[SQUEAKING]

[RUSTLING]

[CLICKING]

AUDIENCE: Hi, Professor Townsend. Can I ask a question about the income factor on the Slutsky equation?

ROBERT TOWNSEND: Yes.

AUDIENCE: Yeah, I noticed that in the last lecture slides, the income factor, you'll note that it was x multiply data x over [INAUDIBLE]-- why it is not minus x modify data x over data y.

ROBERT TOWNSEND: You're wondering about the sign.

AUDIENCE: Yeah.

ROBERT TOWNSEND: Yeah, it is confusing because the way that equation is written, it looks as if we're-- I can't remember. Which way is it written? Minus, right?

AUDIENCE: Yeah, minus x.

ROBERT TOWNSEND: Subtracting the income effect, but the graph on the previous slide says we're adding the income effect. If prices go down and income effectively increases, then you buy more than you would if you held income constant. So the answer to the puzzle is the demand curve figures have reversed the dependent and the independent variables.

That is to say, we usually write dy dx, where y is the dependent variable, x is the right-hand side independent variable, but the demand curves actually put price on the y-axis. So it's the inverse slope, if that makes sense. So the slope dy dx, as written in that Slutsky equation, is the inverse of the slope in the diagram.

AUDIENCE: OK, thank you.

ROBERT TOWNSEND: Yeah, don't worry. I stumble over that myself. I was pondering it, and I actually asked Michael to look at it. And I'd forgotten that strange labeling that we use, so it is a little bit unusual. Thank you for asking the question.

I do always start the class. So you've got us off slightly early, but that's fine. Are there other questions? The reading list for this lecture, it's just one starred reading here, Kreps chapter 7. So again, a lot but not all of the material today is in that text, and it's starred for a reason. I think you will find it to be very helpful to read through it. Let's see. So that's that. And then the study guide on what we did last time.

OK, income and substitution effects-- can someone volunteer to answer this request to discuss Engel's law? Be as clear as possible about the data used, and what you would find if you look at the data properly?

AUDIENCE: Engel's law is the law it says that the higher proportion of your income is spent on food, right?

ROBERT TOWNSEND: The higher is your income, some good expenditure shares go up, and some expenditure shares go down. So the exact statement of it was, I guess, in this figure with the data. And food is a necessity, so the consumption actually goes down or rises less rapidly than income does. So thank you. So that's where it is in the lecture.

One comment I have, this is about expenditure shares. It doesn't say the quantity of food. It says 3 times the quantity of food is going up. So expenditures for food are going up. It could be a normal good. But it's not going up as fast as total income is going up.

So last time I drew what I called an income expansion path, which traces out the locus of tangencies as income rises, holding prices fixed. And if they were linear, then the expenditure shares are constant. That is to say, income goes up. The amount spent on the first good goes-- the amount spent on the second good goes up. Everything is rising proportionally.

In order to get the expenditure share to be not rising as fast for food, you have to have those income expansion paths be somewhat curved so that you don't get a proportional rise in food. And really, answering a question that was asked in class last time, or at least what I thought about afterwards might be helpful.

Another way to say this. We talked about Giffen goods. Oh, well, maybe that's a question. So someone volunteer to tell me, what is a Giffen good?

AUDIENCE: It's when the price and the quantity demanded move in the same direction. So if the price of something goes up, then you want-- then you buy more of it.

ROBERT TOWNSEND: That's right. OK, perfect. Perfect answer. So the point is, the quantity is what we have in mind when we talk about Giffen goods, not price times quantity, not expenditure shares. What creates a Giffen good is a good which is not normal, that is to say, where expenditures are falling, relatively speaking, as we increase income. And the point is there's two effects of price changes.

Engel's law is about income changes. Giffen's paradox is about a price change. But the price change can be broken into two pieces, the substitution effect and the consequent income effect. And the substitution effect does mean that as price goes up, you would demand less.

But the income effect means that as price goes up, you have effectively less income. And if it's a necessity, then you're going to spend relatively less on that good. So that's the-- tends to make the quantity go down, so it's about income effect dominating the price effect with implications for quantities. I hope that's helpful.

Another way to say this is when we look at income expansion paths just changing income, we're not moving the price around at all. And we are assuming somehow you can make comparisons across to different income groups, that they're all facing the same prices. It's just that some people have more income than others, which is effectively what would happen for a given individual if somehow a poor individual were given more income.

But it's all about moving income around in the cross-section or as if, over time for a given individual, holding the prices fixed, whereas the Giffen good has to do entirely with just changing prices. Income changes only as a consequence of the price change, and the Giffen good is about quantity. So I hope that's helpful.

Can someone do the last one here? Compare and contrast the utility maximization problem with the expenditure minimization problem. Essentially, give a statement of duality. What does duality mean?

AUDIENCE: Yep, so the duality just means that these two ways of formulating the consumer's choice problem are, in fact, just identical. Utility maximization says given a budget, what's the highest utility I can achieve? And expenditure minimization is given a level of utility I want to hit, what's the minimum budget I need to hit that? And these two will, they're basically identical problems of each other.

ROBERT TOWNSEND: Perfect. That's a perfect answer. Thank you. I can't add anything. OK, so as always, it's quite apparent these are review questions that you should go over after class and before the next class to make sure you're really following the material to go with the production lecture, lecture 4.

To get reoriented, we have an economy, which consists of preferences, endowments, and technology. Last two lectures, we focused on preferences. Oh, I should have also said, the application at the end having to do with borrowing and lending, in addition to labeling the goods by the date, date 1 and day 2, it was quite apparent that we had also actually plotted the endowment. So that was a bit different than what we had been doing with good x and good y, where we just talked about income.

So in addition to having income i as exogenous in that diagram, we had the endowments, the amount of money they had in day 1, the amount of money they had in day 2, and plotted the endowment point in the diagram. Which is actually, either way is fine. But we did add more information. Anyway, you've seen endowments and preferences. Today, we're going to do technology.

In choice of overproduction plans, we'll explore some possible properties that technology sets can have, as we did for consumer preferences, so on. And then we'll go to the application, which is going to be similar to what we did for households. We're going to have firms maximizing profits given a production possibility set. And we'll take that into an application. And actually, by the end of this lecture, we will actually look at our first entirely specified economy, Robinson Crusoe economy, and begin to talk about international trade.

So this slide is meant to portray technology. It's a picture I took, actually, in Northern Thailand, where I was doing my field research. And if you're really into economics, you'll start talking about land, labor, and capital as inputs into production.

And here you see all three. The land is kind of obvious. He's getting ready to plant. Labor, this guy, who's showing off, by the way, because he's not using any hands. He's not grabbing the handlebars on what you might think is like a lawnmower.

And then this tiller, which they call an iron buffalo, replaced the buffalo that used to do the plowing, and it's just this small horsepower motor driving the treads on the wheels. But anyway, that's the capital-- land, labor, and capital. So a picture is worth 1,000.

Anyway, production, OK-- so firms possess production capabilities, meaning some commodities are inputs. Others are outputs. And the capability is described by a set of possible values, which you can think of as vectors with inputs and outputs, or some people say netput. If we have the number of goods, including inputs and outputs, being of dimension capital K, then a vector has dimension capital K, and the set of all feasible production vectors is called the production possibility set.

Now, we do have a convention, namely, inputs are negative, or at least not positive, and outputs or positive, or at least not negative. And with a couple of exceptions, I'm going to continue that. It is easier to interpret if we designate the inputs and talk about labor and capital and land as being positive into a production function to get rice as the output. But for reasons of very succinct accounting, it's quite useful to have the inputs as negative and the outputs as positive. You'll see more of that powerful notation a little bit later.

Anyway, so for example, if we have three commodities, two of which are inputs, say, labor and steel, numbered 1 and 2 in the vector, and the output being safety pins, then, for example, we have a netput vector of minus 5, minus 8, and 10. And the inputs are each negative, and the output, the steel safety pins, is positive. So not too hard.

More generally, Z could be a production possibility set. Sometimes we call it Y. I apologize. We used different sources here. We can write Z from 1 to K as this vector, where, again, K is the number of commodities. And if we wish, we had a firm J, we could index everything by J and talk about the production possibilities set of firm J and the inputs j1, j2, through jK. And then we drop the j for most of this talk.

Now, this is very similar to what we did for consumption. We had x1 through xL, I think it was. And then we put an i for [INAUDIBLE] needed. So here, we're dropping the j.

Here's a picture of a production possibility set, Z, where we have one input only, and it's negative. The more negative it is, the larger the quantity being utilized is. And likewise, if you use more of it, you get more output.

This entire speckled region is illustrative of the entire production possibility set. You don't have to be on the boundary, which is the maximum output you can get for a given input in order to be in the set. That boundary is typically included in the set, but so are these interior points.

So here are possible properties. And again, just like for the household, we didn't insist that every class of every production possibility set contain each of these. In fact, some of them contradict each other, as you'll see in a minute. So we usually assume it's not empty. I mean, why are we writing it down if there's nothing to do?

Closed means it contains its boundary. That's one way to say it. Another is just that every limit point in-- every sequence that consists of points in the set that converges has a point that is also in the set. That's a technical condition. I'm sure you've noticed I'm not overdoing the continuity part and the closure part. It's a bit more technical than we need, and these properties are pretty intuitive.

No free lunch means that whenever Z is in the set and has no negative components, you can't get any output. So Z is a vector. Z can have positive and negative components. Here, writing this Z vector being non-negative means that no component in the vector is negative, which means they're either 0 or positive. And if they were, say, all 0, then you have no inputs, and you can't get any output. You can't get lunch out of that. And that's the expression "no free lunch."

This is used by cynical economists to mean something else, which means if someone offered you something for free, be careful because there may be [INAUDIBLE]. Anyway, I won't share that too much.

Convexity-- if you have two elements, Z and Z prime, in the set, then the alpha-weighted combination is also in the set. We did this with consumption sets, so this is not surprising. Free disposal means that if you have a vector Z in the production set and Z prime is not higher-- say it's less than or equal to Z-- then that Z prime is also in the set. This is called free disposal.

There are two ways to look at it. It's kind of the same. One-- suppose Z prime had the same amount of all the negative inputs, quantities of all the inputs, which are negative, but at less output. So the positive components are not greater and maybe extremely less. Then you're perfectly free to throw away that extra output and maintain yourself in the production possibilities set.

A little less obvious, if Z prime were say strictly less than Z because the input quantities are actually higher-- and again, the negative makes it look like it's less. Higher inputs means a bigger negative number. Then that's also in the set. Why? Because if you've got more inputs-- say you're maintaining the same output. Well, if the inputs were at all productive, you would have had more output. But you throw away that output. That's the free disposal again.

And then, finally, you can always shut down. Most, I think, production functions, production sets have that property, that 0 is feasible, which is a bit different from saying it's empty. It means 0 is in there, but other things are usually typically also in there.

So three other properties-- decreasing returns to scale. And I'll show you some pictures. If Z is in there and alpha is less than 1, then alpha times Z is also in there. So decreasing returns actually means you can scale back something in the production function, like US Steel blast furnaces. And somehow they can pare that down and make a smaller one. So it may not be true, but that's what the property is.

Constant returns means that you can scale up or down willy-nilly and still have the vector alpha Z in the production set when Z is. And increasing returns goes the other way. If you have Z in the set and alpha is greater than 1, then alpha times Z is also in the production set, which means, in this case, to continue the example of US Steel. If you still has one big blast furnace, they can just as easily have two or three. They can run them.

OK, so let me draw you some pictures. This is decreasing returns, which intuitively comes from the curvature declining of the boundary of the production possibility set. The property says if you're given Y in the set and it's not 0, then alpha between 0 and 1 times Y is also in the set, which these things are.

This is increasing returns, which is, whoa, getting steeper and steeper, and then there's [INAUDIBLE], but it's instructive. So this is the same alpha Y, yes. But the property of increasing returns says if you have Y, say, here and alpha is greater than 1, then alpha lies in the set. So, actually, the error-- we'll fix this, I guess-- is that the alpha Y should have been depicted out here in the production set. All this area is in the production set.

And then constant returns is neither of those extremes. It's just a linear boundary. And it says if you pick anything in the set, say, here, and then draw a line, a ray, from the 0 through that picked point, everything on that ray towards 0 or out to infinity is also in the set. And I guess we should have put alpha down here instead of lambda, but it's kind of obvious. So constant returns to scale is going to have some strong properties, which we'll come to.

One word about aggregating up production sets-- so this is one of the exceptions, where what-- well, this is the classic picture. Inputs on the x-axis. Output on the y-axis. This would be the production set. Here's a point on the boundary of one firm's, labeled number 1 firm here. Also, the exception is we're labeling firms 1 and 2.

This is the production set of another one that actually uses as an input the output that this produces. And I'll come back to that later. In this case, the inputs and outputs are reversed.

And finally, we have an endowment point. So we wanted to add all these things up and get the aggregate set of possibilities. You got to do a lot of vector algebra. But essentially, what's going on here is start at 0. That's feasible. It uses less than endowment. That's fine.

Then find y1 on the first production set-- y1. Then add that endowment back in. Notice this, the parallel line shift. And then, so far, we've been keeping y2 at 0, but let's move y2 to something non-trivial production.

That movement here is the same as this movement here. So at this point, which is, of course, the sum of y1, y2, and z1, is feasible because it comes from these three possibilities. So this larger area out here is the aggregated or aggregate production set.

Now let's talk about the firm's problem. So we're already getting into an application now. There are these prices, p, which the firm takes as given, just like the household took prices as given. And the firm wants to come up with a plan, input-output plan, to maximize its profit. So this capital pi here is just a statement of prices time pk times zk, summing up over all the k's.

And here, again, it looks like everything's positive. What a wonderful world. But some of these z's are required to be negative. And hence, [INAUDIBLE] times a negative object. So there's expenses in here and revenues in here.

And then when we talk about the firm's problem, which is FP, for Firm's Problem, is to maximize the profits by choosing feasible points vectors in the production possibilities set. So this very cryptic abbreviated notation, max p of z, [? proves ?] z and z.

And here's our picture. So inputs, which are negative in sign, outputs, which are positive. These guys here, these linear hyperplanes, represent possible profit levels. Call them isoprofit lines.

When we introduced indifference curves, I said it might be better to call them isoutility lines, and I was anticipating this slide and another one that's coming. So an isoprofit line means as you, say, use more inputs but get more output, you have expenses balancing revenues, and you stay on this line. If you go put in lower inputs, so for the same prices, you would have higher profits.

The slope of any isoprofit line is just price of output divided by the price of input, like p2 over p1. And of course, we're depicting this tangency as the maximum profit line. You could be in the interior, but then you would not be maximizing.

So let's call little pi of p to be the maximized profit, not just any old profit, which was capital pi a minute ago. This is the best choice of production vector z in the production set, given price is p. p is a parameter here. This is, quote, partial equilibrium, so we index the solution by p. Of course, we'll do experiments where we vary p.

There are some profits-- properties of the profit function, namely, this profit function pi is homogeneous of degree 1 in the price vector p. Homogeneous of degree 1 means that if you scale prices up by some proportion, like doubling them, then maximized profit should also double.

The reason that should be obvious, at least when you think about it, if we're doubling prices, we're not changing the slope. So we're not changing the maximizing choice. But we have changed the levels of the prices, and the levels means even though the z's aren't moving, the prices are moving. If prices have doubled, then it must be that profits have doubled and constant returns to scale. Homogeneous to degree one, sorry-- homogeneous to degree 1.

Profit function is continuous. It's actually true much more generally than you might think, but I'm not into emphasizing continuity too much these days. Profit function is convex. That's an np. So that's an interesting property. So let's look at that one.

The profit function is convex. OK, what do we want to show? If we have two p's-- two prices p and [INAUDIBLE] the associated profits, pi of p and pi of p prime, then we take the linear combination of the prices, say, p double prime, this alpha weighted sum, the associated profits should be not greater than the weighted sum of the profits. So I said that in words. It's probably just a whole lot easier to do the mechanics.

So we start with two prices, p and p prime. We pick an alpha. We consider p double prime to be the alpha-weighted sum. We let z be the solution to the firm's profit maximization problem at p double prime. So in other words, that special maximized profit pi of p double prime is just p double prime times z because z is the solution.

Now, z, therefore, was feasible. It was in the production possibility set. So even though we changed the prices, that z remains feasible, which means whatever solves the problem at p, it cannot deliver a lower number than this and may deliver a higher one. The profit-maximizing solution at p might involve a different z different from z. And likewise, the profit-maximizing solution at p prime can only be not lower and maybe higher than p prime evaluated at this z.

OK, so now we do the math. We want to show that pi at p double prime, which is, by definition, p double prime at z. Remember the starting point. We're starting at p double prime, and we're told that z is the solution. So this is p double prime z.

But p double prime written out is just alpha p, 1 minus alpha p prime-- that was the definition-- times z. And now we invoke each of these two weak inequalities and say, well, alpha p times z is alpha pz less than pi of p. So the inequality is less than or equal to, and we get the same thing for p prime.

Sorry, I think the ratio of my words to what's on the slide-- I must have talked for three minutes. But there's only four lines here, so I apologize.

Alpha function is convex, maybe weakly, but it's not concave, or not strictly concave. It may be convex or strictly convex. OK, so we'll come back to that momentarily.

Homogeneity-- two more properties, and back to convexity. So if f is a real valued function, mapping rn into r and k is an integer, we say f is homogeneous of degree k if whenever you scale the arguments of the function up by-- the arguments of the function v by alpha, the value in the domain goes up by-- in the range goes up by alpha to the power k. So homogeneous of degree 0 would mean nothing changes. Homogeneous of degree 1 we just did for profits, where k is equal to 1, and this is the most general concept.

Another property of the profit function-- it's continuously differentiable. So I'm going to show you an expression with the derivatives. If you take that function, which has as a parameter the price vector p, and differentiate the maximized profit with respect to, say, the k-th item, pk, holding all the prices-- doing it locally around the maximizing value p star, then the solution is the input. You get the input back. Derivative with respect to pk, you get the input zk.

A quick, maybe misleading, intuition for that is you write down profits-- p1, z1, p2, z2, et cetera, et cetera. Take a derivative at the maximizing solution The derivative is with respective pk. Where does that show up? It shows up in the pk, zk term, so it seems like this ought to be true.

Actually, what it's using is the convexity. So this thing is called Hotelling's lemma. Here is the actual profit function, but now with all the prices. It's actually got a star on everything but the k-th price here, no star. But we're going to evaluate profit at star k. So we're right here on this profit function. This would be the maximized profit in general, and also the maximized profit in particular when the price of the k-th input or output is pk.

What is this n? And this is like a tangency? It's a straight line. And what is this straight line? That's just what I was saying in words. That's just the evaluation of inputs and outputs at this price vector star, and it starred everywhere. OK, so pk doesn't have a star here. But because we're plotting everything is a function of pq, but there is a particular pk, namely pk star, where we can evaluate that linear function.

Now, the point of the tangency is that the curved function, the profit margin, and this linear function have the same slope. So if we want to know how the profit function is changing as we change pk locally, we can just as easily show how this straight line is changing as we change pk locally because they have the same slope.

And how is the linear function changing as we change pk? Well, it's just changing linearly. So that's the derivative, namely zk. So that's the proof of this. So we did actually use the convexity of the profit function, and the picture is strict convexity. It actually works more generally, but it would be a pretty hard picture to interpret.

Now, this is also illustrative of something else called the envelope theorem, and I'll just show you the definition. Oh, and remember the strategy here. Tools-- so every so often, we bump into a concept that generalizes. In this case, we just went over some example of something called the envelope theorem.

So let me give you a statement of the general theorem. If fact, [INAUDIBLE] [? set ?] and let t be some parameter of the problem. We have a function f mapping elements in the choice set in that parameter into some real valued objective function, namely like this. So put max here if you're more comfortable.

So f is the thing we want to max. We want to max it by choice of x. But we have a parameter in the problem called t, so we denote the solution, the maximized value of f, at t to be V of t, the value function at t.

And what is X star? X star is just the set of all maximizers X. It could be unique. There could be many. But each one of them, if there are many, has to have the property that evaluating the function f at that maximizer x gives you V back. And again, everything is indexed by t.

Now for the big E. If you took the derivative of the value function with respect to t, even though t is a parameter, it will be the derivative of the underlying function you're trying to maximize at the optimized solution. That is for an x in the set of maximizers at t.

So one thing to say this is an example of that, right? We wanted the profit function, which is parameterized by not t, but p, and we wanted to know how it would behave at the optimum if we were to slightly change the parameter p. And we just took the derivative of the objective function at the optimizing solution, which here is what this line is at, the optimizing solution. And the derivative was just z.

Well, at the risk of confusing the issue, there was this chat going on about lambda as the margin utility of income in the consumers problem. And, again, income was a parameter of the problem. And nevertheless, we wanted to know how utility would change-- in this case, utility f-- how utility would change as we change that parameter income i.

And the answer from the envelope theorem is it is the derivative of a maximized objective function-- that is to say, of utility substituting in the budget constraint-- at parameter i. So that's consistent with this language of the marginal utility of income. In the chat, there is a great discussion about that Lagrangian, which is so powerful it's mysterious, because the Lagrangian has those Lagrange multipliers types constraints, not just the objective function.

So the chat was, it's not really marginal utility because it's a Lagrangian. It didn't just write down the utility term, et cetera, et cetera. But you can see how it-- hopefully get a little more intuition about that problem from the envelope theorem.

OK, so, as I said before, in many applications that we're clear about what's an input and what's an output, we could have, for example, n possible inputs and m possible outputs, and then we'd be writing these vectors as negatives for the n inputs and positive [INAUDIBLE] m outputs [INAUDIBLE] some degrees of freedom, depending on the production function. The other things being 0. There are things other firms can do this firm cannot.

So we could maximize if-- even easier. There's a single output and capital N possible inputs. We can talk about a production function. Rice is a function of labor, holding capital and land fixed, for example. Or we could have multiple vectors x-- land, labor, and capital. And we write down this output function, this production function.

We will have prices of the output p, say, taken as given. We'll have the prices of the N inputs, W1 through wn taken as given. And we can write down the profit maximization problem, now with a bit more intuitive signs. It's revenue, price times output, minus the cost of the inputs that were used to create that output at the price is wn.

If f is differentiable, then we can look for the first order conditions. That derivative's equal to 0. Say, in particular with respect to the nth input, we'll get p times the derivative of [? f of ?] xn over here, and we'll get wn here. That's just expression.

So the shorthand way to say this is the value of the marginal product or the marginal revenue product is equal to the price. That's economist speak. If you think about what would happen when this is, say, greater than the right-hand side, that means one more unit of input costs you wn. But that input applied to the production function gives you a marginal output, and that marginal output evaluated at the price gives you the marginal revenue.

So if the left-hand side is greater than the right-hand side, you can't be done because you get more out of increasing your costs. You more balance your costs by increasing your revenue. And likewise, when the inequality goes the other way, you've gone too far and used up too much of the inputs. So this is a very intuitive expression.

So now let's talk about cost minimization. Instead of solving the global problem and finding the maximizing profit, maximizing inputs and outputs, we're going to break the problem into pieces and take the output y as a given, like something exogenous. Got to do it. Got to achieve it. But let's achieve it in the cost minimizing way.

So here's a statement of the problem. If this V of y is the input set that delivers y as the output, taking y as given, we can choose any input vector, as long as it gives us the output of y. So let's minimize the cost, the dot product w times x, which is the cost for achieving that output y. And let's call that the cost function. It's the cost of producing y in the least possible cost way, given input, price, vector w.

The price of output doesn't matter here because that's a given anyway. That is to say, price is not only given. Output is given. So revenue is not in question here. It's just the cost of financing it through the inputs. So we'll call the solution conditional factor demands. And in principle, they're a function of these parameters W and y.

So let me draw you a picture. In this case, there's going to be an output level. That's to think of it as a single good. q is the level. And we're going to plot-- and there's two inputs, x1 and x2. And we're going to plot here the set of all inputs that give you the same output. That's that curve.

Since the quantity is the same, namely q, we've got a name for this guy. It's isoquant, which is the final instance of iso this and that. So isoquant for same quantity as we move along this isoquant.

And what are these straight lines? Our expenditure curves-- they're down here. w times w1 x1 plus w2 x2 equals c, where c is some level of cost. The higher ones are associated with higher costs. To minimize [INAUDIBLE] one, move towards 0.

So the solution is here. It's the minimized cost that allows you to be on the isoquant. In fact, the way it's drawn, and the other point on the isoquant would give strictly higher costs. You with me?

So you can also characterize this tangency. The slope of this cost line is w1 over w2. The slope of the isoquant is like a marginal rate of substitution, except we call it the marginal rate of product transformation. It's the rate at which you can give up the use of input 2 and use input 1 in such a way as to maintain quantity constant. So this is very similar to the indifference curve that we had in consumer theory. And this characterization of the tangency is similar in terms of the ratios of the two tangent lines being the same.

AUDIENCE: So the thing that's getting toggled in this case is the budget line, right?

ROBERT TOWNSEND: The thing that we're trying to minimize is cost. So this, instead of being a budget line, it's an isocost line.

AUDIENCE: Yeah, yeah.

ROBERT TOWNSEND: Yeah, and we're trying to move this way to make it as low as possible. Yeah, so that's another difference. When we did utility max, we had a budget, and we were trying to maximize. Although, we did in the duality lecture talk about minimizing the cost of achieving a certain level of utility, and this is very close to that. So here, we're minimizing the cost of achieving this level of output.

OK, let me just say in words, this is a particular q. I could have picked another one, which would minimize the cost of achieving that q. And then we would get cost curves, which are going to show up momentarily. But that was too much clutter to put on the diagram, so we didn't do it.

Instead, we put on this, the idea of changing the ratio of costs. So say Disney Company producing movies uses humans, animators, and computers. [INAUDIBLE] the year, computers got cheaper relative to the wages of humans. So in the early 1900s, 1990s, the slope of this line would be the cost of computers divided by the cost of animators-- relatively flat.

By today's standards, the line is steeper. So the cost of computers, well-- it's-- always have to remember, it's like p1 over p2. So the cost of animators relative to the cost of computers has become a lot steeper.

And of course, you're seeing different tangencies A and B. Naturally, Walt Disney Company would switch in the direction of using the cheaper input, namely, computers at the expense of using more humans. So later, when we get into international trade, we're going to talk about-- we're going to talk about changing factor prices and wages and so on, so this is motivation for that when we get there.

Cost function is homogeneous of degree 1. If you double the inputs, you double the cost. Same argument as before. We're not going to change the optimum. We're just changing the prices of the inputs, so the cost doubles.

The cost function is continuous, and the cost function is concave. Just like we did that the profit function was convex, we're getting that this cost function is concave. And the argument is very similar. And instead of chewing up time, I'm actually going to break precedent and leave this for you to do, but it's exactly the same argument that we went through with the convexity of the profit function, except the signs are-- the inequalities go the other way.

OK, and then we get Shepard's lemma, which is a version of Hotelling's lemma focusing on costs, and hence only on the inputs prices. If we differentiate the minimized cost of producing output level y with respect to an input price w-- I pride myself on the ability to go through these slides and remove all the typos. But here, I'm kind of wishing that I had had a wn on here because it's xn.

So if there's only one input, it's fine. If there's multiple inputs, then there's a vector, and we should be differentiating with respect to a particular item in the vector, namely, the nth input. So probably better to have written this as wn and have an xn over here. Anyway, hopefully, we'll take notes and correct that. It's not wrong. It could have been written better.

But the real point is you've seen this before. This is a version of Hotelling's lemma. It's called Shepard's lemma. But Shepard's only dealing with cost function. Hotelling was dealing with the whole profit function. So all that stuff about the envelope theorem, blah, blah, blah, that all applies here.

Cost curves-- so we've got the total cost of producing q as the minimized cost of producing q. The marginal cost is the derivative of that. And the average cost is total cost divided by the quantity q. That's the cost on average of the various units of q.

Of course, profits would be just the revenue minus the cost. This is the minimized cost function suppressing the reference to w as if it were fixed. And take the derivative of this, of course with respect to q, we're getting a price on the left-hand side and the marginal cost on the right.

So again, this has an intuitive-- you've probably seen this in your principles of economics course, that when the firm is optimizing, it should pick a quantity q, where price is equal to marginal cost. And the reason is, of course, if price were greater than marginal cost, you get more in revenue for one more unit than you sacrifice in terms of costs and vice versa when it goes the other way.

In particular-- we're going to get a strong property here-- when we have constant returns to scale, this production function you've seen before, production set. And the cost is going to be linear. It just scales up and down as you vary q. And in other words, the marginal cost is equal to the average cost, and we want to find the point when you're optimizing the price p is equal to the marginal cost, which is also equal to the average cost.

So any price less than that, you're like, no, I'm not doing it. And any price greater, they would go to infinity and make an infinite amount of money. And at this critical price p, which is exactly equal to the constant marginal cost, because the slope is the same on this line everywhere, they jump from 0 to any old q that you want.

Let me show you the intuition of what's going on here. Again, keep track of what's in the numerator and what's in the denominator. This is supposed to be p1 over p2. The slope of this isoprofit line is p1 over p2. If p2 is going up, this isoprofit line slope will be going down.

So let's compare the three diagrams here. So here, given a slope associated with p1 over p2, profits just get higher and higher as you move further and further along the boundaries. So 2 is the output. So as p2 goes up, p1 over p2 is going down. So we get this isoprofit line, which is actually coincident with the boundary of the production set.

And finally, for-- so the mistake here is this should say going down. I'm constantly catching myself that it wasn't working that the slope is moving in the direction of the opposite of what I'm saying. What we want here is a price where they go-- or infinity over here. They're indifferent at this price of output. And at a lower price of output, they, don't want to do anything.

So [INAUDIBLE] nothing [? pro ?] profit. And that's where these lines are steeper, where these isoprofit lines are steeper. And they're steeper because it has a slope of p1 or p2. And the price of output p2 has gotten so low that these lines are getting so steep that the best position is zero profit, so I finally set it right that time. And I did not catch this typo before. We'll fix that.

The point of the diagram is constant returns to scale has some very special properties. Prices have to be just right. If price ratio is just right, then any production is fine, and they're making zero profits. So they don't care. They will do something if demand makes them do it.

But on these two extremes, this isn't feasible because it's infinite profit. And here, they would make losses, so they minimize their losses and go to 0. So you get these very extreme solutions.

If you have an equilibrium with positive profits from a production set which displays constant returns to scale, you know that the price ratio has to be coincident with the slope of the boundary of the production function. When we get to equilibrium, we can determine prices from the production side, at least on the premise that something's being produced. That's what I said.

This is a property of constant returns to scale. Namely, we might just as well solve the cost minimization problem for an output of 1 and get the ratio of the inputs that way. And then, if you want to scale production up or down, y being greater than or less than 1, that's fine.

You'll change the level of the inputs. They're going to scale up and down as well by the same scalar you're using to scale y, but the ratio isn't going to move. But that's a property of constant returns to scale.

Now, again, I had decided this in advance and, indeed, have my eyes on the clock here. There is a proof here. So if you like, take some time and try to write out a proof by contradiction. I think you'll find that to be a useful exercise. In other words, prove something like suppose you changed y and the input ratio were to change, then that would lead to a contradiction, something inconsistent with constant returns to scale.

All right, here's Robinson Crusoe. It's like a one-island, one-man economy. If you don't the story, this means nothing to you. I think he actually had a helper called Friday. But anyway, just think of Robinson Crusoe being a one-man economy, and we plot the production possibility set with Robinson's input, getting food as output, picking fruit off the trees, whatever.

And what's the maximal solution for Robinson? It's to find the utility function which has the highest tangency. So these lines are the curves of indifference curves, and the maximal utility level is [INAUDIBLE].

So this is kind of a big moment in the course, even though the concepts probably look pretty transparent. We just combined utility maximization with production. And I'm telling you, the optimum point for the whole economy is this tangency level.

And, actually, if we drew a line that were tangent to both the production boundary and the indifference curve, that could represent both a profit line for Robinson Crusoe, the producer, and a budget line for Robinson Crusoe as a consumer. So for the first time, we're actually looking at a competitive equilibrium for the whole economy. It's kind of stupid, though, because Robinson's all alone. There's no one to trade with.

But suppose he could. So suppose Robinson Crusoe could produce two goods, some agricultural output, but also some tools or some manufactured good on the x-axis. So the set of production possibilities is this quarter circle area with its boundary. And if Robinson Crusoe would be still intent on maximizing utility, we'd find where the slope of the-- where the-- we'd find the highest level of the indifference curve and the tangency.

There's stuff going on in the background here I'm not showing you to make it simpler. I'm not talking about his input. [INAUDIBLE] whether I'm going to work in agriculture or work in manufacturing, and I'm not showing you that work level. But it is producing this set of production possibilities of the two goods, and this would be the utility maximizing point A.

Now, suppose all of a sudden, he has the opportunity to trade with another island. And suppose we represent the isoprofit line in red. So we want to be as far to the northeast as possible, maximizing profits, but staying on the production possibility set. So that maximized level of production is here, B, and consumption is here, C.

Now, this is a much more interesting economy because there's some external agent here. [INAUDIBLE] Robinson Crusoe and the rest of the world, that other island. And they're now trading with each other. And Robinson Crusoe has decided to specialize more, though not completely, in manufacturing goods and to give them up in trade to the other island in order to acquire more agricultural output.

So we're seeing exports and imports as we move from B to C. And, obviously, Robinson Crusoe is doing better utility wise from trade than that autarky. A might stand for autarky. So this is our first glimpse at the gains from trade. With the one-man economy, at least, Robinson Crusoe can be made better off by trading with another island.

This is a related picture. Suppose we have two countries, the UK and the United States, and their production possibility set look different from each other. We're plotting wheat and cloth in both countries, and we might have an autarky that the UK's maximal utility position is here at DUK, and that the US's no-trade position would be at DUS.

Now, you do see that this is deliberately chosen so the slopes of these lines are different, and that suggests there are gains to trade. In fact, what's depicted here is almost common price ratio, where you could think of it as the prices in the rest of the world. But this may be a bilateral country example, but somehow the prices are fixed for each country by competitive forces. So the slope of the dashed line is the same for both countries.

And what happens? Well, the UK decides to export wheat and import cloth. I thought it was going to go the other way, actually. And the United States chooses to-- did I just say that wrong-- chooses to-- oh, yeah, I did say it wrong.

The US chooses to produce more wheat and export part of that increased production back to the UK. The UK was here in autarky, when it faced with the possibility of trade, it produces much more cloth, and then it exports that cloth to the US. Cloth and the UK kind go hand-in-hand because people think of the Industrial Revolution and manufacturing of textiles, which was kind over big deal in Britain, and the US and Kansas and so on having lots of wheat. So it's meant to be a somewhat realistic example.

Anyway, this is the climax of the slides today because you're seeing an application of production theory, in particular, a profit maximizing theory, but now integrating with utility maximization theory and looking at an equilibrium of an entire economy, at least consisting of two countries. When we talk about tariffs and taxes and so on, these dashed lines will not be of identical slope. They're going to have different slopes in different countries, depending on which good is subject to a tariff. Imports from China coming into the US, for example, are subject to a tariff. And we'll get there, but we're getting very close now.

So there is more to this part of the lecture on production, but I preplanned to stop here to say, next time, when we start, we'll finish up this production lecture on Leontief and Google, Facebook page rank, I mean, and the great Japanese earthquake, and then we'll-- which are fun applications for production. And then we're going to go to lecture 5 and talk about the consumer theory and then combine them all over again with some techniques on linear programming. OK, see you soon.