd stop.

But you're not. You're take an interesting course 14 course instead, where we tell you that while this may be the correct definition of accounting profits, accountants simply say, you add up the revenues you make minus the costs you incur, and that's the profit. But economists say, wait a second-- that's not right, because we also have to account for opportunity costs, as well as cash costs.

Let's do a simple example. Let's say you're going to start a website design firm

when you graduate. And all the firm's going to be is you, some slave programmer you're going to hire for \$40,000, and the computer they're going to use.

And moreover, let's say you already have a computer sitting around. That's your computer. It's in pretty good shape. So you could just have the slave work on that computer.

And you can just basically say, look, I'll pay this guy \$40,000. I'll think up all the ideas and supervise them. And then that's the way the firm's going to work.

Let's say you do that, and at the end of the first year, you have sales of \$60,000. Then your accounting profits are the \$60,000 that you earned minus the \$40,000 you paid your employee. So you've made a \$20,000 accounting profit.

Why does that make no sense. In fact, why are economic profits, the true profits from this enterprise once you account for opportunity cost negative. What do the accountants miss that economists get? Yeah.

AUDIENCE: [INAUDIBLE] something better.

JONATHAN GRUBER: The opportunity cost of your time. You just spent a year doing this. You could have been off making a zillion dollars, like you guys all will. So we missed the opportunity cost of your time.

Let's say you could have graduated and gotten a job at \$80,000 a year. Well then, actually, the cost of the year is not just the \$40,000 you paid the programmer. It's the \$80,000 that you forgone by working that year on your company.

What else? What's the other opportunity cost? Much smaller, but still relevant. What could you've done with the computer?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Well, you could have done what?

AUDIENCE: You could have sold it.

JONATHAN You could have sold the computer, right. Now, used computers aren't worth that

GRUBER: much, but let's say you could have gotten a grand for it. Well, that's an opportunity cost.

If you would just gone on and used your-- now, you might get all the utility from the computer, whatever. Let's put that aside. But you just gave it to this guy. You never touched it, literally got rid of it. It was like you didn't have the computer that year.

Then you could have sold it. Let's say you could have sold it for another \$1,000. So now, you've paid him \$40,000. You've given up \$80,000 of earnings and \$1,000 on the computer, so your true costs are actually \$121,000. So actually, you haven't made \$20,000. You've lost \$61,000.

And that's why we don't care about accounting profits. We care about economic profits, the difference being economic profits account for opportunity costs-- very important to remember. Yeah.

AUDIENCE: The profits, economic profits will be higher than financial profit or typically [INAUDIBLE]?

JONATHAN It all depends on-- it depends on there's no way to assign it. yeah.

GRUBER:

AUDIENCE: You've been talking about sunk costs and opportunity costs. Is it like [INAUDIBLE] economics to consider time to be a sunk cost, like if you spend a lot of time on something?

JONATHAN Well, no, it's not. I mean, in some sense, a sunk-- basically, sunk costs is sort of an irreversible fixed cost. If you spend time on it, yeah, in that sense, from today's perspective, it's sunk. But the point is, the more important concept is that time is money, that if you spent your time running this business, that's time you could've spent doing something else.

GRUBER:

Now, one thing we'll come to in a couple lectures, you might say, well, hey, I would've taken the year off and screwed around. So in fact, there was no \$80,000 cost. What if instead of running this company, the only thing I want to do in my life is run the company. My second choice is watching TV.

So then there's no opportunity cost. That's wrong. And we'll teach you why in a few

lectures. I don't want you to tell me why now. I want you to think about why, even if you would have spent the year watching TV. No, I don't. I said, no. Put your hand down.

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: I knew you're going to try. I love, I love that willingness in this class to answer questions. This one I want you to just think about, because we're going to come back to this in a few lectures. Now, let's go on and talk about maximizing profits.

Maximizing profits-- now that we defined what profits are, how do you maximize them? We know how to maximize a function. If π equals R minus C , then the maximum $d\pi/dq$ -- you can't control the price. All you can control is how much you produce. Your only control variable is q , little q .

So $d\pi/dq$ equals dR/dq minus dC/dq . Know what dC/dq is. We defined that last time. What's that? What do we call that, the change in costs with respect to an increment in quantity?

AUDIENCE: Marginal cost.

JONATHAN GRUBER: Marginal cost. So we know this is the marginal cost, but for a competitive firm, what is marginal revenue? So that's marginal revenue, the amount you work on the next unit you sell. For a competitive firm, what is their marginal revenue? What is the amount?

AUDIENCE: Price.

JONATHAN GRUBER: Price, which is given to them, because they're a price taker. Firms, they don't have to think about a complicated concept here. This is all complicated, with tons of math. I hated that last lecture, all the math. This is hard.

This is easy. It's price. You're given-- if you're a price taker, you're given a price by the market. So marginal revenue for a perfectly competitive firm is just price.

So you maximize profits when price equals marginal cost is the profit maximizing point. Profits are maximized when price equals marginal cost. That is, you want to produce until what you get from the next unit equals what you spend to make the

next unit. Yeah.

AUDIENCE: What if you make-- what if you move it to cost a bit less and you sell it for below the market.

JONATHAN GRUBER: Let's go to an example. So let's go to figure 7.2. Here we have the cost function that we derived last time. C of q equals 10 plus $5q$ squared. You remember that from last time.

And then we have a revenue function, and I assume the price per unit-- we have a revenue function where I'm going to assume the price per unit is 30. For our example, I'm going to assume P equals 30. I've just made that up. Once again, that comes from God in these perfectly competitive firms.

They have no idea where this comes from. It's just a price. We'll talk later where it comes from, but for now, it's just a given thing. Let's say it's 30.

So the firm's cost function is of the form-- is graphed on the left-hand side here. The revenue function is simply a straight line, which is 30 times q . So if they sell one, revenue is 30, two, revenue is 60, et cetera. And the profits are simply the difference.

So when we maximize, how do we maximize profits? Well we want to graph-- what we want to do is graph the profit for each additional unit sold. So for example, when you go from selling no units to selling one unit, at no units, your cost is what? If you sell no units, what's your cost? Somebody raise their hand and tell me. Yeah.

AUDIENCE: 10.

JONATHAN GRUBER: 10, not 0, because you have the fixed costs. Those are paid no matter what.

GRUBER: Remember, those are fixed in the short run. We'll come back. This is very important.

So if you zero, your costs are 10, your revenues are what? 0, so your profits are negative 10. If you produce one unit, what's your cost? 15. What's your revenues? 30. So you make a profit of 15.

If you go from one unit, now you're producing one. You want to know, should I produce the second unit? What's the marginal cost of the second unit?

Well, we know what marginal cost is with this function. Marginal cost, we just differentiate this with respect to q and we get $10q$ is marginal cost. So we know the marginal cost of the second unit is what? 20.

What's the revenues from the second unit? Well, that's linear. It's 30. So you're still making-- you're making a profit of 10 on that unit.

Now, we go to the third unit. So right, now after two units, you are still making profit. Now we go to the third unit. It's hard, because we're discretizing a continuous example, but let's go to the third unit.

The third unit, what's the cost? What's the marginal cost? 30. What's the marginal revenue? 30.

So you're at the profit maximizing point. You have climbed the hill. You have maximized your profits. What's your total profits at that point?

Well, your total profits is 3 times 30, which is 90, minus 10 plus 5 times 3 squared. Minus 10 plus 45 is 55. So your total profits are 45. That's your maximum profits. That's the most you're going to earn.

Now, coming to the question that was asked a minute ago, or a related version-- wait a second, if I sell one more unit, I still have positive profits. So my profits were 45 selling three units. The next unit, I'm still in the black.

Why not sell more? Why not sell one more unit? Why not sell that fourth unit? Yeah.

AUDIENCE: Because it's going to cost you more than it's going to give you.

JONATHAN GRUBER: Right. Because we always make marginal decisions in economics. We always ask what's the next step I should take. And the next step is a losing step, because the fourth unit, what's your marginal cost? 40. What's your revenue? 30.

So the fourth unit has a negative 10 profit, so you don't want to make it. So profit maximization is a hill-climbing exercise. I like to think of it, and I describe in my videos, I think of it as it's like you're climbing a hill blindfolded. All you know is whether you're stepping up or stepping down.

And you need to figure out when you get to the peak. If you step up, if your next

step is upward, you must be short of the peak. If your next step is downward, you must have passed the peak. It's just keep going till your next step leads to a flat part. Your next steps, right in front of the other. That's what profit maximization is.

The key thing is you can-- the key thing about the blindfold is you don't need to see the big picture. All you need to know is, is my next step increasing my profits or decreasing my profits? If it's increasing, I'm doing it. If it's decreasing, I'm going backwards. That's all you need to know, is just think about putting one foot in front of the other. Does that next unit make me money or lose me money?

So that is how we think of profit maximization. Yeah.

AUDIENCE: So what's the next [INAUDIBLE]? You just stop it or--

JONATHAN GRUBER: Well, that is the optimal production. That's pinning down how much you want to produce. Remember I said at the biggest lecture, the problem with producer theory is we have some set of relationships between how much we produce and costs. It doesn't tell you how much you should produce.

This tells you how much to produce. You have now solved the firm's problem. This one extra-- this imposition of the perfect competition constraint has allowed us to finally solve the problem. We have now solved it, and we've said that the firm, given this cost function and given a price of 30, a firm should produce three units. Done.

Just like before, we said, we had this many pizzas and cookies. Now we're pinning down q . Once we pin down q , we can, of course, go back and pin down l and k , because l and k are a function of q . But this is where it's one step harder. So this extra step we have to take is we have to impose the market condition to get to the little q we want to produce. And in this case, that's 3. Questions about that? Yeah.

AUDIENCE: So you know how in the beginning, you said that in a perfectly competitive market, the producers are price [INAUDIBLE]?

JONATHAN GRUBER: Yeah.

AUDIENCE: So I'm having trouble understanding why all of the guys who are selling Eiffel towers wouldn't get together and [INAUDIBLE]. Let's just make all of the Eiffel towers more

expensive.

JONATHAN GRUBER:

Great. Great point. That's exactly-- we call that a monopoly or an oligopoly, depending how you want to think about it. They essentially could monopolize. And we're assuming that doesn't happen here.

And I'll talk later about why that's way harder than you think. So we're starting-- once again, in economics, we always start with simplifying assumptions to draw general lessons. Then we'll make some more, and then we'll come back to the more complicated, real-world examples. But it turns out virtually everything we learn here is still going to hold. Yeah.

AUDIENCE:

So if--

JONATHAN GRUBER:

We can go back to it. So now let's ask, how big is the firm's profit? How do we measure the size of the profit?

Profit, if we think about profits per unit, profits equals revenue minus costs. So profits per unit, profits per q , equals revenues per q minus costs per q . What is this term, cost divided by quantity?

What do we call that? Average costs. And how much do you get per unit? You just get the price.

So dR/dq and R/q are the same. In a competitive market, marginal price and average price are the same. It's the price. So this just says profits are equal-- profits for q are equal to price minus average cost.

So your per-unit profits is price minus average cost. And we can see that in figure 7.3. Figure 7.3 shows our cost curves for this $10 + 5q^2$ function we derived before.

And you can see we have an average total cost curve, which first declines and then increases. We have a marginal cost curve. And then we have an average cost curve, average total cost curve. And then we have average variable and average fixed.

Now, we already announced that the optimal production level was three, where price hit marginal cost. What's our profit at three? Our profit is the difference

between that price and the average cost of three units over all the units we produce.

So the profit is actually 30. So there's a discreteness problem here. I have to work this out. The problem, I said 45 before-- it's 45 mathematically. It's 30 in this because of the continuity problem of the sort of discreteness between two and three.

So the bottom line is, in a continuous example, this would be the same. And the bottom line is the profit, is the difference between the price and the average cost times the number of units sold. That's going to be the bottom line profit.

Average cost here is 10 over 3 is 3. Let's see, let's make sure we've got the math right here. Average cost is 10 plus 5q squared.

So average cost is 10 over q plus 5q. q is 3, so it's 3.33 plus 15, which is 18.33. Yeah, no, my math was right with 45. This graph is wrong.

So our average cost is 18.33. What's the revenue at that point? It's 30. 30 minus 18.33 times 3 is 45. So it is 45.

Is that right? 30 minus 18, 12 point-- no, that's wrong. No, 30, no, I got that wrong. You're right. It's 35.

So it all comes from the discrete-- it's all about the problem of a very discrete function versus a continuous function. The bottom line is this-- did I get right? Yeah, 11.67 times 3 is 35.

So the bottom line is, we got 45 from a hill-climbing exercise and 35 from this graph is basically a discreteness problem, that in some sense, in this continuous function, they give you the same thing. The bottom line is the profit is equal to-- we'll be clear on this in the problem sets and exams, what we're looking for.

The bottom line is the profit is the difference. The profit per unit is the difference between the average costs, and the price that you get for that unit. Now, with that in mind, let me ask you the following question-- what if I imposed a tax of \$1 per unit?

What if I imposed a tax of \$10 per unit? What if I imposed a tax of \$10 per unit? A \$10 per unit tax, what would that do to the cost curve?

Somebody, without looking at the handout and cheating, somebody tell me what's the new cost curve if I impose a \$10 per unit tax? What's the new cost curve? What was the old cost curve?

The old cost curve was $10 + 5q^2$. What's the new cost curve?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: That's the answer that almost everyone always gives the first time, but it's not right.

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: $10 + 5q^2 + 10q$, because I said per unit I charge \$10. If I'd said a fixed tax of \$10, you would've been right. It would've been $20 + 5q^2$, but it's not, because it's a per unit tax of \$10. Yeah.

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: We'll see that in a second. It's just an easy way to explain it. So that's the new-- with a tax of \$10 per unit, the new cost function is $10 + 5q^2 + 10q$. Oh, good, we didn't show that in 7.4.

So what does that do? Well, we can go to Figure 7.4 and see that. Here, the marginal cost and the average cost have both risen by 10.

The old marginal cost was $10q$. Now, it's $10q + 10$. So marginal cost is now $10q + 10$.

Average cost used to be $\frac{10}{q} + 5q$. Now it's $\frac{10}{q} + 5q + 10$. So the bottom line is both marginal and average cost have shifted up by 10.

What does that mean? Well, first of all, it changes optimal production. If marginal cost is $10q + 10$, and our optimization is set price equals the marginal cost, now we're setting $30 = 10q + 10$.

$30 = 10q + 10$, which says the optimal q drops to 2. Subtract 10 from 30, divide by 10, you get the optimal q . Now that falls to 2, so now, you only want to produce two units, because the marginal cost is higher but the price is the same.

You're going to produce less.

So now, you're only going to produce two units. And you could see that in the graph, that your marginal cost hits the price now at two units, rather than hitting it at three units. You see that in Figure 7.4. Your average cost is also higher, so you're going to make less profits per unit. So you're producing fewer units and making less profit per unit.

So the entire rectangle is shrunk. The rectangle has shrunk. Yes, the entire rectangle shrunk. It shrunk because you're selling fewer units.

And it shrunk because you're getting less surplus per-- you're getting less profit per unit. Remember, profit per unit is price minus average cost over quantity or price minus average cost. Since average cost is higher, you're getting less profit per unit.

You're selling fewer units at less profit per unit. So that tax has significantly lowered your profit. Questions about that?

Doesn't have to be a tax. If I'd said suddenly, whatever you're producing, there was a change in what it costs to produce, any change in what something costs to produce in the production function can ultimately affect how much you sell and you profit. So one excise you can show yourself at home, imagine that I had said there was a fixed tax of 10.

Actually, let me ask you this question-- imagine I had said there was a fixed tax of 10. How would that have affected your profit? How would that affect the amount you sell and your profit? It's a fixed tax of 10, not 10 per unit. How would that affect how many units you'd sell and what your profit is? Yeah.

AUDIENCE: Both [INAUDIBLE] would have less profit.

JONATHAN GRUBER: Exactly. You'd sell-- why would you sell the same amount?

AUDIENCE: Wouldn't affect your MC.

JONATHAN GRUBER: It wouldn't affect your marginal cost. Remember, profit maximization is where marginal cost equals price. If I have a fixed tax, that's an effective marginal cost,

but it does lower your profits, because you've still got to pay that extra tax.

So here, your profits get doubly hit. You sell less, and you make less per unit. With a flat tax, you would sell the same amount, but you'd make less per unit. That's an important distinction. I think I'll add that for next year.

Now, one other important point-- the firm has one other decision to make here that we haven't covered yet, which is that it has to decide whether to shut down. Now, remember I said there's no entry and exit. That means you can't literally leave and go somewhere else, but you can just walk away.

You can just literally say, look, I'm not in this market anymore. I'm not going to go set up shop somewhere else, but I'm not in this market anymore. So the question is, when would you do so?

So I'm going to get to it. Hold on. Is it a question or answer? If it's an answer, I don't want it yet. If it's a question, I'll take it.

So basically, suppose the price in this market fell to \$10. Should you continue to participate in the market? Well, if the price falls to \$10, what is the optimal level of production?

AUDIENCE: [INAUDIBLE]

JONATHAN GRUBER: Zero. So does that mean you should walk away? No. What are your profits? If you sell zero, what are your profits? Negative 10.

If you shut down and walk away, what are your profits? Negative 15. What's profits? Profits is revenues minus costs.

Costs at zero-- I'm sorry, you don't produce zero. That's bad. You produce one unit. I got the answer wrong. If price equals marginal cost-- yeah.

AUDIENCE: [INAUDIBLE] with the tax imposed, or are we--

JONATHAN GRUBER: No, without the tax. I'm sorry. That's what I confused. We're back to without the tax.

The price in the market drops to \$10, so you produce one unit, not zero. I shouldn't

listen to you. I should look at my notes. One unit, not zero. You produce one unit.

At that point, what are your revenues producing one unit? You get 30. Yeah, this example is messed up. I'm going to have to come back to this. We made this-- this example is wrong. I'm going to have to come back to this one.

So the key point with the shutdown is basically, you only want to shut down if you're actually going to lose more money by staying in the market than you get by exiting the market. So let me stop there. I'll come back and fix this.

I made an error on this. I'll come back to fix it at the beginning of the next lecture. And then we'll talk about finishing up short-run profit maximization. So let's pause there and come back.