[SQUEAKING]

[RUSTLING]

[CLICKING]

JONATHAN GRUBER: OK. So today, we are going to turn from the consumer side of the story to the producer side of the story. Once again, in the first lecture, we talked about supply and demand curves. The last few lectures, we developed were demand curves come from. Now we'll develop where supply curves come from. So we're going to flip and explain where supply curves come from.

And to do so, we're going to start by talking about production functions. Now, we all know what firms-- what consumers want. They want to figure out the bundle of goods that makes them as well off as possible given their budget constraint. What do firms want? Firms want to maximize their profits. So just like consumers want to maximize their utility, firms want to maximize their profits.

And firm profits will depend-- so they want to maximize profits, and profits will depend on three things. They'll depend on consumer demand, they'll depend on the costs of production, and they will depend on competition.

Now, the first two, the first and the third are things which are largely out of the firm's control. Not totally out of the firm's control. We'll come back to how things like advertising might affect demand and how things like monopolies might affect competition. But in the standard model, view them as out of the firm's control, take them as given. What the firm controls is the cost of production. That's what they can-- that's what's under their control.

So from the firm's perspective, maximizing profits is the same as minimizing costs. The way to maximize profits is to minimize costs. And that's how the firm produces as efficiently as possible, by minimizing costs. Once again, in the real world, there's complications, but in the model-- and what will explain most of the real world, we just say firms, their goal is to minimize costs.

Now to understand how to minimize costs, we have to understand where a firm's costs come from. And to understand that, we have to understand what firms actually do.

So in this course, firms are simply black boxes for much of the course. They're simply boxes where inputs go in and outputs come out. And in the middle is something called a production function. That production function is the magic that makes-- is the magic of the firm.

And technically, what a production function is, what a production function is, it's a description of what is technically feasible for the firm to produce. That's a production function, what is technically feasible for them to produce.

And essentially, what this production function does is it takes as inputs two things. It takes as inputs factors of production. It takes factors of production, what we call factors of production-- I'll come back to what those are in a minute. Take factors of production. It puts it through a production function, which is a black box, will express a mathematical function. And out comes the production of goods and services. And that's the way we model the supply side.

So what are these factors of production? These factors of production are what we call inputs into the production process. The inputs into production are of two types-- to make life easy-- once again, obviously in the real world, there's a million types of inputs into production. But to make life easy, we'll start with a model with two types of inputs into production, labor and capital. Labor and capital.

Labor, it's clear what labor is. It's the amount of hours that people work at your firm. So labor is literally how many hours are devoted towards work at your firm. Hours are provided by individuals in the labor market.

So now, people on the other side of the equation. Instead of people wanting pizza and cookies, now people are inputs into production because they produce the labor. So people produce labor, they work at your firm. And the price they demand for that labor is the wage. That is the price that they demand for their labor. Where does the wage come from? We'll talk about that in a few lectures. Right now, we're going to take the wage as given.

Just in previous lectures, we took the price of cookies and the price of pizza as given. And we'll tell you later where they come from. Now we're just going to take the wage as given, I'll tell you later where that comes from. By proceeding in steps, we make it feasible to model this in a reasonable way.

Capital is more confusing. Capital is more confusing. This is basically-- think of it as everything else that goes in production. The machines, the buildings, the land you build on. Think of it roughly as machines. Think of this black box as being workers and machines come together to produce something. The workers we call labor, the machines we call capital.

And the price of capital we're going to call the rental rate. In other words, think of whatever you use in production as rented. You rent a machine. You rent a building. In some sense, workers, the wage is the price you pay to rent an hour of a worker's time. That's what the wage is, it's the price for renting an hour of worker's time. And the rental rate is the price for renting a unit of capital.

Now later on, we'll talk about where that comes from, the fact that firms actually own machines, et cetera, but for now, this makes it easy and allows for a parallel analysis between workers and machines. They're both things you rent.

Now, the production function is a function q that basically converts-- I'm sorry, it's a function f. I'm sorry, it's a function f that converts labor and capital into outputs. So the production function f converts labor and capital into output. Just like the utility function-- and once again, there's going to be enormous parallels to the consumer theory. That's why I drilled into you developing the intuition of consumer theory. That same intuition is going to carry over here to producer theory.

So just like the utility function is a mathematical representation of how goods turn into happiness, the production function, which, in some sense, is easy to understand, is a mathematical representation of how inputs turn to outputs. So easy to understand. The utility function is a vague thing. A production function is the technology. We all know that. It's the 2007 contest. You've got some inputs and you produce an output. So a production function is just the technology that translates the inputs to the outputs.

Let me highlight one thing. Little q I'm going to distinguish here-- and if I get this wrong you can catch me. Little q refers to a specific firm. Big Q refers to the market. So I did my supply and demand curves before I had big Q? That refers to the whole market. Little q refers to a given firm within the market. And we'll go over that throughout. OK.

Now, these inputs can be what we call variable Or fixed. And what do we mean by that? A variable input is an input that can be easily changed. A fixed input is one that is hard to change, at least in the near-term. So the key distinction between whether an input is variable or fixed is whether you're talking about the short-run or the long-run.

The short-run is defined as the period of time over which only some inputs are variable and others are fixed. The long-run is defined as the period of time of which all inputs are variable. That's it, that's the definition. I'm not going to tell you the short-run is six months and three days and four hours. And the long-run is past that. I can't tell you that. I can't tell you what they represent, except theoretically.

Theoretically, we define the short-run as the period of time of which some inputs are fixed and some are variable. The long-run is the period of time over which everything is variable. Think of the short-run as days and the long-run is decades. And so where in between days and decades things lay in depends.

So what we're going to do is we are going to assume that labor is variable and capital is fixed. That is, labor is something you can adjust at any point in time. You can always fire somebody and hire someone the next day. Now in reality, that's difficult, but in theory, you can.

Capital is harder to change. You've got to build the new machine, you've got to build a new building, et cetera. That's slower to change. So we're going to call labor the variable input and capital the fixed input. So in the short-run, we get that labor will be variable and capital will be fixed, but in the long-run, they're both variable.

Now, of course, once again, in reality, this isn't so distinct. This isn't so clean. You can't just fire workers right away and replace them the next day, especially in this labor market with a lot of labor shortage. And you can build machines pretty quickly. So the short-run/long-run distinction is obviously not perfect, but it allows us to essentially distinguish the concept of things you can change quickly versus things you can't. OK, let me pause there. There's a lot of definitions. Questions about that? Yeah?

STUDENT: Shouldn't q just be L plus k

JONATHAN GRUBER: No. No, it should not. q could be L plus k. That would be a particular form of f. But just like you can't tell me what form my utility function is, I can't tell you what form your production function is. This is some rich function of how you translate L and k to q.

So if q is just L plus k, that would be a particular production function, which would imply at another unit of labor, another unit of capital deliver the same thing. But what if a machine is 5 times as good as a worker? Then it'd be L plus 5k. What if what if the effect of machines and workers varies based on how machines and workers you have? That it's a richer function.

So just like utility isn't just cookies plus slices, it's some function of cookies and slices, production is some function of labor and capital. Good question. Other questions? OK.

So now, let's talk about the short-run. Let's talk about short-run production, and how do we think about a firm's short-run production decisions? How do we think about a firm short run production decisions? Well in the short-run, capital is fixed, so we can write our production function as q equals f of L and k bar. Capital is fixed in the short-run. That's definitional.

So what the firm is trying to ask is, look, I can't change their machines I have, that's fixed. In the future I could change it, but for my decision today, I've gone to work today-- September 25, I showed up at work. I have a decision to make today, which is how many workers do I want today?

Well, the key input to that decision is going to be the marginal product of labor, which is how valuable is the next worker in terms of delivering output? Which is going to be dq dL. That is, what is the marginal effect of the next worker-- remember, this course always works on marginal decisions. What is the marginal effect of the next worker on output produced?

We are going to assume that in the relevant range, dq dL is negative. That is, we are going to assume diminishing marginal products. This does not mean diminishing total product. This does not mean-- once again, like with consumers, more is generally better, more workers are better. But each worker, given fixed capital, adds less to output.

Now, this is hard to think about and harder than-- it's easy with pizza slices to have this intuition, but same intuition as pizza slices, so let's do an example that illustrates it. Let's think about digging a hole. You want to dig a hole and you have one piece of capital, a shovel.

And that can't be changed in the short-run. All you can change is how many people you have working on digging the hole. So you hire the first person and they're digging. You hire a second person, clearly you'll get faster and more productive digging. But will it be twice as fast? Well, maybe if you can get them to alternate shifts. One works 12, one the other 12.

How about the third worker? Will it be 3 times as fast? Maybe, but probably not. By the time you're the hundredth worker, there's only one shovel, what are they going to do? They're sitting around, and you can rotate them really quickly, and you're certainly better off with 100 workers than with 90 workers, because you can rotate it more quickly. But clearly, the 100th worker is not as valuable as the second worker in terms of digging that hole.

So the idea of diminishing marginal product is given that you only have one shovel, given your fixed capital, each additional worker does less and less good. It's not quite as simple as an intuition as diminishing marginal utility, but I think it's sort of natural. And it's not always true. Like I said, you can imagine the second worker might be as good as the first worker. But in general, each additional worker will add less and less value if they're working with the same number of machines.

Now in the long-run, you can take more advantage. You could add wheelbarrows and more shovels. And you could add lots of capital and then make each worker more productive. But in the short-run, with only one piece of capital, that one shovel, there is nothing else you can do. Yeah?

STUDENT: That negative sign means d squared q over dL squared is negative, right? Not dq over dL--

JONATHAN GRUBER: I'm sorry. Yeah, I'm sorry, it's diminishing. I shouldn't say negative, it's diminishing. My bad. It's diminishing. My bad. It's not negative, it's diminishing. Sorry about that. So we've diminishing marginal product. dq squared and dL squared is negative, but dq dL is positive, each additional worker adds value, it just adds less and less and less. Thank you for catching that. Other questions?

So diminishing marginal product in the short-run. That's all we're going to say about the short-run. The more interesting case is the long-run. So let's talk about long-run production. The long-run is more interesting because now, you can vary both capital and labor. So now, the problem is going to look just like the consumer choice problem.

Consumer choice, we never said you couldn't affect the number of pizzas, you can only change cookies. That's what I've just done here. I've said you can't change pizza, you can only change cookies. And just as in that case, given a fixed amount of pizza, each additional cookie does less and less.

Now we're going to vary both. And guess what? We're going to choose a familiar form for the production function. Let's assume q equals the square root of L times k. Let's assume q equals the square root of L times k, a familiar form. But once again, a useful form in this context just as it was in consumer theory.

And let's go to figure 5-1. This illustrates a set of combinations of k and L which deliver the same amount of q. Look familiar? It should look familiar. These look like indifference curves. But in this context, we call them isoquants, because we like making up cool words. Isoquants. An isoquant is the combination of k and L that delivers the same output.

So all the features of indifference curves that we liked apply to isoquants. Further out is better. They slope down, they don't cross. And basically, what it's saying is, it's about opportunity cost. That basically, as you have more capital for a given amount of money you want to spend and less labor, that's going to affect how much you produce. If the amount you produce stays constant, you switch capital and labor, you stay in the same isoquant. If you can produce more, you move out to a higher isoquant.

Now, here's a good question. It relates to the earlier question that was asked about k plus L. What determines the shape of these isoquants? What determines their shape? What underlying notion? It's a hard question. Who here knows the intuition? When are these isoquants going to be more bendy and what are they going to be-- when are they going to be straighter and when are they going to be more bent? What's going to determine that? Yeah?

STUDENT: Depending on how many different inputs go with each other?

JONATHAN GRUBER: Exactly. It's going to depend on the substitutability of the different inputs. The more substitutable, the straighter is the isoquant; the less substitutable, the more bent is the isoquant.

So for example, let's look at figure 5-2a. Here, we have the production function that was suggested earlier. I like to think of this as a production function between this is one where goods are perfectly substitutable. I'll take this production function between Harvard graduates and Beanie Babies. Perfectly-- Nothing on that? Come on, guys. Work with me here.

This is the-- they're perfectly substitutable. And I don't mean to offend the Harvard-- the Harvard students in the class, they're smart, they came over here to take classes, so they're exempted. But basically, the idea is these inputs are perfectly substitutable. So it doesn't matter if you have a unit of labor or capital as long as you have-- just matters the total number of units you have in that term is your production. That would lead to these straight, downward-sloping isoquants.

On the other hand, consider something like figure 5-2b on the other side of the page. These would be non-substitutable inputs so that the production function is q equals the min of x and y. For example, think of these cereal and cereal boxes Doesn't matter how much-- for every amount of cereal you got, you got to have a box. If you only have one box, it doesn't matter how much you have, it's just going to sit there. You need more boxes.

So just like the indifference curves for this case were left shoe, right shoe, this is-- you can think of cereal, cereal boxes, things which are non-substitutable. They can't substitute for each other at all. And in that case, you get these squared isoquants. Questions about that intuition?

So it's the degree of substitutability which determines the shape of these isoquants. More generally-- will generally be in between. As I said, nothing's ever at the extremes. We'll generally be in between these two cases. And we'll have an isoquant that looks like figure 5-3.

And along that isoquant, we have the slope, which is the marginal rate of technical substitution, just to really mess you up. Now with consumers, your marginal rate of substitution and marginal rate of transformation., MRS and MRT. Now I'm going to have MRTS, which is the marginal rate of technical substitution.

That is-- the marginal rate of technical tactical substitution is the rate at which effectively you can trade off labor and capital to produce the same quantity of good. And essentially, we can relate this, just as we can relate the MRS to marginal utilities-- yeah, I'm sorry, go ahead.

STUDENT: [INAUDIBLE] marginal rate of what substitution?

JONATHAN GRUBER: Marginal rate of technical substitution. Marginal rate of technical substitution. Remember, when we developed the marginal rate of substitution, i used a little proof to show you how it relates to marginal utility. So we can use the same proof here, which is let's ask what happens if the marginal increase in labor and a marginal decrease in capital so that you stay on the same line?

Well, we know the-- so we're asking where dL times MPL, the marginal product of labor, plus dk times MPk, just as a marginal product of labor, there's a marginal product of capital. Same intuition's diminishing. One worker, each additional shovel does less and less good. Same concept. There, that's equal to 0 along that indifference curve. It's an incremental step along that indifference curve.

So solving this, this says that the slope of that indifference curve, dk dL, is equal to the MRTS, is equal to minus MPL over MPk. The slope of the indifference curve is minus MPL over MPk. Once again, should look familiar from consumer theory. It's the ratio of the marginal products that determines the slope of the indifference curve.

So for example, let's take our production function we're using here, which is a square root of L times k. That's our production function. Let's solve for the marginal rate of technical substitution. Well, we know what's the marginal product of labor. The marginal product of labor is the derivative of this production function with respect to L.

Take the derivative of that production function with respect to L, and you get 0.5 times k over square root of k times L. The marginal product of capital is the derivative of this production function with respect to k. You take that derivative, you get 0.5 times L over square root of k times L.

So that the marginal rate of technical substitution is the ratio of these two-- the opposite of the ratio, which equals minus k over L in this case, once again, I don't want to confuse people, this is the-- for this production function, that's the marginal rate of technical substitution. just like for this utility function, I got a particular marginal rate of substitution. That is-- I'm just showing you how to develop it for a given production function.

And that gives you something like figure 5-3. Figure 5-3 shows you the marginal rate of technical substitution for this production function. And you can see that at, for example, a point like A, you have of MRTS of minus 4. That means that you have very little labor and a lot of capital. That means you give up four units of capital to get one more worker. Why? Because marginal products are diminishing.

When you've got four shovels and one worker, you'd happily give up shovels to get more workers. So you want to move to the right. Likewise, at point C, you've got four workers in one shovel. That's doing you no good. You'd happily give up a worker to get a shovel, so you want to move to the left.

So basically, that is how we end up with this marginal rate of technical substitution. Questions about that? OK. Once again, parallel to the consumer theory.

Now, I'm going to talk about two things which are not parallel to consumer theory. Let me just say by way of introduction, producer theory is one step harder than consumer theory. It's all the fun stuff for consumer theory plus one additional step, and we'll talk about that step starting two lectures from now. But for right now, it's parallel to consumer theory, but here's a couple of wrinkles that arise in the context of production that you don't get in the context of consumption.

The first wrinkle is what we call returns to scale. Returns to scale. Returns to scale are the notion of what happens when I increase all inputs proportionally? That's an interesting question for consumer theory, but it's interesting for producer theory because I want to say, what happens if the firm gets bigger?

When you say in your mind, what happens when a firm gets bigger? You're thinking, what happens when it just-- when everything increases? When all inputs increase proportionally is our notion of how a firm gets bigger. And so that's what we call return to scale. Scale is basically when you scale up everything or scale down everything equally.

So, the question is, how does output respond? For example, it could be that when you double inputs, you double outputs. It could be that f of 2L, comma, 2k equals 2 times f of Lk. We call this a constant returns to scale production function. Double output, double inputs, double outputs. It's sort of a natural intuition. But it doesn't have to be the case.

Indeed, typically, we assume firms have we call diminishing returns to scale-- decreasing returns to scale, I'm sorry. Decreasing returns to scale. Where f of 2L, comma, 2k is less than 2 times f of L and k. That is, when you double inputs, you get less than twice as much outputs. Why is that? Well, think about the idea that basically, if you're still trying to dig a hole, there's one guy with one shovel, you go two guys with two shovels, they're in each other's way, they can't quite work as efficiently as one guy with one shovel because you're still trying to dig one hole.

So it's a decreasing returns to scale, doubling the enterprise. Or if you go to four guys with four shovels, they're definitely going to be just beating each other, the way they have to sit around for a while. So, doubling the enterprise doesn't double production. Of course, we can also consider cases of increasing returns to scale where f of qL, comma, 2k is greater than 2 times f of L and k. Increasing returns to scale.

That might arise through something like the fact that when the firm gets bigger, it can specialize more. So maybe as a firm gets bigger and bigger, it can get better and better at what it does, and production can become more and more efficient.

So, this is an important concept because basically, this is how we think about economies growing. This is the core of any model of economic growth. If you take macroeconomics, it's all about economic growth. The core of any model of economic growth is what is the returns to scale for the economy?

Typically, economists do not believe that there is everywhere and always increasing returns to scale. We typically believe in returns to scale may increase initially as you specialize as a firm, but eventually you have to decrease. And the reason we believe that, or at least, believed that, was that if there was always and forever increasing returns to scale, how many firms would there be in the economy? One. Because it could always just do better by getting bigger and bigger and bigger.

Now it turns out, for search engines, maybe that's not such a bad description. We'll see in the trial that's going on now where Google may get broken up. So it may be that increasing returns to scale may last longer than we thought as economists. But typically, we think of firms as maybe initially having increasing returns to scale, and then eventually as they mature, having decreasing returns to scale. Questions about that? Yeah?

STUDENT: So if a company like Apple, for example, innovates and they come up with a new phone every year, can they do [INAUDIBLE]?

JONATHAN GRUBER: Well, I'm going to get to that in a second. Not necessarily, and I'm going to get I'm going to get to that in one second. I wanted to-- because we haven't talked about innovation yet, and that's a missing piece of this that I want to turn to next. But other questions?

So I'm going to turn to the last slide. Before I do, I just want to review where we are. So basically, just like we have a utility function we have a production function, it translates your inputs into your outputs. In the short-run, labor is fixed, capital is not. This is just-- the only reason I tell you short-run cases is to develop the intuition of diminishing marginal products.

The interesting case, the long-run, where it's just like consumer theory. Just like consumer theory, you're trading off pizza and cookies to make you happy. Here, you're trading off labor and capital to make a good. And it's the same mathematics we saw before. And likewise, just as we had marginal rate of substitution determining the slope of the indifference curve, we have these marginal rate of technical substitution determining the slope of the isoquant.

Now, let me come to the question here and talk about a last topic, which is productivity. One of the most exciting topics in economics is thinking about productivity. One of the most famous-- and the idea of productivity comes actually from a very famous application from 1798 of the idea of decreasing returns to scale or diminishing marginal product.

Thomas Malthus was a very, I think, probably pretty depressed philosopher given all the stuff he wrote back in the 1700s. And he said, look, think about the production of agriculture. There's essentially two inputs. There's people and land. The people harvest the land.

He said the difference between capital and land is land is always fixed. Even in the long-run. There's only so much land. We got the Earth. Unless we go colonize another planet, which I don't think he was thinking about back then, there's a fixed amount of land, so that means that the marginal product of labor must be everywhere diminishing. That means that eventually, we're going to starve. There's only so much land.

As population grows, each worker is producing less and less and less on the same amount of land, and eventually we starve. So the principle of diminishing marginal product-- he didn't put it that way, but that was intuition. Basically, the fixed amount of land, an ever-growing population, each worker produces less and less and less until you can't feed yourself, until you produce less than is actually needed to feed yourself and people start to die off.

So Malthus actually said that there would be cycles-- the world would be marked by cycles of starvation. Essentially, population would grow, we'd be unable to produce enough food, a bunch of us would die off. Then we'd have a small enough population, we'd have marginal product of labor that could support production, and we'd go back to producing food again until we all died off again. Really not a very happy picture of the future.

Now, since Malthus wrote that book, the world population increased roughly 1,000% Roughly 1,000% increase in world population, and we are fatter than ever. So what did Malthus get wrong? I mean, not that there's not starvation. And let's be clear, there's horrible starvation around the world. Although let me also be clear, the famous economist, Amartya Sen, effectively won the Nobel Prize for an important observation, which is that there's never been a famine in a democracy.

So we think about famines and starvation all over the world, remember, that's a political failure, not technological failure. There's never been a famine of democracy. Famines only happen when corrupt governments don't spread the food they need to the people who need to get it.

So fundamentally, we have enough food despite the fact population has gone up 10 times. Gone up 1,000%, so 100 times. So what is it? What's happened? What's changed? What did Malthus mix miss? Yeah?

STUDENT: Technology.

JONATHAN GRUBER: Technology. Malthus mixed the fact that we are missing a key component of how we write down our production function. Because there's actually productivity improvements. In agriculture, what were they? Well, fertilizers, machines.

Basically, in the long-run, we have-- and eventually, disease-resistant seeds, better land management practices. He missed all the changes-- he was right that the amount of land is fixed, but he missed the fact that the way we use that land will get more and more productive over time. He missed all the technology that was going to come. He missed the productivity improvements that we were going to see.

Actually, world food consumption per capita is actually up since 1950 despite the fact that world population has grown. So what this says is really, we should really think about writing production functions as q equals A times f of L and k where A is a productivity factor. Basically, there's an extra piece which makes the world more productive.

This is not just true in agriculture. Let's think of one of the most famous examples, which is car production. Basically, before the 1900s, cars were made like other customized goods-- a single person or a couple of people would make it together. And then Henry Ford introduced the assembly line where basically, you had interchangeable parts, and you had a conveyor belt or assembly line where people could just specialize in putting this-- screwing this bolt in over and over and over again, and they could do it much faster, and they can specialize in it.

Now this seems pretty obvious at the time, but-- it's pretty obvious now, it was pretty radical at the time, he actually cut the cost of car production in half almost overnight. That was a massive productivity improvement. The same number of workers could produce many more-- same number of workers and same number of machines, you could produce many, many more cars.

Now, but this innovation didn't end with Henry Ford. The Indian company Tata produces the Nano, which is a $2,500 car. How is the Nano so cheap? Well, it's basically much lighter than other cars. It's much smaller, but they do things like putting the wheels way out on the edge instead of under the car to create more room for people to sit in. And they minimize the parts that are used to make it easy to exchange and easier to repair.

Innovation is going on, and of course, we're seeing perhaps one of the most major innovations in our automobile industry in decades, which is the electrification of the automobile fleet. And if you don't think Tesla's more innovative than an old Ford, you haven't seen a Tesla. Despite we think about Elon Musk, it's a pretty cool car.

So basically, we are seeing continuous productivity improvements, and that's what the standard model misses. Is that basically-- now this has become a really important area in economics because ultimately, A-- or productivity improvements become one of the key determinants for countries in their standard of living.

So with standard of living, think of it as how much stuff do we get to consume per person? Think of that as our standard of living. How much stuff do we have per person? What's going to determine that? Three things. How hard we work, how many machines we have, and how productive we are with those inputs. Those are the three things that will determine how much stuff we have.

What's great about productivity is these two are expensive. Working harder is costly. You've got to work harder, it sucks. Capital, you got to build a machine, you need to use stuff to build a machine. Productivity, boom, you got more. A brilliant idea, and suddenly, boom, you got more stuff for everyone. So productivity is the magic that makes economic growth happen almost costlessly.

And as a result, productivity is the central determinant of our standard of living. And it turns out, productivity has changed. From World War II until about 1973, US productivity grew at about 2.5% a year. What that meant was working just as hard with just as many machines, we could each consume 2.5% more stuff every single year just magically through the growth and productivity.

However, from 1973 until the 1990s, that fell to less than 1% of year, which meant if we wanted more stuff, we had to work harder, get more machines.

Now, one question people were asking was basically-- the IT revolution started in the late '80s, but productivity didn't increase. People were saying, well, where's the IT revolution? There's a great saying that says "you can see the IT revolution everywhere in life except the productivity numbers."

But then suddenly there was a boom. From the mid '90s to the mid 2000s, productivity jumped up again, and we thought, aha! The IT revolution is here. Productivity rose to about 2.3% a year. But after 2005, it stopped and productivity is back down again to a little more than 1% a year. So we're basically not really growing, we're not really innovating and developing and allow our standard of living to grow as fast as it used to.

This raises three key questions. The first question is, why-- why not long-lasting effects of IT. ET? Why do the IT revolution cause this brief burst and then eventually kind of wear off? Is it because basically all our innovation is watching TikTok and not in making stuff? I don't know, but that's one question we have, which is why is this incredible-- you guys have no idea how different life is now than it was 25 years ago in terms of the internet and everything you have, but somehow, we're not that much more productive as an economy. The question is, why not?

The second question this raises is where does productivity come from? Where does A come from? And here, this is a rich field of study because I said it was free, it's not free, it's not magical. Productivity is through real things. In the case of Henry Ford, it was through a brilliant insight.

So one place a comes from is education. As we educate people to make them smarter, make them more creative, they'll think of cool new ideas. Productivity comes from firm research and development. As firms do research and development, stuff that many of you may go on and do with your lives, they think up new ideas. And they develop new products.

It also comes from public research and development, like the National Science Foundation. Some of you may eventually get funding from the National Science-- may already have or may eventually get funding from the National Science Foundation for education. The National Science Foundation funds a lot of research. The National Institutes of Health fund a lot of research.

This turns out to be particularly critical because that's a lot of the basic research that gets done. Firms do what we often call applied research. They take basic ideas and turn them into products. But the basic research is often funded by government.

And so that's where that comes from. So basically, there are things we can do as a society to increase our productivity-- we don't just sit back and let it happen. Having more educated population, more investments in R&D. Research and Development, by the way, is R&D. Those can really make a difference in terms of how fast we grow.

Third question. Third point. How should we spend productivity? And here's what I mean by that. Imagine you have a cool new idea that can make things more efficient. Take Henry Ford's idea. He could do two things with that. One is he could make more cars. Two is he could give people more time off. If A goes up, you can reduce L, and with the same amount of output. There's no reason an increase in A has to go into more q. An increase in a could go into less L, which people like. They like time off.

One thing we learn in this class, the hardest thing about teaching a class at MIT, the single hardest thing is in the real world, people actually don't like work. I find this hard to believe, guys, but in the real world, people would rather not be at work than at work. And basically, people would rather say, well, gee, if you can make cars quicker, why am I working so hard?

And the answer is we don't know. There's no clear question about how to spend productivity. Indeed, the US and Europe have made radically different choices on this front. If you look at the decades since World War II where productivity has risen, the US has largely put that into more goods. Europe has largely put a lot of it into more leisure.

So a starting worker in Europe gets six weeks off. A starting worker in the US gets two weeks off. Basically, people in Europe work from about 25 to 60 and then stop. People in the US work from 20 to 65. We work a lot more hours per week.

The US has decided, we want-- well, it's been decided. I don't want to say it's necessarily a conscious decision, but our society is one which has evolved to be-- that all the proactive improvements go into more stuff. Whereas in Europe, a lot of it goes into more time off. What's the better thing to have? Good question.

When I was young, I probably thought stuff. Now that I'm older, I might think time off. But that's something we need to be conscious of, that our society-- that societies can take different paths.

The pandemic may have shook this up in the US. The pandemic was the first chance that people had to actually be at home and say, wait a second, maybe I don't want to work that hard. And there was a huge reduction in labor supply after the pandemic. A lot of people retired early. People aren't working as hard. People want to be in better jobs. People want to work from home. So the pandemic really shook things up in terms of people maybe starting to value their leisure more than maybe they did before. So that's the third question.

The fourth is who gets A? That is, which group in society benefits from more productive economy? So basically, if you look at the period from World War II to 1973, the average US income grew by 2.5% a year because of productivity. It turns out, pretty much every group in society saw that increase. The poorest people saw their incomes go up about 2.5% a year, the richest people saw their incomes go up to about 2.5% a year.

Since the mid 1970s, that has changed radically. Since the mid 1970s, virtually all of the benefit of increased productivity has gone to the richest members of society, the people that own the capital, and not to the workers.

Basically, if you look at the purchasing power of a typical wage today, it's not much higher than what you could buy with that wage in 1980. That is, the wage itself is higher, but goods cost more. We'll come back and talk more about inflation and things. The bottom line is, the average worker can't buy that much more stuff than they could in 1980.

Meanwhile, the wealthiest have gotten incredibly wealthy. The share of wealth in the US, the share of income controlled by the top 1% of people in the US was 26% in 1989, and it's 35% today. Even more, in 2021, in the coming back from the pandemic, in the year after the pandemic, as we readjusted to life, the top 1% of earners saw their incomes increase by 9% while everyone else saw their incomes decrease. So top 1% of earners were the only group that actually saw their income increase. Everyone else saw income decrease because inflation ate into people's wages.

The bottom line is, we also, A doesn't magically go to everyone. It's not like a more productive economy, everyone benefits. That's determined by a set of decisions made by individual actors in the economy and by the government. And that's what we'll talk about later in the semester, is how does the government think about those decisions? Or how do we think, as a society, about how to spend the extra resources we get when the economy gets more productive?

And that comes to the topic of economic fairness, which we won't spend nearly enough time on, but we will come back to towards the end of the semester. Questions about that?

All right. Why don't we stop there? And there's no section Friday, right? No section Friday. Remember, the first problem set is due next Friday, and the section you go to next Friday is the one you're going to stay in for the semester. OK, thank you.