[SQUEAKING]

[RUSTLING]

[CLICKING]

JON GRUBER: OK, today, we are going to continue our discussion of producer theory. Remember, to set the stage, I said producer theory is like consumer theory with one extra step. With consumer theory, I gave you a budget constraint. I gave you preferences. You could solve for what consumers we're going to do. You were done.

With producer theory, I can't do that. With producer theory, I give you a production function, which is like utility function. I give you input prices, which are like the prices consumers face for pizza and cookies. But there's no income. In some sense, the income is determined by the firms production decision.

Firms get to choose how much they want to produce. So instead of one relevant tangency between isoquant and an isocost, we end up with many relevant tangencies, which creates our supply curve. If you remember, the supply curve is the firm's marginal cost curve above the point where it might shut down.

So what we did is we said we developed a supply curve. So basically, supply curve tells us, at any price, how much a firm wants to produce. We match that with the demand curve, and we get equilibrium. We'll dive back into that next lecture.

But today I want to talk about an important feature of the competitive assumption. Remember to nail down the supply curve, we drew that supply curve only for a particular market assumption. Remember, with supply theory, what's going to drive the outcome is going to be the assumption of how we assume firms compete.

We are now working within a particular assumption on competition, which is the extreme assumption of perfect competition, where, basically, consumers are totally indifferent and fully informed. And therefore, each firm faces a perfectly elastic demand curve. The market demand is isn't elastic necessarily. But each firm faces a perfectly elastic demand curve, OK?

We're going to now extend what we did last time from the short run to the long run. In the long run, we developec a short run supply curve. In the long run, we're going to add one wrinkle, one, basically, extra assumption to perfect competition.

So to talk about perfect competition in the long run, we're going to add one extra assumption for the long run, which is we're going to assume free entry and exit. That is in addition to the assumptions made last time of full information and identical goods, we're going to add one more assumption, which is free entry and exit, that firms, in the long run, are free to come and go.

That assumption, which seems pretty harmless, firms seem pretty free to come and go as they like, will actually lead to an important conclusion. What we'll find is in the long run, when do firms enter? Firms enter when there's a profit to be made.

When do firms exit? Remember, because there's no shutdown in the long run, it's just about whether profits are greater than 0 or not. So when do firms enter? If profits are greater than zero. When do firms exit? If profits are less than 0.

What does that mean? In the long run, profits are always 0. So we're going to show you how, in the long run, perfect competition delivers the stark implication that long run profits in the market are 0. Not short run, short run, you can make money.

But in the long run, if you're ever making money, more firms are going to enter until profits are driven to 0. So to see this, I want to go to a particularly old example relevant from my lifetime, not yours, which is the market for personal computers circa 1990. And I draw this now-- this is the first of many times in this class we're going to go back and forth between side-by-side diagrams.

So figure 8-1, I'm going to go through slowly. But please, if it's not clear, ask me questions. Now, we're going to look at the market for personal computers. This is just when personal computers were starting to take off. There were only a few firms.

There was Dell, Gateway-- you probably hadn't even heard of them. They were a big personal computer firm back in the day. There were a few firms that produced PCs. But it was a pretty new market in 1990.

And what we're going to assume is we're going to assume this is a perfectly competitive market. And we'll assume that all the producers are identical, that any new PC company is identical to the ones that came before it. You can already see why these are silly assumptions. And we'll talk about that later. But that's the kind of assumption we make in perfect competition.

Now, so what you have is, in the left-hand side of figure 8-1, you have a firm's production decision. You have their marginal cost curve. So let's just say the firm is Dell, one of several firms that are producing PCs in this market.

You have their marginal cost curve, which is upward sloping. You have their average cost curve, which is a parabola, because, basically, they first pay off their fixed costs. And then they have rising marginal costs. So remember, the marginal cost curve intersects the average cost curve at the minimum of average cost. So that's the firm's production decision.

Now, I want you to carry that over. So where's the firm can produce in the short run? They'll produce where the market price, little P1, equals their marginal cost. So once again, focusing on the left-hand side of figure 8-1, the firm will choose to produce at the point, little Q1.

They'll make a profit because marginal cost is above average cost at that point. They'll make a profit of the entire lightly dotted area called pi1, OK? Questions about that?

This is the firm's production decision. They will choose price equal to marginal cost, which happens at point little Q1. That will deliver profits, which is the difference between price and average cost times the number of units produced. And that will be the lightly dotted rectangle pi1.

Now, if we take that over to the market, with only one firm in the market, they will then have some supply curve, S1. In the market, that supply curve, if the price is little P1, and the supply curve is S1, then the market produced is big Q1. So basically, big Q1 is the quantity produced in the market.

But now flipping back to the left-hand side of the diagram, what's happened? There's lots of profit. What does lots of profit mean? With free entry and exit, that means more firms will enter the market.

As more firms enter, what happens to the supply curve? We learned this last time. As more firms enter, the supply curve flattens. Because there's more identical firms, the supply curve gets flatter. It's just from S1 to S2.

So the supply curve flattens. At a flatter supply curve, now, at a flatter supply curve, you will now get that there'll be more produced at a lower price. So you'll get Q2, big Q2. But the price will fall to P2.

What happens when the price falls to P2? We'll now go back to the left-hand side. When the price is P2, Dell now only produces little Q2. Because they produce where price equals marginal cost, they produce little Q2. And that leads to a much smaller profit, pi2.

So once again, let me go through the math. They choose, originally, little Q1 at P1. That, you shift to the right, creates a supply. You've got a supply curve that yields profits of pi1. Due to those profits, more firms enter.

The supply curve flattens. As supply curves flattens, the price falls. Shifting back to the left, that delivers a lower profit. Note that little Q2 is to the left of little Q1. But big Q2 is to the right of big Q1. Why is that? Yeah?

AUDIENCE: More firms.

JON GRUBER: More firms. Remember, because little q is what Dell is producing. Big Q is what the market's producing. Well, but there's still profits. So what's going to happen? More firms are going to enter.

And firms are going to enter until profits are 0. Because as long as those profits are positive, firms can enter. So we've shown how entry drives profits down. But there's still profits being made, pi2.

So some other company will say, hey, this is still a good business to be in. They'll enter. And they'll enter until profits go to 0. Because if profits are greater than 0, there's an incentive to enter. So firms will enter until profits go to 0.

Now let's imagine a different market, the market for mainframe computers. You guys don't even know that term. When I graduated MIT in 1987, our graduation speaker was Kenneth Dale-- I'm sorry, Ken Olsen, who founded DEC, which was the big, main-- He was an MIT guy. He started this company called DEC.

They were a dominant, huge company because they made these huge mainframe computers. And he gave this speech in 1987 at my graduation. So the first lesson of graduation speeches, if any of you should ever give them-- I haven't given one. But if ever you should have this position is you don't spend the whole speech just talking about how wonderful you are.

Well, that's what he did. He spent that speech on how wonderful he was, how DEC was incredible, how awesome they were. Five years later, they were basically out of business. Why? Because they got wiped out by PCs. Because people went out-- maybe 10 years later, they're essentially out of business because they got wiped out by PCs.

Well, so let's look at the market for mainframe computers. And let's look at a big company that produced mainframe computers, IBM. So mainframe computers, they were just big. They were what we call servers today.

But we've gone to transition almost back to that model. But we went to a decentralized PC model. Now we're back to a model where there's PCs, just a node onto the cloud.

I think in that area, you had a terminal. It wasn't a standalone machine. It was a terminal onto a mainframe. So it's a different version of what we have today. So IBM was making those.

IBM was poo-pooing PCs just as much as DEC was. And they were originally at S1. Theirs was a supply curve, S1, with a price Q1, with a quantity big Q1. Now what happened was people stopped wanting mainframes. So they were losing money.

IBM had been invested big in mainframes. They didn't believe the PC was going to work. And they were losing money because people didn't want mainframes anymore. So at their original supply curve, S1, which involved IBM and DEC and all these companies, there was money being lost.

So if we go to IBM's case, so basically, S1 was the market supply curve. S1 intersected demand at big Q1 at a price of P1. If you take that over to IBM's curve, P1 intersected the marginal cost curve at a point where they produce q1, little q1, on the left-hand side figure 8-2. So I've now moved to the left-hand side figure 8-2.

IBM has been producing little q1. And little q1, that's their profits, they're negative because marginal cost is below average cost. They're losing money.

Now, they're still above the shutdown point. Look, they haven't reached the bottom of their marginal cost curve. If the price fell far enough, they'd shut down. But it hasn't. They're still producing. They're just losing money because, basically, remember last time, as long as their variable costs are below the price, they still want to produce.

So in this case, they're still producing. The variable cost is still below the price. They're just losing money. What happens when you lose money? Firms exit, like DEC.

Firms exit, that steepens the supply curve. Just like firms entering flattened supply curve, firms exit, that steepens the supply curve. You steepen the supply curve, that raises the price. And at this point, I've drawn it, raised the price to P2 where IBM is basically making zero profits.

So the way I've drawn it, we've gone from Q1 to Q2. That's raised the price to P2. At P2, price equals marginal cost equals average cost. And profits are 0.

This is the lesson that we derive, which is entry will happen till profits are 0. And exit will happen when profits are 0. So what does that mean? Profits are 0 in the long run.

So what that means is that the long run-- if you go to figure 8-3, the long-run supply curve is flat. Each firm will be pushed to produce to the point where price equals minimum average cost. Each firm will be pushed to produce to the point where marginal cost equals average cost.

So basically, entry and exit will drive profits to 0. We know that profits are only 0 when marginal cost equals average cost. If marginal cost doesn't equal average cost, there's profits or losses. So profits go to 0, profits go to marginal cost equals average cost. And we know that only happens at minimum average cost.

That's the point where marginal cost equals average cost. There's your little proof. Basically, profits go to 0, which means marginal cost equals average cost, which only happens at minimum average cost. So in the long-run, every firm produces at minimum average cost.

That is, every firm produces at maximum efficiency. They minimize their costs so they maximize their efficiency. So the long-run supply curve is flat at the minimum average cost. Remember, we assume firms are identical, so every firm is identical minimum average cost.

So the long-run supply curve with perfect competition with identical firms with free entry and exit, the long-run supply curve will be at the firm's minimum average cost, and profits will be 0. OK? Questions about that?

Now, you might say to me, Jon, that seems a little bit odd because profits aren't 0. Apple makes ungodly amounts of money. Tons of companies make ungodly amounts of money. So I bought your pizza cookies thing. But this is getting ridiculous.

I can't buy a model where companies don't make profits. There wouldn't be a stock market. What's going on? Well, two things are going on.

The first thing that's going on, the easiest answer is, remember, the companies can still make profits in the short-run, OK? So there's still short-run profits. But you look at many companies, General Motors has made profits for decades and decades and decades. They blipped up and down. But it's not like they only make profits in the short-run.

There are plenty of companies that have run a long time making profits. So it can't just be that they're around because they make profits in the short-run. So basically, what's going on? What's going on is that there's not free entry and exit.

OK, well, no, I'm sorry. Let me back up. That was so extreme. What's going on is there's a number of complications to this simple model. And we got to dive right into them because this is really an unrealistic prediction.

It's a good extreme. It's a good teaching example. You will come back to it in problem sets and exams to remember that in all the assumptions of perfect competition, long-run profits are 0. But in reality, they're not. So let's talk about why.

There's really three reasons why long-run profits might not be 0. The first is limited entry. I assumed the profits of the PC industry got driven to 0 by entry. But that's not true because, in particular, there exists in the world things we call barriers to entry.

There are costs, which are barriers to entry. And in particular, the ultimate barriers to entry is sunk costs. Because once someone's in the market having paid the sunk costs, if you're looking at entering-- OK, think about you being a doctor. If you make doctors an economic decision, you want to go look at how much money you'll make compared to the cost you'll sink into med school.

Well, if profits of being a doctor aren't very big, you're not going to sink $300 grand into med school. You're going to say, forget it. So as long as profits aren't too big, you won't go become a doctor. So profits can't be enormous. But they can be driven not all the way to 0.

They're driven to the point where unless you can pay off your sunk costs, you don't want to do it. So basically, entry will happen, not freely, but entry will happen only if you can make enough money eventually to pay off your sunk costs or otherwise you won't do it.

There's lots of these barriers to entry that basically exist in the world. With pharmaceuticals, its patents. Even in the little market that I described, our perfectly competitive example market outside the Eiffel Tower, maybe if you try to set up another blanket, someone will come beat you up. And you've got to decide, is it worth the punches I take to make the money?

And if the answer is profits are truly 0, you won't take the punches. You just won't set up your blanket. There's basically lots of barriers to entry in the real world. If there's barriers to entry, then firms will not enter till that point the profit goes to 0. They'll stop at a point where profits are still positive. So that's one example of why you might end up with positive profits in the long-run.

The second example is a little more complicated, but important, which is firms are not identical. And in particular, firms may differ at what their minimum average cost is. And if firms differ what their minimum average cost is, you can end up with the lowest average cost firms making money.

So basically, you can end up with the lowest average cost firms making money. So to see this, let's look at figure 8-4, which is confusing. But it gives an example. If there's time in section, we'll do a more complicated example. But this is a simple example.

This is an old example from a real empirical study that was done a long time ago of the marginal cost of producing cotton in different countries. So here we're just assuming flat marginal cost curves, OK? Just assume not upward slope. Just assume it's just a constant marginal cost, a flat supply curve, a constant marginal cost of producing cotton.

You might say, well, in that case, obviously, with free entry, profits are 0. But in fact, it turns out that cost is different in different places. It only costs $0.71 a kilogram to produce cotton in Pakistan because, I mean, I'm sure that's changed now, but labor is cheap, things like that. It's good growing conditions for cotton.

Argentina, their lowest cost they can produce cotton-- their minimum average cost is $1.08. Australia is $1.15, all the way up to the United States where it's $1.56 to produce a kilogram of cotton. OK, so just assume in this case, bear with me, think of marginal cost as just one number. It's a constant number, OK?

And that marginal cost is higher in the US than it is in Pakistan. Now, what happens in this market if the demand for cotton is 1,000,000,000 kilograms per year? Who makes the cotton and how much profit do they make? And what's the price?

Let's say demand is 1,000,000,000 kilograms per year. What happens? Who makes the cotton? What's the price? And how much profits do they make?

Anyone? You got this. You know the answer. Yeah? Pakistan makes it. At what price? They make it at $0.71. And therefore, what's their profit? 0, right?

Price equals marginal cost equals average cost. If marginal cost is flat, average cost is flat. So price equals marginal cost equals average cost. So in a world, that's the example we just did. In a world of perfect competition, profits are 0. Yeah?

AUDIENCE: Why don't they make it at a high level? Because on this graph, Argentina is the highest. So why wouldn't they make it for a [INAUDIBLE] or so?

JON GRUBER: Well, basically, that's a great question because presumably, within Pakistan, there's multiple companies competing. And if one company tried to sell for higher than $71, the other one would come in and sell for $71. It's weird to think of a country-level supply curve. But that's the logic.

But now let's say the demand curve is such that there's 5,000,000,000 kilograms demanded. If demand is five billion, what is the market price? If the demand for cotton is five billion, what's the market price? Yeah?

Right. This is a supply curve. This red line is supply curve. Demand equals supply at 156. So what are profits in the US at that price? What are profits in Pakistan? $0.85.

So in a world where, to meet demand, you need to draw on less efficient firms, the more efficient firms can make money. This is a confusing point. Let's go back and remember.

What drove the last example, I'm saying, was the assumption that all firms were identical. That led to this flat supply curve at long-range average cost. But in fact, when firms aren't identical, some are more efficient than others. The more efficient firms can still make money even in the long run, as long as demand exceeds their part of the supply curve, they can still make money.

So another reason why there could be profits in the long-run is because particularly efficient firms can make profits if demand is sufficiently high. Questions about that? Yeah?

AUDIENCE: If the demand is at the higher one, where the US has to make more--

JON GRUBER: Yeah.

AUDIENCE: --why doesn't Pakistan just make more?

JON GRUBER: Oh, well the point is Pakistan will make more. But think of Pakistan's supply curve as such-- it's a great question. Think of Pakistan's supply curve as such that after they've made 1.7 million kilograms of cotton, their cost just goes through the roof.

So that's the way we think about it. But you're right. They could make more and make even more profits. All right? So that's the second reason why you could end up with long-run profits.

The third reason is quite different. The third reason hints at something we'll start talking about in a few lectures, which is that input prices may depend on quantity. They do, input prices depend on quantity.

That is, as you produce more, the cost of buying inputs could rise. So to see this, I'm now going to skip ahead to something we'll talk about much more in a few lectures, which is, let's look at the market for labor. Remember, so far, in this class, we've taken the wage of labor as just given from God. But now we're going to talk about where that comes from.

It comes from the market for labor. So let's go to figure 8-5. Figure 8-5 is the kind of figure we're going to start working with in a few lectures. It's the classic supply demand figure, quantity on the x-axis, price on the y-axis.

The difference is the quantity is not the number of cookies. It's the hours of work or number of workers. And the price is not the price of a cookie. It's the wage. The wage is the price of labor.

So this is a demand and supply curve for workers. Why not? You can draw a demand supply for anything. So like any demand supply curve, we have a demand for workers. The demand for workers is downward sloping, OK?

Why is the demand for workers downward sloping? It's downward sloping because as the price of labor goes up, firms want less of it. We know this. If you go back to your underlying production theory, if you solve for those isoquants and isocosts, we know as the isocosts shift, that's going to change your production mix. You're going to end up wanting, as labor gets more expensive, you want fewer workers.

So therefore, if the price of labor goes up, they want fewer workers. That's why the demand is downward sloping. Why is the supply upward sloping? Can anyone tell me? Why is supply of labor upward sloping? Yeah?

AUDIENCE: That's half the wages for workers.

JON GRUBER: Exactly. The assumption is-- and I'll teach you in a few lectures, it's only an assumption. I like the celebration there. That's good. Good answer.

It's only an assumption. But the assumption we're making here is, as there's more wages, more people want to work. As wages go up, people want to work harder. More people want to work. They want to work harder. Supply curve slopes up.

And you end up in some equilibrium, L1, W1. Now, imagine that people suddenly want more of a good. Well, if they want more of a good-- and that's in some other diagram. That's not in this diagram. It's the labor market diagram.

There's some other diagram if people want more of a good. If people want more of a good, you need to hire more workers. How do we represent that here? Well, that means there's an increase in demand for workers. That means there's a shift outward in the demand curve to D of Q2.

Once again, people want more rivets or whatever we're producing. To produce more rivets, you need more workers. That's a shift outward the demand for workers. What does that do? Well, that increases the number of workers. But it also raises the wage, the price of labor.

With an upward sloping supply curve, you're going to find that if you want more workers, you're going to have to pay more to get them, right? So think of the firm's decision. I want to make more rivets. I need more workers.

Well, I'm in equilibrium. My current level of workers, they are happy. At the wage I'm offering, I'm getting the right number of workers. If I want more workers, I have to raise the wage.

Well, if I raise the wage, that feeds back to my marginal cost. And that leads to something like figure 8-6. So basically, when I was producing before, I was at MC1 and AC1. Now, I want to produce more. And I was producing at little q1 and price P1.

I want to produce more and produce little q2. Well, to produce little q2, I've got to raise the wage. If I raise the wage, that raises my marginal average cost. So I shift to the new MC2 and AC2, and my costs are higher.

So now you actually see an upward sloping supply curve, that as the wage goes up, costs go-- I'm sorry, as you produce more, your costs go up because you have to pay more for your inputs. So what happens is my marginal cost shifts from MC1 to MC2. So I have to sell at a higher price to break even. To get to that point where marginal cost equals average cost, that efficient point at E2, I've got to charge a higher price.

So now you see an upward sloping supply curve. So this is a perfectly competitive market. Notice, firms are making zero profits. In both e1 and e2, the firms making zero. Little e1 and little e2, profits are 0. How do we know? Because marginal cost is equal to minimum average cost.

Yet, even though they're making zero profits, when you want more goods, the price has to rise. And that's because they have to hire more workers. And that's the key distinction between cases one and two and case three. In case one and two, barriers to entry, our firms non-identical, leads to long-run profits greater than 0.

But in case three, long-run profits are still 0. It's an important distinction. In case three, long-run profits are still 0. Firms are still producing at the minimum average cost that every firm is. It's just the supply curve is upward sloping. So I said, this is the complication, which doesn't explain well. That's why I backed off and said there's long-run profits.

There's still not long-run profits. But there is a long-run upward sloping supply curve. So what I concluded a few minutes ago is a long-run, flat supply curve meant zero profits. Here, we're actually a long-run upward sloping supply curve and 0 profits.

Why? Because we'll have to pay more for labor as we produce more. So now, this, once again, still unrealistic because we know profits aren't 0 in the long-run. It's just showing you can't assume, as someone might ask you to in a problem, assume that because the long-run supply curve is upward sloping, there's profits to be made. That's not necessarily true.

Here's a case where there's no profits. In these cases, the long-run supply curve is upward sloping and there are profits being made. OK?

That's a lot. That's confusing. Let me pause for questions about any of that stuff. Yeah?

AUDIENCE: The input pipeline on [INAUDIBLE], maybe the supply curve on that, would that both be labor market?

JON GRUBER: Right now, the labor market supply curve is upward sloping by assumption. I'm talking about creating upward sloping in figure 8-6. So you compare your figure 8-6 to figure 8-3, OK? In figure 8-3, you see that every firm produces their minimum average cost.

I could have drawn a whole bunch of firms here. They'd all be at that same price, or a whole bunch of quantities that are all the same price. Now, what I'm saying in figure 8-6 is, no, as the quantity goes up, you get a higher price, even though profits are still 0. OK, good question. Yeah?

AUDIENCE: Is there a way to measure the effects of one and two versus three kind of like [INAUDIBLE]?

JON GRUBER: Once again, we don't have income substitution effects here because the income substitution effects come from underlying consumer preferences. This is a firm decision. So we're not getting income and substitution effects.

The way you'd care about one, two, versus three is ask, are there long-run profits? So if you want to ask, which is the right case, there's an easy test. Are profits 0 in the long-run? Profit 0 in the long-run, then either our base case is true or this-- so we have four total cases.

We have our base case. And then we have these three complications. Both the base case and this third complication delivers one prediction, which long profits is zero. These two deliver a different prediction, which long profits are greater than zero.

Now, differentiating between each of those pairs of cases, that'll involve additional testing. So for example, differentiating between the base case and this case, you differentiate empirically by asking, do wages rise as quantity rises? In these cases, you differentiate by asking, well, are firms identical or are there sunk costs, et cetera?

Good question. Other questions. Yeah?

AUDIENCE: When you look at economics, how do you tell economic opportunity cost from [INAUDIBLE]?

JON GRUBER: That's an awesome point. Apple's $18 billion per quarter profits, not bad, are accounting profits. What are economic profits? That's a great question. And that is much, much harder.

Now, to be fair, at one level, it's harder. It's not like accounting profits are easy because Apple gets to make up their numbers. So let me give you a simple example of what Apple does. Apple, by the way, makes $18 billion in reported profits to their stockholders. To the US government, they report profits of 0.

They report $18 billion to their stockholders. But to the tax authorities to whom they pay taxes, they report profits of 0, essentially. Why? Well, because they play a game where they pretend they have a company in Ireland with two employees.

But Ireland has low taxes. They claim that all their profits are earned in Ireland. How can they do that with only two employees? They claim that their profits come from intellectual property. Intellectual property is housed in Ireland with those two employees.

So accounting profits are not necessarily right either. Accounting profits are not so clean. But economic profits would be very, very hard. Economic profits or economic opportunity cost is also very hard to measure. There is measuring firm profits.

And that's why, in some sense, many of us would say, well, look, there's an easy answer to economic problems. There's tax firms that make all the profits. The problem is, what are profits? And it's very hard to measure both economically and for accounting reasons.

Great question. Other questions? Yeah?

AUDIENCE: In the figure 8-5 where you talked about how the demand curve shifts or the cost of the wages shifts based on the demand curve shift, which is in relation to the market demand. But then does the supply curve ever shift to adjust for the increased cost?

JON GRUBER: Well, no, but the supply curve is determined-- that's a good question. Supply curve is determined by the workers preference for how they want to work. Nothing's changed that. The supply curve is just how much more do I want to work when the wage goes up?

And that's just fixed. That could change. And we'll talk about that in a few lectures. In fact, we'll tell you in a few lectures why it might not even slope up. It might slope down. And that's about it.

Then our big income and substitution effects fan here. We'll have something to talk about. Question, yeah?

AUDIENCE: There's an assumption of perfect competition to also workers. Wouldn't that mean that in the long-run, all workers would make no profit?

JON GRUBER: Wonderful point. And we will come back to that as well when we talk about-- we'll do a whole lecture on how workers make their decision how to work. And we'll come to that point. OK? Great questions.

All right, well, let's follow up on the Apple discussion to talk about something that-- once again, let's interject the real world here for one second and ask, do firms actually maximize profits? We assume that. We assume firms maximize profits. There's no reason not to, or the logic is pretty impeccable.

But in fact, we see a lot of things that don't look like firms are maximizing profits. So for example, in 2010, Amgen lost 3% of their market value and 7% over the past five years. Yet their CEO made $21 million.

In 2008, the CEO of Abercrombie and Fitch, whose stock had dropped 71%, was paid $72 million. That's a lot of money based off good abs, right? And then we have the case of Leo Apotheker, who was CEO of Hewlett Packard for 11 months. And it was so disastrous, they fired him after 11 months, and he made $13 million.

So the question is, if firms really maximize profits, how can we be paying all these outrageous amounts of money to employees that don't necessarily seem to be doing such a good job? And there's two answers to that. One answer is, well, look, we know one answer is the opportunity cost. We don't know what would happen if those CEOs hadn't been in place.

Maybe they would have lost 120% of their value. Maybe they would have done even worse. Maybe these highly paid CEOs are worth it.

Look, a great example is Michael Jordan. There was an academic study which measured the value of Michael Jordan. Michael Jordan was the highest paid player in basketball at one point at $32 million a year.

Someone did a study based on the ad revenues that are generated by increased viewership when Michael Jordan was on TV. And based on that, he was worth $100 million a year. So he was actually underpaid at $32 million.

So basically, the point is, just because someone's paid a lot doesn't mean they're not worth it. OK? We can't just look at someone paid and say they're not worth it. That's one explanation.

However, that one doesn't seem to work. And the way it does the work is a super clever test that was done by a professor at Sloan Antoinette Schoar and now a new professor in the computer science department, Samuel Mullainathan. They did a very cool test.

Actually, this is Mullainathan, Marianne Bertrand in Chicago. What they said is, look, let's look at what happens to how much CEOs of oil companies get paid when the price of oil goes up. When the price of oil goes up, oil company profits go up. But it's not due to anything a CEO did.

The price of oil is determined by a world market. No one CEO can have any influence on the price of oil. So if the profits go up because the price of oil goes up, CEOs shouldn't be rewarded for that. They did nothing. And yet they find that when oil prices go up, CEOs get massive raises in oil companies.

That can't be that they're doing a particularly good job. That's got to be something else. So what's the something else? The something else is to recognize that firms are not just some black box little q function or f function. Firms are actually a complicated set of what we call agency relationships.

What do we mean by that? We mean that when you think of the firm in this course, we're thinking about something like someone working for themselves or a pair of partners where the workers are the owners. But most of the goods in America, not most firms, most of the goods in America, are produced by giant corporations.

And the key hallmark of a corporation is the separation of ownership from control. Corporations are owned by a disembodied set of stockholders. But the decisions are made by employees, managers who work at those firms.

And that separation of ownership and control leads to a fundamental challenge, for which my former colleague, Bengt Holmström, won the Nobel Prize, was one of several Nobel Prize for documenting this fundamental tension, understanding it, which is when you have the owners different from the managers, the managers don't have necessarily the right incentives to do what's profit maximizing for the owners.

So here's the example I like to use. Imagine you're a corporate manager. And this is a world you may be in someday. But imagine you're a corporate manager at a pretty big but not massive company.

It's big enough that you don't yet have a private jet. But all your buddies at the bigger companies have private jets. And they rag on you mercilessly, like, you loser. You don't have a private jet. You're nobody, dah, dah, dah, dah dah. OK, so you're like, I need a private jet.

The problem is you've done the math, and it's not economical to have a private jet. You can fly first class anywhere you need to go for less than it costs at a private jet. But you've noticed that if you change things a little bit so you flew only at the busiest times and only the most expensive airlines, and you did the math, then the cost of flying commercial would be more than the cost of a private jet.

So you make a PowerPoint. And you show the owners, look, we could save money by a private jet. Now, the owners could take the time to reverse engineer your calculations and figure out that you cooked them a little bit. But they're not going to.

They're a bunch of disembodied people owning tons of stocks. They're not going to take the time. So yeah, you make sense. Get a private jet. And now you can hold your head up high with your buddies because you got your private jet.

That is not profit maximizing. Profit maximizing was to fly commercial. But it's what made you happiest as the manager. What makes the manager best off is not necessarily what makes the profit the highest.

And in that situation, you could end up with profits not being maximized. Now as firms recognize this, and they set up something to deal with this. They have something called a board of directors.

What the board of directors are, it's a set of representative of the shareholders that keep an eye. The idea is, look, you've got millions of shareholders in a company. You've got this manager. So the administrator of the company elect a set of 20ish people who keep an eye on the manager.

They should be doing the calculations. They should be watching it to make sure the manager is serving their interest. The problem is, boards of directors just don't do a good job. Why? Because the people on the board of directors is largely chosen by the managers.

And it's just their corporate buddies. It's a lot of retired CEOs who like having private jets and think you should have one, too. OK? So board of directors doesn't end up doing necessarily what they should.

A great example was Richard Grasso was a great American success story. He went from working in the shipping department at the New York Stock Exchange to running the entire New York Stock Exchange. Great success story.

When he retired, there was a lot of controversy because he got $187 million as a retirement package. It's pretty good. And there was a lot of press about, that's outrageous.

And they went to the board of directors, and director said, well, we had no idea. We were just told it was this retirement package with a bunch of options and stock pieces and stuff like that. And it looked pretty good to us. We couldn't unpack it. We had no idea that it had $187 million.

Because it's in the managers incentive, it's his whole life's work to maximize his happiness. The board of directors, it's something they're doing part-time while they go play golf the other times. OK?

So at the end of the day, the separation of ownership and control is real and costly. And it means that firms might not be doing things to maximize profits. And that's led, once again, the whole point of this course, is to not solve all your economic problems, but to pique your interest in the field.

This has led to a huge part of the field of what's called corporate finance, which is, how do you set up contracts to align the incentives of owners and employees. This led, for example, to the revolution in stock options. Many of you will know about stock options. You're hopefully hoping to get them from your company you go to.

The idea is that when you get hired, they don't only pay you in cash. They give you some stock. And what's the idea of stock options? It's not actually stock. What they do is they give you what's called an option, which means that you have the right to buy the stock at today's price or at some strike price.

Why is that valuable? Well, it's valuable because if the stock goes up, then you can buy it at today's price and sell it at some future higher price. So stock option, the idea is it's something which is worthless to if the stock value goes down, but valuable if the stock price goes up.

So let's say you go to work today for a company where the stock value is $100. They say, we're going to allow you to buy 1,000 shares at $100. If the stock price goes up to $200, you buy those 1,000 shares for $100 and sell them for $200, you just made $100,000.

And the idea is, well, that will give managers the incentive to raise the value of the company. That will align the incentives. And that's why they became so popular.

The problem is they've led to disasters. And they've led to disasters because of the exact kind of accounting discussion I was having a few moments ago, which is that managers can do actions which raise the short-term stock price, but aren't necessarily in the long-term interest of profitability of the company. And the classic example of that was the company Enron, which went bankrupt when you were young.

Enron basically had a bunch of managers who had a bunch of stock options. What they did is essentially set up a Ponzi scheme where they had Enron sell its subsidiaries. They had Enron sell an unnamed subsidiary that was also owned by Enron, a good.

And it looked like Enron was booking big sales. But of course, Enron was also booking the purchases. But it was also some of the subsidiary. And Enron stock price went through the roof. And these guys got super rich and tried to run away from the company.

Some did, some got out, but some didn't. It got exposed, and the company went bankrupt. But I don't think any of them went to jail. Some of them got penalized a lot.

So basically, the problem with stock options is they work in theory. In principle, they can lead to even more bad behavior, which you can end up with people trying to then go further to try and manipulate the stock price to maximize their income. The point is that this is a very hard and fascinating enterprise. And there's always room for thinking creatively about how do you align the interests.

I'm going to talk about the interests of the workers. We're only talking about the interests of one rich guy versus another rich guy, OK? There's also the interests of the workers and how do we think about contracts that also bring in the interests of the workers is another interesting question.

It's a fascinating area and one which you can learn more about if you study the field of corporate finance and think more about those areas. OK, questions about that? All right, let's stop there. And go see the Arkells tomorrow night. And I'll see you guys in class on Wednesday.