

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT M. TOWNSEND: Let me remind you from the reading list where we are. We did the introductory lecture, of course, about economic science. And last time, we did our first nuts and bolts lecture on preferences, consumer theory, with the associated starred readings. And today, we're going to do more the way we ended the last class on consumer behavior.

I'm showing this reading list because I just want to remind you that these starred articles do exist. The Nicholson and Schneider chapter 5, from which this is drawn, although not entirely. And also, this Jensen and Miller paper, which I will say a few words about today because it fits into this lecture, but also, as I probably mentioned too many times, that's also a big part of your first problem set and involves getting that paper, reading the yellow markings as highlights, accessing the data that these guys used, and running through some of their calculations. So that's the part of the starred readings for this lecture.

So in terms of the study guide, Consumer Theory, as I was alerting you last time, there's a pretty good thorough review of what you should have learned or are learning from the topic of last time. And I will not go through all of this, but I will ask a few questions, get a few volunteers, maybe even call on a couple of you. So this one, give a plausible example of a commodity point that could not be in a typical consumption set.

And I'll take volunteers on that. It's partly remembering the slide where I showed examples of consumption sets. Does anyone remember? The question is-- yes?

AUDIENCE: One infeasible consumption site is consuming 25 hours in a day or something like that.

ROBERT M. TOWNSEND: That's good. Yep, that was one of the diagrams where the amount of leisure available for work or leisure, the sum could not exceed 25. There was another one, which you might not have thought of given the way this is worded, which is consumption of bread in New York at noon and consumption in Boston at noon of bread.

And you're not allowed to have bread in the interior because you can't consume simultaneously in both New York and Boston. So that was another example. But your answer was great.

So can someone volunteer to give me the definition of local non-satiation?

AUDIENCE: That means that there's in that-- so if you have a bundle, in any neighborhood of that bundle, there is some other bundle that's preferred to the bundle you originally had. There are no bliss points.

ROBERT M. TOWNSEND: That's right. That's an excellent answer. Yeah, I can't remember whether we did X or Y.

But for any X in the consumption set, there's some other Y, no matter how close we are to X, which is preferred to X and-- well, and Y has to be in the consumption set. So both things have to be true. It has to be available in this little ball or circle around X, and it has to be preferred.

And you gave a good counter-example. If there were a bliss point, that means there's something that is maximal utility among all possible points. So if we do a circle around that, everything else would be inferior.

If you went beyond the bliss point, northeast of it, you have too much of all the goods. But then, ironically, you also have local non-satiation because, given any point X , there's another point Y , which involves less than both goods. That would be strictly preferred because you're moving to the bliss point.

So anyway, the point of this little discussion is that local non-satiation and global non-satiation are two different things. Thank you. OK, let's-- what does it mean for a preference relation to be represented by a utility function? Given a preference relation, is the utility function that represents it unique? Yes or no, and explain.

AUDIENCE: Yeah, so we say that a utility function represents a preference if, for example, for any consumption bundle, X , any two consumption bundles, X and Y , if X is preferred to Y , then U of X is greater than U of Y . And the utility function is not unique. It just has to respect the given ordering of the preferences. So we could have U , or U squared, or U cubed, and all of those would work.

ROBERT M. Excellent. Perfect answer. Any other comments or questions about that?

TOWNSEND:

AUDIENCE: I had a question. On this slide, it had said that it was unique only up to an order preserving transformation. What exactly does that, I guess, mean?

ROBERT M. Well, if you had three points, ordering them from low to high in terms of preference relation-- say, X_1 , X_2 , X_3 ,
TOWNSEND: where X_3 is preferred to X_2 is preferred to X_1 , all we need to do is to assign a higher number to the best one, a lower number to the intermediate one. So something like 15, 5, 1 preserves the rank ordering. Those are utility number assignments. But we could multiply by 2, and then we have 30, 10, 2. But we're still assigning numbers from high to low to preserve the same rank ordering. Does that help?

AUDIENCE: Yeah, thank you.

ROBERT M. OK. Yes, go on.

TOWNSEND:

AUDIENCE: Do you eventually need to, I guess, reintroduce some concept of magnitude-- or not magnitude, but have something beyond just ordering if you want to start talking about gambling or expected value?

ROBERT M. Yes, absolutely. And we're going to do that in two lectures. So there, the utility function will-- the ordinal
TOWNSEND: representation matters. And the question will be, how much more utility do you get or how much less, depending on an adverse event?

So that the curvature literally represents your willingness to take gambles. Whereas with an ordinal function, a linear function is fine. Linear function when you're facing gambles means you're risk-neutral, and you just rank based on expected values.

That's all a little bit early. It's a great question. And we're going to get there in two lectures.

We got time for one more here. What does it mean for preferences to be homothetic?

AUDIENCE: The MRS is just depending on the ratio of two goods, right? Not depending on the absolute quantity of each good.

ROBERT M. TOWNSEND: Yep, perfect. Wow, I'm very happy, very impressed. Thank you, all. That's tremendous.

OK, so again, these are just somewhat randomly selected. I do go over these before class to help me decide, since I can't ask them all, which one I would ask in class. There are some of these that say draw a graph or so on and so forth.

So I didn't ask about those. But none of them are supposed to be terribly hard. But it is always a good review, at a minimum.

OK, so today is a lecture on income and substitution effects. Or really, it's about consumer-maximizing behavior. And so when we solve the consumers utility maximization problem, we get individual demand functions. It should say function with an S because if there are n goods, little n , then there's a demand for good 1, good 2, good n , and so on.

And what are the arguments in these demand functions? The entire vector of Prices, denoted P_1 through P_n , and Income, denoted as I . And you can think about prices and income denominated in dollars or some unit of account.

We're going to be spending time looking at how the price of good k changes when we change the price of good k . But of necessity, other goods will be changing as well. It's at best, the implicit happenings in the diagrams.

So here is a little bit of a repeat but expanded from last time. We've got two goods denoted, X and Y . And this household has a certain amount of income, I , say, I_2 . If income were I_2 , to this would be the budget line, which has a slope of P_X over P_Y minus. And this is the tangency which is the maximum utility point denoted X_2, Y_2 , associated with income I_2 .

What this diagram is featuring, though, is not a single maximization. It's how the maximization point changes if you increase income. So you would move from X_2, Y_2 to X_3, Y_3 , for example. And likewise, if your income were less, if this household's income were less, they would move to Y_1, X_1 .

So you could connect these dots. And as only income is varying and the price ratio is held fixed, we could call the connected dots an income expansion path. And I'll show you some examples of that momentarily.

Here is a different experiment. Did we miss something? Well, in this one, you'll get lulled into thinking that, as income changes, the amount spent on both goods has to change, increase. But that's not necessarily the case.

So here, again, we're going from I_1 to I_2 to I_3 and tracing out the quantities from ZY_1, ZY_2 to ZY_3 . And here, the quantity demanded of good Y is increasing, but the quantity demanded of good Z is actually decreasing. So there's nothing terribly peculiar about the shape of these indifference curves.

We didn't have to struggle too hard to get this phenomenon to happen. But anyway, this phenomenon has a name. Namely, the good Y over here would be called a normal good, and good Z would be called an inferior good or a necessity because the amount that you consume of it diminishes as income goes up.

So here are those definitions, a good for which the demand of good i goes down as income-- capital I -- goes up is said to be an inferior good. It doesn't have to happen over the entire range of income and prices. But if it happens locally, it's called an inferior good locally. And likewise, if the derivative with respect to income is positive, then we call, in some range, the good to be a normal good or noninferior good.

So there's something called Engel's Law. And with the caveat that this is not one household, it's actually comparing across households that differ in their annual income. So here, we're going from an annual income of \$225 to-- this is written in 1920, by the way.

Don't be too alarmed. There's been a lot of inflation. So this, in real purchasing power, was more than you thought, \$300-- or this range, 450 to 600 or 750 to 1,000.

And what you can see is the percent spent of the budget on food declines as we move to these income groups with higher and higher income. So in some sense, if all consumers were alike, and some had more income than others, we'd be looking at what an individual household would do as its income changed. And in this case, food would be an inferior good. You're saying, what's going up? Clothing, not by much.

Lodging, not at all. What really expands as income goes up is services, education, health, legal, and other services. In fact, in the US economy, looking over decades, the fraction of services produced in GDP just keeps going up, and up, and up, which is reflective of the fact that, as households have more income, they demand more services than less food.

So that's pretty dated, 1920. So I went on the web and looked this thing up, Engel's Law today. Here, we're dividing households into quintiles, five groups, an equal number of households in each group. And we plot the fraction of so-called luxuries and necessities for each group. Now, this is a bit problematic in the sense that necessities are defined as goods or services where consumption is proportionately lower as a person's income increases.

So again, necessity or non-normal goods or as we defined it in the last three slides. Not too surprisingly, the highest income groups, the top 20% of the population, the proportion of expenditure on luxuries is a lot higher than on necessities. And that gap lowers until you get to the lowest group, in which case, the fraction spent on necessities is a lot higher than luxuries.

Now, when I grabbed this, there was an whole article about how poor people are extravagant because they're still spending money on luxury goods. Well, I mean, narrowly speaking, yes, that's true. But again, it's a little bit harsh because it doesn't mean they're driving around in a Mercedes Benz. It just means they're eating some fraction of what they spend on are normal goods.

AUDIENCE: Wait. Can you-- is there a little bit of issue with how those two-- I guess, how the two lines are defined? Because the way you defined them, it seems like you said luxuries are just what rich people buy, and necessities are what poor people buy. So doesn't this kind of get confirmed just by how you define it?

ROBERT M. TOWNSEND: Yes, I think that's exactly right. It's almost tautological. I mean, the way they defined it, it has to be true that this blue line is going up and up as you go across the quantiles. So I agree with that comment. But I still think the graphs-- how much is it going up, what's the difference, and so on, could go up just a teeny little bit, in which case, that blue line would not move much. Other questions?

Well, one takeaway from this is different households that differ in income do not spend the same fraction of their money on goods. So let me take a chance at the risk of-- preview of coming attractions, this is way back in lecture 18, where we're going to talk about aggregation, which has to do with whether we can represent demand functions as if they came from a single maximizing representative household or weaker, that there is a group of households, like the lower quintile, and they're all alike. And we could take a representative low-quintile household and talk about the demand as if it were coming from that one person.

What is a necessary and sufficient assumption for us to be able to do that kind of aggregation? It is that these income-- in this case, denoted wealth expansion paths-- are linear and have exactly the same slope. So these guys will spend-- and I'm not trying to go through the entire lecture. I'm just really trying to motivate you that we looked at some micro data, and later, we'll make some assumptions. It's good to put the microdata about the assumptions on the table at the same time and compare them.

What's going on here is that, if these income expansion paths are parallel, then when we increase income for one household, its expenditure would go up. But because this line is parallel, when we decrease expenditure for another one, its expenditures will go down in exactly the same amount that this guy's went up. So those income effects are completely offsetting. And that turns out to be necessary and sufficient to add up all of income and drive demand as if we've given all of income to one household.

All right. Well, it was with some trepidation that I showed you this picture, but I thought you might like to know what's coming. And we'll review all of that when we get to 18.

OK, so let's talk about income and substitution effects now. Instead of changing income, we want to change a price. We're going to focus on good X. We're going to look at an initial point, which is Y^*, X^* . And then we're going to lower the price of good X and not touch good Y.

One way to quickly convince yourself that the diagram is correct is that this point over here on the Y-axis is taking all your income and devoting it to the purchase of good Y. The price of Y hasn't changed, so this point hasn't changed. But the price of X is going down. And logically, the set of things you can buy is expanding.

So this budget line shifts to the right, but not in a parallel fashion. So what is the new purchase? The new purchase is a tangency, Y^{**} and X^{**} . And we're going to break that movement from the initial point to the subsequent point into two pieces-- one, a price effect, and two, an income effect.

So since the price of good X goes down, P_X over P_Y in absolute value is going down. So there's a new budget line with a lower slope that is pivoting on the initial indifference curve, U_1 . So setting aside income effects, which we'll get to momentarily, the natural impact of lowering the price of good X is to increase the expenditure of good X and decrease the expenditure of good Y.

So that's called the substitution effect. And it's a movement from X^*, Y^* to B. But in fact, of course, income-- effectively, purchasing power has gone up.

Even though nominal income has not moved, one good, at least, is cheaper. So their effect of income has gone up, and that's represented by this parallel shift from B to X^{**}, Y^{**} . And that last bit from B to there is called the income effect. So the global movement is from X^* to X^{**} with a corresponding movement from Y^* to Y^{**} . And it's broken into two pieces, this substitution effect, which corresponds with movement along the original indifference curve, and the income effect of having higher effective purchasing power. OK?

A lot of today's lecture is just going over this in various ways. You can get mesmerized with the price decrease. Here is a price increase. They're starting out with a tangency.

The price of good X still goes, in this case, up. Their purchasing power goes down. The budget line rotates from this initial point, inward. So the total demand of the price increase is going down, a reduction in X.

But again, we can break that into two pieces. One is a substitution effect, the movement away from good X-- because it has a higher price relative to good Y-- to point B. And now we take into account the loss of effective purchasing power, which is a shift in this budget line, the dotted tangency line, which is tangent at B, to the actual new, lower budget line.

So here, we denote the combination of the substitution and income effect. This is really the previous diagram in reverse. But I think it's useful to show both.

Of course, we can also start changing the price of Y. Let's not do that right away. I'll come back to that later. But the principles remain the same.

So let's construct some demand curves, as in the equation system we had initially. If we keep lowering the price, moving in this direction, we trace out this tangency. And those will be the demand at the associated prices. So here, we have, say, X prime at price, P prime, X double prime at price P double prime-- and this would have been even better if these things were lined up as they were intended to be-- X triple prime at price P triple prime. Now, again, price is going down, and demand is going up.

Now, just to alert you, I don't know where this got started, with Marshall or whatever. But usually, we plot Y is a function of X, where X is on the Y-axis, and Y, the dependent variable, is on the Y-axis. This is actually plotted in reverse because we're varying price as an exogenous variable and plotting the demand.

No big deal. I don't think I can plot it the other way. Actually, I could. But you'll see why this can get confusing. I just thought I'd point it out along the way.

OK. Now we come to Giffen's Paradox. If the income effect of a price change is strong enough, the change in the price and the resulting change in the quantity demanded could actually move in the same direction.

So let me give you an example that's attributed to Robert Giffen, who claimed to observe this in Ireland, where they had a famine. So the potato crop was failing, and the price of potatoes was rising. And despite the rise in price, people consumed more. So the demand curve was upward-sloping, not downward-sloping.

His explanation was, potatoes are a big part of the budget. When their price goes up, the effect of income goes down. And that income effect dominates. The lower income effect is what lowers the demand for the good.

But in effect, prices go up, and demand goes-- well, income's going down. But this good is a necessity. So you get a negative sign on the income effect, so the demand is going up.

So price and income are going up at the same time. It's a bit of a mouthful. No pun intended.

Actually, this could not have happened. It would require that potatoes be an important part of the budget, that potatoes be a necessity. But what doesn't fit the facts is, how are they consuming more potatoes when, in fact, there was a blight, and there were no more potatoes to be had? So something doesn't add up, literally.

Anyway, all intermediate micro classes typically feature Giffen goods and then move on. And you may be left wondering if it ever really happens in practice. So searching around, I discovered Jensen and Miller, who actually ran an experiment. So this is literally a randomized controlled trial. They went to China, two different provinces, looked at relatively poor households, and changed the price of the staple food item that was a big part of their budget.

Now, in order to make the comparison, you've got to change the price. So it was a randomized controlled trial, meaning, randomly, some households were exposed to a lower price because they were given vouchers so that when they went to the store, the sum of U unpaid plus vouchers equaled the stated price. But they didn't have to pay for the vouchers.

So effectively, from their point of view, the price was lowered, given the vouchers. And the control group did not receive vouchers. So the control group faced the normal market price. And what Jensen and Miller are doing is comparing the control group with the treatment group.

And what they discovered was, in Hunan, in the South, rice was the staple good. In Guangzhou, it's wheat. They did surveys, like Engel, before, during, and after the subsidy on both treatment and control groups. And they found strong evidence that poor households in Hunan experienced this Giffen behavior with respect to rice. That is to say, when the price of rice went down, households reduced their demand for rice.

The evidence was mixed in the other province. And I think I will choose not to say very much about this. There were groups, the very poorest and the least poor that responded in what you might think to be the ordinary way, with price-- with demand down-sloping. But there was a group in the middle between the very poorest and the least poor which exhibited Giffen behavior.

As it's written on the slide, this is a bit of a mystery. Why is it that we're getting this difference across the three groups, or two groups plus one? And I leave that for you because, again, the problem set features the entire article. We try to make it easy by highlighting the text.

But somewhere in that paper is a discussion of why this result, which is a bit hard to interpret here, does happen. The authors did give it an interpretation. So that's advertising to go and grab that paper, which I think I already have. Yes, question?

AUDIENCE: Yeah, so is the demand the amount of potatoes or rice you're consuming, or the proportion that is of your income that's increasing?

ROBERT M. TOWNSEND: It's the quantity. But obviously, they have the prices. So you can look at both. But the definition we've been using of a Giffen good, it's-- well, actually, it never showed. But it is the quantity-- in this case, the quantity of potatoes, not expenditure on potatoes.

All right. OK, so now I want to get to something, effectively, we've already been doing, which is compensated versus uncompensated demands. So the uncompensated demand is the total demand, sometimes called the Marshallian demand from Marshall, who was an early English economist. It shows the relationship between price of the good and the quantity purchased, assuming that all other prices and income are held constant.

The compensated demand, sometimes called the Hicksian demand, shows the relationship between price and quantity. It continues to assume that all the other prices are held fixed. But instead of holding income fixed, utility is held fixed.

So to go back to the earlier slides, we were featuring the substitution effect as the movement along an indifference curve-- i.e., holding utility fixed. So when we looked at how much demand would change one way or the other as a function of the price change, and we went along the initial baseline indifference curve, that was the compensated demand.

So more dramatically, if we fix the initial point and change the price, we would be, for a lower price, pivoting along this indifference curve, U_2 . By construction, we keep utility constant at U_2 . And then plot the demands, which are going up, and we put that on the price quantity diagram, and the demand curve is sloping down. But this is the Hicksian demand at little h to remind us that it's Hicks.

So putting both types of demand on the same diagram, we'd have the normal demand, or the Marshallian demand, and h for the Hicksian demand. So as price goes down, demand goes up, but the demand goes up from two different effects. One is the substitution effect holding utility constant, and that's a point on the Hicksian demand. And then the rest of it is this movement to the Marshallian demand, or ordinary demand, which is the income effect that comes from the lower price.

And obviously, we can go in the other direction. Starting here as a baseline, if we increase the price, we would pivot along the indifference curve and get a lower demand due to the substitution effect. But effective income has gone down, so this negative income effect kicks in, and demand is even less. So kind of intuitively, the Hicksian demand is steeper than the ordinary demand curve because it has included those income effects.

Or putting in another way, the ordinary demand curve is, quote, flatter with a lower slope because there's a more dramatic larger change in the quantity as we change the price. Either the income effect is positive, and demand goes up all the more, or the income effect is negative, and the demand decreases all the more.

So we could say-- be careful about this-- that, in some sense, the ordinary demand curve is more elastic. At least here, it has a higher derivative.

So there's a mathematical expression for the difference between the Hicksian demand and the ordinary demand. Namely, the derivative of demand of good X as we change the price of X, the slope of the ordinary demand curve, is first term, the substitution effect of the price change, P_X , but then this income effect kicks in. And there's two pieces to the income effect.

One has to do with how high is the marginal change in X as you change income. The other is, how much X you were eating to begin with. Think of potatoes, and so on, or those budget shares, either Engel or-- but this is X.

This is the quantity. So it looks, from this, since we're subtracting-- say, the income effect is positive, and X is positive. Then we're subtracting a positive number.

So it looks like this minus this must be lower than this. So it looks like the ordinary demand should have a lower slope if this equation is accurate. But in fact, as I just said, it's the other way around. Puzzle, puzzle, puzzle. What's going on? Well, it's this stupid way that we traditionally plot the dependent variable on the X-axis.

So this is dX/dP , the response of the dependent variable to the independent variable, which is 1 over the slope of this. So this diagram is consistent with this equation. It just takes a little bit of juggling to realize why.

Anyway, this is called Slutsky's equation. And again, at the risk of trying to do too much too fast, but on the other hand, providing some motivation for where we're going, eventually, we'll get to lecture 19 and talk about whether there are retractable restrictions on data, as in the science lecture. We have a model. Does the model fit the data or not? Can you even reject the model? Is it providing some guidance? Or you can always tweak things, so the model is always true.

So this lecture 19 is about restrictions that can be tested. And there's something called Slutsky, same guy, Slutsky's Decomposition this equation is a generalization of the one I was just showing you, with the exception that we put this substitution effect on the left and the regular derivative on the right. So it looks like the signs are changing. This becomes a plus.

This is a generalization in a number of respects. We're looking at the demand for good L as we change the price of another good, good K, not the own effect of one good with respect to its own price, but a cross-partial derivative. Although we still have this income effect over here, and we're still weighting it by the initial quantity purchased.

So it turns out that what we want to do is ask whether observe demands-- this guy, these guys, really-- how demand is changing with prices and income, whether there are patterns which could be rejected as being inconsistent with a model. And the problem is that the restrictions are on these Hicksian demands, and we don't see those. We see the regular demands.

We don't see the Hicksian demands. And I'll say more about this in a minute. It turns out that, with some math, you can actually go back and forth between the regular demand and the Hicksian in demand, and there are restrictions.

So anyway, I'm running an experiment now because I'm showing you things that are more toward the end of the class set of lectures than at the beginning. But since we're introducing Slutsky's equation today, I thought I'd show you that we will come back to it and use it in a subsequent lecture. This is a lot simpler notation. And don't worry about either the income expansion paths or this generalization of Slutsky. That was meant to be a preview of coming attractions and motivation, not something you should be worried about learning right away.

So now I want to introduce something called duality. This looks like an income expansion path. As we increase income, we get more and more points. I think this is identical with the figure I showed you earlier.

But I want to use it for a different purpose. I want to say, suppose we're at this point with this intermediate income, I_2 . And we could be maximizing utility subject to that income level holding prices fixed. And X_2, Y_2 would be the maximum point of utility.

But suppose we did a different experiment. Suppose we wanted to achieve this utility level, U_2 , by spending as little as possible. The amount spent would be captured by these various possible budget lines. So something called a minimizing expenditure problem, minimal expenditure problem, would be to find the budget line, which is lowest, closest to 0,0, and still allows the household to get the utility level U_2 .

So at this point, it's clear. You could think of this as maximizing utility subject to a fixed budget or minimizing the expenses to get to a fixed level of utility. Both things are true and, you can go back and forth. And that is the essential idea of duality.

OK, let's do it a bit more with notation. Let's take a consumer with a consumption set in L-dimensional commodity space with the preferences represented by a continuous utility function. And let's look at the Utility Maximization Problem, UMP for short, which is to maximize utility subject to being within the budget. So this says, choose a point in the consumption set, all of the positive [INAUDIBLE], which maximizes that utility function subject to, the points under consideration have to be such that, when you take the dot product of P and X , P_1X_1 plus P_2X_2 , dot, dot, dot, it cannot exceed wealth or income. So that's the utility maximization problem.

The result of solving the maximization problem can be substituted into utility-- in other words, utility at the maximal solution. And we're going to call that the value function or indirect utility function. Note that the maximized utility is a function of the parameters or the exogenous variables of the problem-- namely, the Price vector, p , and Wealth, w .

Now, in contrast, the Expenditure Minimization Problem, EMP, is to-- given a certain level of utility-- minimize expenditures, the dot product $P.X$, that allows you to achieve at least that utility. Getting more is fine. But in most of the figures we show, this will be at an equality.

But just like you could spend less than your income in principle, here, you could achieve a higher level of utility. But whatever you're doing here in the expenditure minimization problem, you're trying to minimize total expenditures. And we call the solution to that the E for Expenditure-- minimized expenditure at the optimum-- is a function of the parameters of the problem, the same price vector p . But now what's taken as a given is u , not w , but u . So we'll have this notation for value function and expenditure functions.

Expenditure functions and demands, for that matter, have certain properties. It turns out, the expenditure function is concave, and the derivative of the expenditure function is demand. So when you take the second derivative of demand, it picks up the concavity of the level expenditure function.

And that's why that Slutsky equation holds because we're looking at the matrix of second-order conditions, effectively. But again, just ignore me if that's confusing. We'll get there in that lecture 19.

For now, I just want to focus on this relationship, this duality. And there's a number of ways to put it. The notation is surprisingly subtle.

But if you work through it, it's fine. So suppose we're given a continuous utility function representing locally nonsatiated preferences with a consumption set, strictly positive price vector. If x^* star is optimal under the utility max problem for a given wealth at those prices, then x^* is optimal in the expenditure minimization problem when the required utility level is the utility you achieve under the utility maximization problem.

I showed you a picture of that just a minute ago. And the minimized expenditure when you solve the EMP is exactly the wealth that you had in the utility maximization problem. Let me go back.

So I started here. We had income I . Prices were given. You maximized utility, ended up at U_2 , given I and P .

We had a certain optimizing solution here called X_2, Y_2 that solves the utility maximization problem. And then I said, what if we fix the utility level instead and solve the expenditure minimization problem, which is finding the expenditure level earmarked by these budget lines, which is minimal and, yet, allows us to achieve the utility level U_2 . So this budget line, which was a budget line for the utility maximization problem, is kind of the minimized expenditure line for the expenditure minimization problem.

I don't know. Maybe the picture is just easier. But the notation says that. And it actually goes the other way.

If x^* were solving the expenditure minimization problem, when the required utility level is parameter u then x^* is also optimal-- the same x -- in the utility maximization problem, where we give the guys wealth equal to the minimized expenditure. And when they solve the utility maximization problem at that level of wealth, they end up exactly with the u we started with in the expenditure minimization problem.

Relatedly-- and this is the most compact notation of all that's kind of cool. Another way to state the proposition is to say that, given equation 3, when we start, given a price vector p , we are fixing prices p , we give the guy income or wealth, w . We solve the utility maximization problem. We substitute the maximizing demands into utility. We get the indirect utility function.

Now, starting with that level of utility, we do the expenditure minimization problem and the minimized expenditures at that the in particular utility level is W all over again. So it's like feeding back to itself.

Likewise here, we start with the parameterized level of u in the expenditure minimization problem, taking, again, p . Vector p is given. At that level u , we solve for the minimized expenditure, which we called e of p and u .

And if we take that minimized expenditure and substituted it as a wealth term and solve the utility maximization problem at that wealth, we would get that the indirect utility is exactly the u that we started with. So it is a bit of a mouthful. The notation is quite compact, and you may just have a white-out moment when you see it. But it is repeating in different ways in related notation.

Oh, this also means that e and v are inverses of one another in the sense that take e and multiply by some imagined e inverse, and you'd get rid of e on the left-hand side and have an e inverse on the right-hand side. So you'd have the value of function is an inverse of v . Sorry, I said that wrong, actually.

If you take e and get rid of it by taking e inverse, when you take the inverse of the expenditure function, you end up with a value function. And likewise, when you take the inverse of the value function, you'll end up with the expenditure function. OK, I said it right at that time. There's a question.

AUDIENCE: You keep saying indirect utility function-- or, I guess, what do you mean when you say indirect?

ROBERT M. TOWNSEND: It has a label. I mean, v looks like value. You say, value function. But it's called the indirect utility function in the literature.

AUDIENCE: OK.

ROBERT M. TOWNSEND: That's the name that we give it. So Hicksian Demand is solving the expenditure minimization problem. And we have the Hicksian demand or demand correspondence if there's more than one solution. And then the diagram, so there's been a unique solution.

This looks like it's a new slide. But actually, right from the very beginning, where we were talking about price effects and income effects, and we defined the price effect as a substitution effect, we were moving along a given indifference curve. So we were really talking about Hicksian demands all along.

And again, for the third time, we have a different way of stating the problem, that the ordinary demand is equal to the Hicksian demand evaluated at a utility level, which is the indirect utility of the utility maximization problem. And likewise, for the expenditure minimization problem.

This is, as I've been saying, true starting at this point. You start looking at it one way or the other. You end up with exactly the same point. But if we start actually taking discrete changes, things are a bit different. And in fact, we were actually doing that at the beginning, right? We were talking about a non-trivial decrease in the price of X.

So let's talk about discrete changes and compensated demand. We start with p and w . We have the maximized demand, x , for that given w . But we change p from p to p prime. Then w prime would, say, be the minimized expenditure at the level of utility, u , the initial level of utility at the baseline but with the new different price vector, p prime.

And we call that new wealth as opposed to old wealth. Call it w prime. So in words, that level of income or wealth w prime is the one that, under the new price vector p prime, would minimize the expenses at p prime of achieving the old utility level. Remember, when we change the price, we just rotate it along the indifference curve going one way or the other. That's all this English and all this notation means.

We're just compensating the agent for the change in purchasing power. But of course, demand doesn't stay constant because the new demand, x^* , will be different as the prices are different. And the Hicksian demand includes the price substitution effect, is defined by it. So the Hicksian demand, at the old utility level and prices p prime, is different from the original baseline purchase.

Anyway, if you understood the income and substitution effects at the beginning, that's all this slide is saying, where we have nontrivial discrete changes in the price. Here's the picture again. But be forewarned that, somehow, to make it more interesting, we're going to change the price of good 2 instead of good 1. And we are going to increase the price. So the original baseline budget is from w to w over p_2 . Why just w ?

Well, there's an assumption down here that we only care about the price ratio. So let's denominate everything in units of good on the X-axis. p_1 , the first good, equals 1. So we don't need to worry anymore about the price of the first good because it's 1. And the question is, what's the price of the second good relative to 1? In this case, the price of the second good, p_2 , is going up.

So you would go from w over p_2 attainable to w over p_2 prime attainable when you go from a price p_2 to p_2 prime. This is like the earlier diagram, where we fixed a point on the Y-axis and pivoted on the X-axis. But this is the other way around because we're changing the price of good 2 and not the price of good 1.

OK, so we start here, and we have an income, a lower budget line. Now we want to break that into two pieces, namely the substitution effect and the income effect. What is the substitution effect? At this higher price of p_2 , it's p_1 , over p_2 . If you remember the slope.

So as p_2 goes up, that slope goes down. So we're going to rotate along this indifference curve and effectively maximize utility, which is giving you this point. What's the maximum utility you can get or the minimized expenditure? Rotating along this indifference curve. And then the income effect is moving from here, a parallel shift back to here.

So our initial level of wealth, w , was on this good 1 axis, and our new level of wealth is this minimized expenditure at the new prices, p prime, given utility level u . So we've gone from w to w prime. That's the change in wealth in terms of good 1. Initial wealth, new wealth in terms of good 1, this must therefore be the change in wealth.

And that's the change in wealth, which is positive, which you may have been wondering about, which is necessary to compensate the agent for the fact that the price of good 2 has gone up. We have to compensate the agent for the adverse price change by giving him or her more income. And this is that level of income.

OK, so as I keep saying, but it's probably been more or less an abstract statement, I'm really interested in real economies, where we have something going on over time, or something with the geography, or something having to do with uncertainty, and gambles, and states of the world. So having memorized or shown you all those figures with some very interesting concepts and a little bit of data along the way, one may, nevertheless, walk away with the impression that we're talking about a static world, where we have rice and other goods or potatoes and other goods. And we're just comparing wheat to apples as we change prices.

But no, the framework applies to any goods, including goods that are earmarked by time. So suppose we have two time periods only. And we'll call consumption C_1 -- consumption and date 1, and C_2 , consumption and date 2.

And we're going to give income as an endowment, meaning how much income they have in the first date, 1, and how much income they have in the second date, 2. Using income is a bit tricky here because we used that before in the static optimization problem. So think of it as money, then-- money in date 1, money at date 2.

Now, what is the budget line? Well, if they just eat their money, they would have consumption equal to m_1 and m_2 . But the way this is drawn, they don't want to do that. They want to move in one direction or another.

So what is this budget line? This budget line reflects the possibility of giving up money tomorrow and getting more consumption today-- that's borrowing-- or giving up money today and getting more consumption tomorrow-- that's lending. So the slope of this heavy blue line represents the tradeoff of the price of goods today for the price of good tomorrow, which you knew all along but probably forgot. Why? Because the slope of the budget line is p_1 over p_2 -- p_1 being consumption at date one, p_2 to being-- sorry-- the price of consumption at date two.

But the economics of the problem doesn't write p_1 over p_2 . It writes r as an interest rate. Why is it the interest rate? Because if you gave up some of your money today and put it in the bank, you would get the principal and interest back tomorrow. So $1 + r$, which is also the slope of this line, reflects the tradeoff between consumption today and consumption tomorrow as made possible by this interest rate. So then, as per the entire lecture, we try to change a price.

Over here, let's increase the interest rate. So it's still possible to eat your money in both periods. But now that the interest rate has gone up, the budget line is steeper.

For a given surrendering of one unit of consumption today, you get a higher $1 + r$ tomorrow. So this paler blue line crosses through the endowment point, but it's steeper. And as drawn with these indifference curves, the household would move from this point to this point.

So let's talk about income and substitution effects. Consumption, tomorrow, at date 2, goes up and up. Why? Two effects, and they're both moving in the same direction.

The increase in the interest rate effectively means they have more income. They can take their endowment of money tomorrow, and it will be worth more. So there's an income effect associated with the price change.

Not only that, there's a substitution effect. Why? Because as the interest rate goes up, the opportunity cost of eating today goes up. It's more painful to eat today because you could get even more tomorrow by putting the money in the bank at a higher interest rate. So the substitution effect moves them with these nice, smooth indifference curves away from consumption today toward higher consumption tomorrow.

So the income and substitution effects in this diagram work in the same direction for good 2, but it's not true for good 1 at date 1. Or at least, it could be ambiguous. Why? There's an income effect, which should increase consumption at date 1. But there's a price effect, which should decrease consumption at date 1. And those can, in principle, move in opposite directions.

In this case, date 1 consumption is going down, so the substitution effect is dominating the income effect. But it didn't have to go that way. You could have drawn the indifference curves and gotten an increase in the consumption of date of that good.

OK, let's start out with a different but related diagram. Suppose they have money today and tomorrow as endowments. And this guy is actually doing what? Originally, borrowing, getting more consumption today than the money he or she has at the expense of paying it back with interest tomorrow, therefore having less consumption than the money that he or she originally had.

So we'd end up with this tangency. And now we still increase the interest rate. That makes this pale blue line steeper. And you can see consumption is moving inwards. So actually, consumption of both goods in the diagram is going down.

But let's think about income and substitution effects. Because they were borrowing originally, as opposed to lending, intuitively, an increase in the interest rate is a bad situation for them. They were borrowing, and now the interest rate is higher. So the income effect is negative. And that could show up with both goods.

On the other hand, the substitution effect should still push them in the direction of higher consumption tomorrow. That's the substitution effect. So for one of these goods, consumption today is unambiguous. But consumption tomorrow is ambiguous, despite the way it's drawn, because the income and substitution effects work in the opposite direction.

So take a look at those. There's a lot going on. In fact, that's kind of the point, that when we move from apples, wheats, and so on with prices and income to an applied problem, like what's happening in the US with these lower and lower interest rates, well, if we look at consumer behavior, there's income effects.

There's substitution effects. You can go in different directions and be in conflict. And that's an important part of understanding the impact on the households and the impact on the economy.

And it's actually a lot more work to go through this slide and the previous one than you might have imagined, even though, say, you were totally sure you could shift these curves around in whatever way you wanted. When we get to the economic application, we kind of do it all over again with the new realistic labels. And it brings that tool to life.

OK, well, that's the last slide for today. That's the end of this lecture. All right, thank you very much.