[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: So take a look at the reading list. And you will see there's various things that are starred on there that are on today's lecture two. There's readings by Kreps and by Nicholson and Snyder. The lecture does not overlap entirely, but it is close in many respects. So I urge you to take a look at. There's one Kreps and two Nicholson and Snyder.

I also mentioned last time, but it might have been abstract, the review sheets. This is lecture one from last time. It's a review sheet-- pretty comprehensive, but also quite specific.

So you should be looking at these things before, during, or especially after the lecture, even if you think you understand the material. Take a hand at this. I certainly have included questions almost verbatim like this on the midterms, so you might as well do it in advance.

And I'm going to read out just one or two of these. Obviously, we don't have time to do too many of them. I would say I'm almost picking at random, but not quite. So the fact that I pick some and not others does not mean that some are more important than others, but only I want to practice with you.

So for example, this one-- Angrist distinguishes randomized trials from quasi-experiments or natural experiments. Please explain what he means by this distinction and give one example from each of the two groups, meaning from quasi-randomized controlled trials and one from quasi-experiments. I know you're new to this, so I hesitate to call on you. So it would be wonderful if someone were to volunteer to answer this question.

STUDENT: I think I could make a go at it.

ROBERT TOWNSEND: All right, please go ahead.

STUDENT: Oh, yeah, so an actual experiment is randomized controlled trials where you actually randomly take a group of people, randomly select some to be a control group, some to be the experimental group. And then you apply some change to one of them, to the experimental group, and track the outcomes. A natural quasi-experiment would be more if something happens-- for example, one state raises the minimum wage and another state doesn't. And you can track a difference.

So it wasn't an actual experiment, but you can still-- but it sort of behaves a bit like an experiment. So you can try to tease apart the causality there. And then I can't remember the rest of the question. Sorry.

ROBERT TOWNSEND: No, it's fine. Do you remember an example from any of those?

STUDENT: From a natural experiment, that would be like the Moving to Opportunity. I remember that being mentioned, where people were effectively randomly chosen to be given housing opportunities to move. And then that was tracked.

ROBERT TOWNSEND: Excellent. OK. Yeah, and the other RCT we did was the paying bicycle messengers temporarily higher wages. RCTs are where we are with these vaccines, just to have an obvious analog to the widely used RCTs in medical trials. But we like to frame questions that way.

Thank you, Charles. And I also went and read some of the things you recommended last time. Thank you. Let's see, let me pick another one.

So Varian speaks of big data as consisting of machine learning and data science. What do those techniques accomplish and what do they fail to accomplish? And again, let me ask for volunteers.

STUDENT: So I remember you talking about what they fail to accomplish, where it's difficult to actually construct the explanation behind the phenomena that you observe. Instead, it just gives you good predictive power where you can construct accurate models, but not necessarily explain them.

ROBERT TOWNSEND: Excellent. That's exactly right. The example was just because we see more police in the community, that doesn't mean it causes crime. And I said something about the example's a bit more problematic these days than it was when Varian used that as an example. But it's still the case that we need to out the causality whether it runs from crime to police or police to crime.

All right, so there are other important questions in here, but I'm a little leery about taking too, too much of the lecture time. So these are called review sheets. And I think you'll find it is not hard. It really is just a matter of keeping up with the lecture and thinking about this afterwards, along with those starred readings.

So today, this is about consumer choice. Remember, we're studying economies, but we need to build and analyze the building blocks, one of which is consumer behavior. So that's what this lecture is about.

The whole economy consists of consumers with preferences, and endowments, and consumption sets, and firms with technology. We're going to focus on consumer choice over possible bundles they might want to buy or eat. We're going to describe some assumptions that characterize what's known as rational, as opposed to irrational consumer choice.

And then we're going to go to this application, which is maximizing utility for a given budget. Let me just say now, because I may not remember to do it later, when we get to this application of maximizing utility subject to a budget, it's what we call "partial equilibrium." We take, as Matzkin would say, some things as exogenous to the consumer choice-- namely, their income and the prices that they face.

Those are like data or givens. And then we're looking at what they choose to buy or to eat. It's partial equilibrium in the sense that we're not trying to explain where the prices come from or where the income comes from. But when we put this building block together with the other ones, with other households and/or with firms, then we'll have what's called "general equilibrium," which is the endogenous determination of prices, and so on.

OK, so let's call capital X here a consumer's consumption set. Two elements of the set called x and y are potential consumption bundles. You should be aware and a bit leery that the notation changes from time to time.

Sometimes in the slides that follow, we're about to say x is one bundle and y is another one. Many times we'll go more abstract, and talk about a vector x of dimension n where they're all x's, x1 through x n. And indeed, when we start to talk about more than one consumer, like consumer i-- that's his or her name-- we can talk about the consumption set x sub i.

The bundle x i is still a vector, but it has a little i subscript in front of the 1 to n commodity labels. But there's no point in being overly detailed when we don't need it, so we'll drop the i mostly today. I don't think it appears anywhere. And as I said for the graphs, we're going to be using x and y quite a bit.

Now, one warning-- should have something flashing on the screen-- stop. Pause-- which is there's something mind numbing about the notation. It's useful. We will build up all the tools with it. But we'll lose sight of the applications.

So I will bend over backwards to pop in a slide here or there and say some words about how this can be used. But the point, although abstract now, is that the artillery where you're going to create applies to wheat and rice, but it also applies to locations. It applies to dates and time. It applies to shocks and states of the world.

So once we do it one way with this powerful but boring notation, we can use it in many actual applications and villages. With the exception that here, we're actually spelling out the consumption set. So you can see that. Can you see it when I'm moving the cursor around?

All right, I'll have to work on my tool kit. So there's two goods here-- bread and leisure. And the consumption set is all of this. It has an upper bound.

Why? Because leisure is measured in hours, and there's no more than 24 hours, a day. So that's one example of a consumption set.

Here's another-- there's two goods, x1 and x2. x1 can take on any value-- integers, rational, or irrational numbers from 0 to infinity. But x2 takes on discrete values-- integers-- 1, 2, 3. So all these lines, including the 0 line, represent elements in the consumption set.

Here's an example with some limited time and geography. You can consume a piece of bread, or bread in general, in Washington at noon or in New York at noon. But you can't be in two places at the same time. So in this case, you have the x and y-axis being the consumption set.

And here's a subsistence. If you don't eat enough-- and we'll talk about starvation in medieval villages soon-- you die. So you have to have either four slices of white bread, or four slices of brown bread, or something in between. But you can't go less than that. Those things are not in the consumption set.

So that's the set in which commodity bundles lie for the house. So what about their preferences? We're going to begin with this primitive, which looks like a weak inequality, but it's actually script or squiggly, to denote that it's a preference relation, which is an ordering.

So if you give me two elements of the consumption set x and y, this ordering will tell me whether x is weakly preferred or at least as good as y in the consumer's ordering. From that primitive, we can create strict preference for any two goods, x and y. If x is strictly preferred to y, that means that x is at least as good as y, but y is not at least as good as x.

Indifference can also be constructed from this primitive. Namely, if x is indifferent to y, then x is at least as good as y, and y is at least as good as x. This preference relation then represents the preferences of each individual-- in this case, as I said, we've deleted individual i. This would be true for all of them. Preference relations obviously can be different across different households.

What do we mean by "rational?" Rational choice means that those preferences are complete, transitive, and continuous. And let me just go over them briefly.

"Completeness" means that any two bundles x and y are ranked relative to each other. So given x and y, either x is strictly preferred to y, y is strictly preferred to x, or the household is indifferent between the two. And the point here is that any two bundles x and y can be ranked and compared to each other. It's a complete ordering.

Transitive-- if x is at least as good as y, y is at least as good as z, then it must also be true that x is at least as good as z. So we will use this later. It's like an internal consistency criterion. You may only see choices between x and y and y and z, but if the consumer is rational, you can infer something about the way they would behave when they're faced with a choice between x and z that we actually didn't see them make.

And finally, we have continuity, which intuitively means if we have x strictly preferred to y, and we have some other bundle x prime which is close to x, since x and x prime are close to each other and x is preferred to y, then x prime is also going to be preferred to y. In other words, the preference ordering isn't jumping around hugely as we move around the commodity space.

So we start with this underlying preference ordering, but we can do a representation as a utility function just by preserving the rank ordering. So for example, if x is at least as good as y, then we come up with the utility function u, which maps the entire consumption set into real numbers in such a way that we give at least as high of a number to x as we assign under u to y.

One thing to note is that many utility functions would work. If we start with this primitive, we could have 2 times u or some multiplicative function of u. It's going to preserve the ranked utility numbers across the bundles. So this is referred to as an ordinal rather than a cardinal utility function.

In the history of economics, utility functions were invented by Jeremy Bentham in England. Jeremy left in his will that after he died, he wanted to be embalmed and preserved, and that once a year he should have tea with his friends. So Jeremy still-- or at least when I was last there-- is sitting in kind of a glass container in the lobby of one of the buildings at University of College of London.

There was some vandalism. Someone removed his head at one point. But anyway, I always think of Jeremy when I think of utility functions. And forgive my morbid sense of humor.

Here's a theorem-- any preference relation that can be represented by a continuous utility function is rational. OK, that's a lot to take in. Any preference relation that can be represented as a continuous utility function is rational.

Well, "completeness" means we're assigning utility numbers to any bundle and any pair of bundles. So pick two, look at their utility numbers, and we have the ranking implied by the utility. Some utility numbers are higher than others.

So we can rank everything. That's why when we have a utility function, the ranking is complete. The ranking is transitive because if the utility number of a is greater than or equal to the utility number of b, and b's utility number is higher than c, then we have transitivity under the ordinary weak inequality.

So the one implies the other, essentially. In this case, we're going from the existence of-- oh, and we're given continuous, which was the third property of a rational preference ordering. So that's embedded in the utility. So obviously, two bundles which are close to each other are going to have very close utility values.

There's something very deceptive about this lecture because no slide is particularly hard, I think. Of course you should alert me if otherwise. But yet somehow, when we go through all of them, you'll be empowered with the tools of consumer theory. So it is definitely worth doing this well.

Local non-satiation is another property. It says that if you're given an x in the consumption set and an arbitrary small number epsilon, there exists some other element y in the consumption set so that when we're within an epsilon neighborhood of x, we find this y. And y is preferred to x. In other words, it's a little weaker than global, which is on the next slide.

Local just means you start with x, draw a circle around it in Euclidean space-- two-dimensional, or a sphere, or in higher dimensional objects. And you can always find something in that area, however small, that is strictly preferred to the starting point x. Eventually, we'll get to competitive equilibria or entire economies.

Are competitive equilibria optimal? The answer may be no if we lose this local non-satiation. So it will sneak back up on us. In other words, there's no bliss point. This is the dismal view of economics, I suppose, that more is always preferred to less, and people just can't get enough.

Here's a more global view of it, which is with two goods. We start with this bundle here. This whole hatched-in area has more of one or the other or both goods. And it's the weakly or strictly preferred set relative to that.

In fact, if they value both goods, then even if you keep one constant, you're increasing the other. So it's a strictly preferred set. This, which has less of one or the other of both relative to the starting point, is the worse-than set.

So back to rank ordering, we start here and say, what about a bundle here? Well, that's strictly preferred. Start here, what about going southwest? Then the original bundle is strictly preferred.

But we leave some question marks in these off-diagonal elements, essentially. Because it has more of one good and less than another. So a priori we're not able to say what the consumer would choose without providing more details.

But we can provide those details if we go back to, say, the representation of the utility function. So we pick a bundle, and it's associated with a given level of utility. Remember Jeremy Bentham. And then we ask what other values in this consumption set would generate the same utility level as the starting point?

And that's the entire set of bundles on this curve. So this is called an "indifference curve," in the sense that these two bundles-- x1, y1 and x2, y2-- under the consumer's preferences are indifferent to each other. It might have been better to call it an "isoutility curve," because that language is even closer to what it is.

And later in the class, when we get to production functions, we can talk about isoquants. And later we can talk about isocosts. And this is isoutility-- same utility. Or traditionally, it's called the "indifference curve." Now, this is drawn with a certain diminishing slope, and I'll come back to that momentarily.

Well, of course, there's more than one curve. Because under a utility function, it's complete. So everything has a ranking.

Let's jump up here-- utility levels are higher. Jump up here, higher still. So as you move northeast, we're moving in the direction of strictly increasing utility.

In fact, stuff is going on between these curves, too. It's just utility levels that are between U1 and U2. So you typically graphically just draw one or a finite number. We don't try to fill the whole thing out because then it would just go black.

Let's try this diagram. This is indifference curves that are crossing. Note that A is northeast of B, has more of both goods X and Y. So A must be strictly preferred to B.

However, A is on the same indifference curve as E, the indifference curve labeled U1. And E is on the same indifference curve as B, the indifference curve U2. Under transitivity, if the consumer is indifferent between A and E and indifferent between E and B, then by transitivity, the consumer should be indifferent between A and B.

But we started out by saying that the consumer strictly prefers A to B. So we end up with a contradiction. Hence, indifference curves can never cross.

So let's come back to the slope. Let's define the marginal rate of substitution, the willingness to give up one good for another, as the negative slope of the indifference curve. So that is, the marginal rate of substitution, is how much of Y you have to get when you give up 1 unit of X while keeping utility constant at U1.

So that curve and this number tell us something about trade-offs that a household would willingly do to be indifferent, or potentially, later, to even gain. So if you go back to this slide, this is a diminishing marginal rate of substitution. The slope is getting less and less.

It's obviously negative because when you get more of one good, you're willing to give up the other. The issue is how much. If you start off with a modest amount of X1 and a relatively large amount of Y1, and you move down this curve, the marginal rate of substitution would be this slope.

But if you gave the household more of X and less of Y, the slope becomes less and less because here, as you move pointwise down, you have more and more X. But X is less and less valued because you have more and more of it. You have less and less Y, so Y is more and more valued as you give up more of it.

And so to make this movement [INAUDIBLE], you require more and more X as you surrender a given unit of Y in order to maintain the indifference on the indifference curve. I mean, down here, it's almost flat. You just have so much of good X that variation of it doesn't change utility very much. So you don't have to change Y. And here, it's almost vertical.

So we normally assume the marginal rate of substitution is decreasing, just as in that diagram. And the intuition is they want things to be balanced. And then we're about to get to notions of convexity and concavity, but I'll wait for that slide.

One thing I want to say right away is that the marginal rate of substitution doesn't have to be decreasing. If indifference curves are linear, as they are here, then the slope is the same. It's a line, after all. So the marginal rate of substitution is not diminishing, it's constant.

And here, we have our classic case, which is the way they're typically drawn. Here, it's the opposite, which is if you start at this point of the L and move in the direction of having more of good X you would gain no extra utility. Likewise, if you started at this point and had more Y, you would gain no extra utility.

The only way to achieve higher levels of utility going from U0 to U1, et cetera, is to increase both, and in the exact same proportions. So these are sometimes called Leontiefian indifference curves, although he really created them in production, and we'll get to that in two lectures. So the degree of substitutability-- these are very curved, these are linear, these are kinked-- is something called the "elasticity of substitution." So these are constant elasticity of substitution curves that are parameterized by the degree of curvature.

And here are some particular functions. This one's called Cobb-Douglas. He was an Economist at Illinois, or Douglas was. Senator Douglas-- I think he was a Senator from Illinois when it was an agricultural state. I'd have to check that to be sure.

So x to the power alpha, y to the power of beta-- perfect complements, which is this one is here, you can see the ratio I was referring to. To maintain utility constant, the ratio of y to x has to be alpha over beta. Perfect substitutes, linear, constant elasticity of substitution, this function. And we will use these for numerical examples, and p sets, and so on, from time to time.

OK, now we get to convexity, and quasi-convexity, and all of that stuff. So a set X, like a consumption set, is convex if whenever you have two separate bundles x and x prime in the set, and you take a linear combination of those two points with weights alpha-- say alpha between 0 and 1-- then this new point is also in the set. So it's a standard definition of a convex set.

A function is said to be convex if whenever you have two points x prime and y prime, and x and x prime and you evaluate the function f, then the weighted average of those two points over here evaluated under f is less than the weighted average in the range of the function. So that's a definition of a convex function.

The set of points on or above the graph of a convex function is convex. So that's kind of linking these two definitions together. I'll show you the picture momentarily.

A convex function and negative of a convex function is concave. To define quasi-convex-- a function f is not simply convex but weaker quasi-convex if whenever you evaluate the two end points, and x, and x prime, and take the max of the two, it's no less than the function evaluated at the weighted average. And likewise, the negative of a quasi-convex function is said to be quasi-concave.

So what are we going to use here? I'll show you a picture where we'll display convexity of a set, convexity of a function, and the use of quasi-concavity. So this is the one that is consistent with the definitions. This is the one that is not.

So this indifference curve is a curve in the space of x and y. And it's convex, actually strictly convex. The set of points with utility greater than or equal to the utility associated with this indifference curve is this hatched area. And this is a convex set. You pick any two points-- for example, this one and this one-- and take a weighted average, and you'll still have a utility level which is greater than the utility is associated with U1.

If you looked at the definition of quasi-concave utility function, we take the utility function which is generating these indifference curves, and say, what's the utility at this point under that utility function? What's the utility at this point under that utility function? Take the max of those two, and then show that the weighted average on the line gives a strictly higher utility level.

So quasi-concave function reverses the inequality here. So the max of the utility at each of the endpoints is less than the utility at the weighted average. So that is tedious, and it's a bit hard to absorb, especially if you're not used to using all of these things.

Anyway, this is the good-looking picture, and this one is ugly in the sense that it violates every single-- this property, this utility indifference curve is not convex. It has some concave portion. The set of bundles on or above the indifference curve, when you take weighted averages of them, you would be below the indifference curve, violating the property. And similarly, if you take this point and this point, take the max of the two, and take a weighted average, the utility is actually now less than, not greater.

That quasi-convexity or concavity of utility is a bit surprising. It has to do with the ordinal nature of the utility function. We don't want to say a function is necessarily strictly concave, although that would imply quasi-concavity.

But we could take another transformation that removes the strict concavity, and it could still be quasi-concavity. It's clear by plotting in the choice space in this way, that quasi-concavity is what we want, although we haven't really talked about choices. We're just talking about rank ordering.

So marginal rate of substitution defined a bit more mathematically-- we have a utility function over x1 and x2. We can call the marginal utility the derivative of the utility with respect to x1 or utility of x1, likewise for x2. Now, if we're moving along an indifference curve, we're changing the quantities of x and y, but we're keeping utility constant.

So dU, or the total differentiable of utility, does not change as we move along an indifference curve. So dU here would be the derivative of utility with respect to X1 times dX1 plus the utility with respect to marginal utility with respect to X2 times dX2. That's this expression.

This is going to be 0 by definition of moving along an indifference curve. Since it's equal to 0, we can solve for dX2. I think you can visualize what happens here.

dX2 divided by dX1 is equal to the ratio of these marginal utilities. Rather than derive this every time, although it's a bit dangerous to memorize things, basically in that diagram, we have x1 on the x-axis, x2 on the y-axis. And the marginal rate of substitution is the negative of the marginal utility on the x-axis and marginal utility on the y-axis.

Another definition-- homothetic preferences-- so preferences are homothetic means that the marginal rate of substitution depends only on the ratio of the goods and not on the total quantity. So a picture is worth a thousand words. Here's an indifference curve, here's a point.

What is the marginal rate of substitution? The slope of that point, of the curve at that point. Here's the ratio of the goods, X2 to X1 on this line. As you move on this line, that ratio stays constant.

We go up here to this higher level of indifference, the slope is the same. So essentially, this indifference curve is a blowup of this one. And you could draw one out here, which is a radial blowup of these two. So you can scale up and down, and just keep the slopes the same as you move, for example, along either one of these two lines.

So who cares? Well it's going to turn out that a household with these kinds of indifference curves is going to respond in a certain way to income, respond in a certain linear way, that's going to allow us to, for some purposes, aggregate up over many diverse households, and act like there's only one representative Robinson Crusoe macro household in the economy. You probably don't remember, but when I went through the great, big, long syllabus last time, the next to last page talked about aggregation. And the lecture on Tuesday will be an example that violates this kind of homotheticity.

Here's a violation. We have a utility function which is linear in one good, x1, which we'll call money. And the log of x2 is not linear, but it is concave.

So we call this the money good because as you inject or take away good 1, utility goes up or down linearly. And it is also true that if everyone in the economy has linear utility functions in x1, then the amount of utility goes down when you tax somebody is exactly the same amount that utility goes up for somebody else who got the subsidy. So again, that's a useful property to have, but it's very special.

Anyway, in this case, the marginal rate of substitution, which is not constant on any rate from the origin, the marginal rate of substitution is constant as we fix x2 and vary x1, so going along horizontally. And you can convince yourself that has to be true. Because if you took the definition, which is the marginal utility of x1 divided by the marginal utility of x2, the marginal x1 is constant because it's linear. And this thing is varying with x2. So if you fix x2 and vary x1, you're not going to change the marginal rate of substitution.

So we have now successfully reviewed basic properties of consumption sets, consumer preferences, and so on. You might want to squirrel away some of those properties because we'll come back to them later throughout the course, in the subsequent lectures. But here they are, essentially all listed now. It doesn't mean they all have to be true. Some of them contradict each other.

Theorems are if we have local non-satiation then a competitive equilibrium is optimal. But we may not have it, in which case competitive markets don't work, and things like that. So these are not requirements. These are options.

All right, but now we get to an application, this partial equilibrium problem, which is to maximize utility subject to a constraint for the first time. So this is the budget constraint. Income on the right-hand side-- I for income-- is taken as given. The prices are of each good.

There are n of them, p1 for good 1, p2 for good 2, and so on. So the sum of the expenditures at prices p over choices x cannot exceed the income I. So this is your budget.

We all have budgets. I sometimes wonder about the US Treasury, its budget deficit, debt outstanding, is now larger than any time since World War II. But households have budgets. And I'm worried about the government, but we'll get to that later.

So suppose you were to double all the prices. We start with this, and we double income as well. So we pre-multiply every term by t, little t, both on the left and on the right.

It's a constraint in the problem, but it doesn't change the constraint by multiplying through by t. Or if you started with t, in there, just divide everything by t and you take it out. So it can't change the math of anything. So that's one way to look at it.

I'll show you in the picture in a minute that no price illusion-- if prices double as in inflation, but your nominal income also doubles, no harm, no foul, at least not for this household. Now, maybe consumers aren't rational and maybe they're sensitive to income movement and not price or vice versa. But for the most part, we're going to assume that these are equivalent constraints from the point of view of the household.

And this is a definition of homogeneity of degree 0 of individual demand. The demand can't change because the constraint doesn't change. And here's the picture.

So this is the set of points which completely exhaust income when there are two goods, x and y, P x and P y. You could move over to this extreme point, take all of your income, and buy good y. How many units can you get?

Well, it's $10 divided by 5, or 2. So we have a maximum amount that would be possible to purchase of good y, and maximum amount you could purchase of good x. We have exhaustion of the budget on this line. Actually in the interior, we're spending less than income. Consumer might not want to end up in the interior, but that has to do with satiation.

What is the slope of this budget line? Well, you can derive this the same way we kept utility constant-- keep income constant, take dI, set it to 0, and then take the derivative of P X times dX plus P Y times dY, and you get dYdX is equal to P X over P Y. So the ratio of this line is exactly P X over P Y.

It is kind of interesting that only the ratio matters. That is to say, setting aside what's going to happen to income, if you double all the prices of both goods, you don't change the slope of this line. In fact, more than that, we could pick one of these goods-- say Y as the numeraire-- set P Y equal to 1, and then we'd only be fussing about the price of good X.

Here it's more general. Everything's in dollars or some unit of account, depending on what country you're in. Could be measured in wheat, goods. Price is measured in terms of how many bushels of wheat you can buy.

So it's arbitrary. You have to specify what it is, but only the ratio matters. That thing about inflation-- it's like doubling prices and doubling income doesn't change this line.

And where does the consumer choose to be? Here, C for choice, maybe. So why? Well, one way to view it-- the household wants maximal utility.

So in the set of choices inside or on the budget line, what point has the highest utility level? It's point C. If you go this way, it's less. If you go this way, it's less.

There's another somewhat interesting way to view this marginal analysis, which is if you start it up here, the slope of this indifference curve U1 is pretty steep. It's negative, but I mean steep in absolute value. And the slope of the budget line is less.

That means this household's willing to give up a substantial amount of good Y to get one more unit of good X as he or she moves away from B. But in fact, given the prices in the market, you don't have to give up that much to gain value in order to be able to purchase X. You give up less Y.

So this is what you're willing to give up and be indifferent, but you are allowed to give up less on the budget line. So you're definitely moving in the direction here of higher and higher utility. And you can conduct the same reverse experiment down here.

So next, let's do a utility maximization problem. And as I promised last time, I'm going to introduce a cool tool. And it's called the Lagrangian.

So if we want to maximize utility subject to a budget, set up the problem in this way-- Lagrangian, which is the French guy that figured all this out, Lagrange, is to write down the utility function plus a shadow price lambda times a repetition of the constraint. So we turn the direct maximization problem, which was this, into a pseudo problem where the constraint is actually in the objective function.

Anyway, if you start taking derivatives of this looking for the first-order conditions for a maximum, the derivative of the entire expression with respect to x k is the way that k will enter and generate marginal utility under U. That's this term.

Plus, we'll pick up lambda. And where else does little x enter? Over here somewhere, P k. It's linear in x k coefficient P k. So we get this thing.

And again, we're looking for an interior solution. We set this thing equal to 0 as a necessary condition for an optimum. Lambda is actually a control variable, too, the way the x's are, just looks a little weird.

We take the derivative of the Lagrangian with respect to lambda, and obviously, we're going to get the constraint right back again. And we're going to set it to 0, which we assumed before that they wanted to spend all their income. We were implicitly assuming before that the marginal utility of income was positive, or in this case, that lambda is positive. So that's why we get the constraint right back.

I'm not sure if you've seen Lagrangians before, but it's just a version of multivariate calculus where you find the first-order conditions to a maximum problem. Note that we can take any two pair of bundles k and j, i and j, or whatever. When you take them pairwise, we're going to have the derivative with respect to good I here and the price of good I here and the lambda.

That lambda is the same when we go J, J. So we can divide the left by the right, and the lambda will disappear. So we're going to get these ratios of marginal utilities, which are, of course, the marginal rate of substitution equal to the price ratio. And that's what we were saying-- maybe I didn't say it-- when we have this tangency over here, the marginal rate of substitution in utility is the marginal rate of substitution in the budget, which is the price ratio. So we're getting that result back now as a consequence of this optimization.

Another danger-- I feel like I ought to have this slide that says, stop. Pause. Do not memorize these diagrams. Just like the notation becomes mind numbing, you get used to drawing a picture in a certain way. And then when someone nasty puts another specification on the table, you don't know what to do with it.

So here I am being nasty. I'll show you n incorporated it into the Lagrangian. So this is a corner solution. Where is the best utility level?

The highest level is here because as you move down this budget line on the line, you get higher and higher levels of utility. But if the consumption set is all of this non-negative [INAUDIBLE] you can't go negative in y. So you stop. This is a corner solution.

By the way, you would get this as well if the utility functions were linear. They're almost linear here, and you have a corner solution. Of course, they could be linear and coincident with the budget, in which case everything on the budget line is an optimum.

Or those linear indifference curves could be really flat, in which case you'll go over here to a different corner. It doesn't have to be linear to generate corners. This is an example. But we need to incorporate it into our maximum techniques.

So what's the general recipe? Before we translate it into the consumer optimization problem, we'll get the technique down. As I promised, you're being gifted with a powerful tool.

So we have a function f over a range x. We want to maximize it. And we have various constraints. We have n of them.

g i of x-- g is a function. There are i functions, each of which, when evaluated at x, cannot be negative. Properties f and g i are real-valued, continuously differentiable. In other words, they're very smooth, and not only continuous, the derivatives are continuous.

How do we solve this problem? Look it up in the cookbook, or memorize it, even better. There's a trick here I'll show you.

The Lagrangian put the objective function, just like we had U. And now write down each and every one of these constraints with these Lagrange multipliers lambda i for constraint i, add them all up. What's the little trick? The g i is non-negative.

We enter it with a plus sign, and then things are going to come out right with the math. It's not necessary. Things work even if you write down the g's as less than or equal to 0, but then you have to struggle a little bit more to make sure you've got the math right, and you didn't reverse the sign.

So anyway, try to do it this way-- g i greater than 0 lambda g i. Write out the first-order conditions. That's just a derivative.

So we're taking the derivative of the Lagrangian with respect to x k is the derivative of the function f with respect to k plus the way in which x k is entering in each of these constraints g i, each pre-multiplied by lambda i. Set it equal to 0, and do that for every argument x k. With me?

So next, write down each constraint again. Write down these multipliers which are non-negative. And this is cool-- this is called a complementary slackness condition, which is the multiplication-- oh, it should be a lambda i on that. That's a typo. Lambda i times g i equal to 0, and then we're going to look for a solution to all these equations.

Now, if lambda i were 0 then it doesn't matter what g is evaluated at x. It's still 0. Well, if lambda is 0, then effectively, it's not entering into the problem, and you can ignore it, and drop it by setting it equal to 0.

If lambda i were strictly positive, though, then the only way you can get this product to be 0 is to have this g i be 0. And that's here.

And this is what happened with the income maximization problem. I asserted-- if we go back now and look-- that the marginal utility of income lambda on the budget constraint was positive. And if it's positive, then we have to set the budget constraint equal to 0, income equal to the sum of expenditures, which is the way we solved that simpler problem to begin with. So what we did earlier generalizes to this more complex.

Now, why are we doing it? Well, let's come back to this problem. Are they necessary? They could also be sufficient. If that function f is quasi-concave, and the constraint sets, these guys, are convex, then finding a solution to the first-order condition is equivalent with finding a global maximizer-- very powerful, and we will use it subsequently, again, to address the issue of whether competitive equilibria are optimal.

So going back to the motivation, which was this picture-- we kind of know that we're going to hit a corner. What comes up with the-- if these were linear indifference curves, and you had this linear budget set, then when you take first-order conditions you're going to be in trouble because everything in the first-order condition is going to be parameters, not control variables that you want to find. So that's a version of being nasty. But like I said, we're working it all out.

So how do we deal with this hitting the corner? What does the Lagrangian find? So we maximize utility subject to the budget, subject to non-negativity constraints. As in that picture, consumption cannot go negative.

We form the Lagrangian mechanically, which says write down the function, write down each of the constraints, lambda on the budget constraint, mu-- I changed notation-- mu k for constraint, this one, the non-negativity constraint on x k, take the first-order conditions. And now we get something similar, except that you pick up this shadow price or Lagrange multiplier on the non-negativity constraint of good k. Whereas before, that mu wasn't there.

So then applying the recipe, we write down the constraints again. We write down the non-negativity of the Lagrange multipliers. We have this complementary slackness condition, which is the product of the Lagrange multipliers with the budgets, and we look for a solution.

If we were at a corner where good 1-- or good x, whatever was in that figure-- were 0, then if we go back here, this would be positive. This mu k would be positive, and we take it out. This would then be less than 0.

Something less than 0 plus something positive makes it 0. We leave the positive thing out, and we get the thing that's less than 0 looks like this. And what this says, essentially, is-- there's a number of ways to read it.

Let me do it this way. Put lambda over here. Lambda is the marginal utility of income. So how much income does it take to get good k? And what are the utility consequences of that?

p k is the number of dollars you have to give up to get good k. The utility of those dollars is lambda. So lambda times p k is the marginal cost of getting a unit of good k.

And what's the marginal benefit? The marginal utility is the benefit. In this case, the marginal cost exceeds the marginal benefit, so you don't do it. In fact, the solution, obviously, is to set x k equal to 0 because at any positive level, the marginal costs are higher than the marginal benefits. The best you can do is to set it to 0.

So Cobb-Douglass utility-- this was an example we had before. Good X and Y to power alpha and 1 minus alpha subject to a budget-- you can solve this. It's really not necessary to go through in class. It's a good example.

You maximize. Using the Lagrangian, you get the first-order conditions. Do some algebra with that first-order condition comparing ratios. This would be the marginal rate of substitution. This is going to be the price ratio.

They have to be equal for an interior solution. And the solution has to be interior because the marginal utility is going to infinity when one or the other of these goods goes to 0. So you never want to go down that far.

We can rearrange these terms. And the bottom line is, p X times X is the expenditure on X, and it's equal to alpha times if I. If you go back to the utility function, alpha is the power on X.

So whenever you see this Cobb-Douglas, you can immediately jump to something you should know, which is the expenditure shares have to do with these powers. The expenditure shares on X would be alpha. The expenditure share on Y would be 1 minus alpha.

We want to see that in notation. p X times X is the expenditure share on X at the optimum. It's alpha times I. Likewise, p Y times Y is 1 minus alpha times I.

So you have your whole budget. You spend the whole budget. And the share of the budget you're spending is dictated by the alpha, and 1 minus alpha, and the prices. So that's constant expenditure shares.

And again, when you go to data, you go take a look. And we'll do that in the lecture on next Tuesday. Now again, we got through all that stuff-- hopefully useful, maybe even interesting to you.

And now you've forgotten about geography. You've forgotten about time. You've forgotten about states of the world. So let me end with two slides-- the lecture about consumer choice, when the consumer decides how much leisure to provide or negatively, how much to work and how much consumption to buy.

So it's going to be one thing to say right away-- leisure is going to be a good thing. Consumption is a good thing. So those axes that were labeled x and y are now labeled c and l, where c is consumption and l is leisure.

So it looks standard now, but it's a much more interesting problem. It's how much to work problem. It's like the bicycle messengers.

The other thing is maximize utility subject to the budget. So this is kind of a twist because instead of having income over here, we've got the wage, w, per unit time times the time endowment that the household is endowed with. So this is as if you took your entire time endowment, worked all of it, and then went back and bought some leisure and bought consumption.

Probably a more obvious way to put this is to take wl, subtract it from both sides. Then you have T minus l. So time minus leisure is the amount spent working. You surrendered it to the market. You don't enjoy it, but you do get the wage for it. And then this becomes income, and income is spent on consumption.

The other thing about this problem is like the choice of a numeraire, everything is measured in terms of consumption. There's no price. You could have done it. You could have price of consumption times consumption, price of leisure, the wage times the leisure.

But we just normalized wages. So this is what's called the real wage. It's how much consumption bundle you can get for a given unit of labor.

And here's the picture. So you're effectively starting here, with 16 hours a day of leisure. I don't know, some sleep is allowed, eight hour bedtime. The remaining is 16 hours.

But you may want to earn some money. So if you move in this direction, you're giving up a unit of leisure, working that unit, and getting the wage for it, which allows you to buy consumption. So the slope of this line is the wage.

In this case, the wage rate is 10. So every unit of work that you supply gives you 10 units worth of consumption. And you could work your entire 16-hour day at a wage of 10, and have 160 units of consumption. But this guy is stopping short of that, where the indifference curve is tangent to the budget line. So this is the choice of how hard to work.

And again, I could have put in a dynamic choice problem. I could have put time-- consumption today versus consumption tomorrow. We're going to have one of those at the end of the lecture next time. As a constant reminder, we could have had choice about locations. And we'll do some of that pretty shortly, as well.

All right, so that's the prepared slides. And we have one or two minutes left for questions. All right, so you have all the slides. Take a look at them, study them.

You also have the review sheet questions. Go back and take a look and make sure you're able to answer them. We'll have a little review next time of what we did today.

And then we're going to go on to explore these expenditure shares, and income expansion paths, and whether poor people consume luxury goods or the opposite, et cetera. So that's for next time. I've already mentioned last time something drawing on a real application from the material we're going to do in the lecture next Tuesday on China, and kind of varying these price ratios of these budget lines by subsidizing consumer purchases.

It's a wonderful problem set, worked out by one of your former classmates, actually, who I hired as a UROP. So you're almost in the position to start to take the look at that, especially after Friday, and certainly after Tuesday. So that's all I have for today.