

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT So today is a review session. But let me start by looking at the class schedule. And today is the review session.

TOWNSEND: So there's no readings for today.

Note the lectures are numbered seven, eight, so today doesn't count somehow. But it does show up on the class schedule. It does not show up as a numbered lecture on the reading list.

I find these, going over randomly selected questions on the study guide, to be very helpful. Taking a bit of time to go over, as we usually do, in this case after the end of lecture seven, is definitely worthwhile. So this question was already asked of me.

So granted, it's kind of a hard question to write and maybe not so obvious how to answer. Go back to the medieval village economy and think about how you would model it in general equilibrium terms. Volunteer.

How do we model the medieval village economy in general equilibrium terms? And he's ready to answer this but we already talked about it. Yes, who's speaking?

STUDENT: I was going to say that it includes the possible consumption set, and the production vectors that are possible, and the endowment for that economy.

ROBERT And the preferences.

TOWNSEND:

STUDENT: Yeah, and that.

ROBERT Yeah, so that's good. So then what aspects of the medieval economy do you remember from figures and so on

TOWNSEND: that look like one of those things? What about production?

Well, either this question is too hard or too easy. I'm not sure which. So one answer I didn't give, the key just now, is we actually thought about the medieval village when we did the dynamics of storage, and starvation, and so on. And we wrote down a production technology mapping inputs and outputs. So we were doing production, and we took a stand on utility. And so that was an example of a medieval village economy, except there was only one person in it. But--

If you go back to those pictures, we have a map with a given person having its plots spread out all over the map of the village. So do you think they have the same production technologies or different production technologies? Anybody can answer. This is a little discussion period.

Well, if it were the same production technology, then why on Earth would they be spreading their land out over the village into different spots? So it seems likely they're different production technologies. In what way do you think they might be different?

I'm going to call on someone. This isn't working. Ewan, can you try to answer that-- in what reasonable way do you think the production technology is different across different plots?

STUDENT: Maybe the multiplier that determines the state of the world is independent, or there are fewer correlations the further out you go.

ROBERT TOWNSEND: Yeah, good. OK. So what in the storage economy was that epsilon might have different realizations for different plots, and then we get into all that stuff about correlations, and covariances, and so on. Yeah.

So as I was saying, the key at the beginning, this is kind of a difficult question. That's why she asked me at right when we logged on at 2:30. It's difficult-- why?

Because you're so used to seeing me show you slides of things, like here's the model, and then you're, like, OK, I guess I better learn this model. And I just put you in the different position of asking you to choose the model. That is a very interesting skill and it is not something you get to practice very often.

That's why I think this question is hard in the sense that you're not used to doing that. But that's the way economists have to think about things. We make decisions.

It's not like God is giving us preferences, endowments, and technology. We're kind of making it up, hopefully for good reason. All right, so that's why I earmarked that question today.

For the set of feasible allocations in an Edgeworth box economy, which allocations are Pareto optimal and which not? And if we had the whiteboard, I would ask you to draw it, but I'm not going to be that ambitious. So can someone just tell me in words? How about Paolo. Can you take a stab at this.

STUDENT: Yeah, sure. So the set of all Pareto optimal allocations inside of the box are the set of points for which the indifference curves for person one and person two are tangent to each other because then you can optimize. There's no lens-shaped region for you to do something that's better. And everything else is just like a non-Pareto optimal allocation.

ROBERT TOWNSEND: OK, that's great, excellent answer. And maybe if we get that far in our review session, all the way back here to lecture seven, we can take a look at some of those pictures. Here's another one.

What is the relationship between maximized linear welfare functions and Pareto optimality?

STUDENT: So whenever we maximize a linear welfare function, the resulting allocation is guaranteed to be Pareto optimal. And I think conversely, any Pareto-optimal allocation optimizes some linear social welfare function.

ROBERT TOWNSEND: Perfect answer. Now I'm a happy camper. Thank you. You guys do really well on these more precisely worded questions. That's a very good answer. Thank you so much.

And let's see-- that's probably enough of that. So I'm going to scroll through these lectures with two goals in mind. One, as I go through the slides, sometimes pretty quickly, please let me know if you have a question. Even if you didn't plan to ask it, you might think of something. It might come back to you when you're looking at the material.

And secondly, I will try to summarize a bit in key spots and draw some connections of one lecture to another one. So that's the plan.

The substance of this first lecture had to do with economics as a science of experiments. And I covered these four different authors. Let me just not dwell on it.

The point was economics as an experiment, in the sense of physics and combining various disciplines to do this. And Koopmans is making this point that it's much more than just measurement. And constantly in the course, I am showing you measurements of things, facts, and statistics, and so on.

But like Koopmans, that's not enough. We want to have a good sense of the context of the problem or economy, but we want to go on and understand behavior and then make predictions. I guess the best illustration of our doing this so far might be the medieval storage problem, where we wrote that we took the medieval village we were discussing earlier, talked about preferences, endowments, and technology.

We did some calibration to put in serious numbers representing risk, and returns, and so on. And then we simulated the economy and pretty much got it to match the data. So that was, I guess, the most full-blown application we've done so far in this tradition of Koopmans, although Koopmans goes on to talk about what he thinks is the real goal, which is policy.

And we did not do any policy simulations in that. We didn't change something and ask if we could remedy the problem, for example. But we will be doing much more of that as the class goes on.

Matzkin has this language about exogenous and endogenous variables. And we've, again in that medieval village economy, I think I was using that terminology-- things that are taken as given versus things we want to explain. Angrist was talking about randomized controlled trials and comparing them to quasi-natural experiments.

This first example was paying bicycle messengers a higher wage. And a couple of comments on that-- when we get to utility functions and utility maximization subject to budgets, we actually had a picture of someone choosing how much labor to supply as a function of the wage. So this should not be abstract anymore. And we actually did look at an RCT in China that had to do with changing the price of a staple commodity, which was featured in problem set one.

So what am I doing? What on Earth am I doing right now? I'm not trying to give this lecture again. I'm trying to tell you that we're delivering on it.

I mentioned these concepts in the very first lecture, was probably quite abstract. But now we're already beginning to be able to fill in with the material that came afterwards. Yeah, I think this is more issues of modeling.

Lucas artificial economies-- I'll go back to that dynamic lecture where I said they could be feeding grain to the animals or turning it into beer. And I said we would ignore that. That represented a choice so that we could focus on what arguably were the two most important storage devices. So we constructed an artificial economy to compare to the real data.

And Lucas mentioned computation. And you can see through these lectures that I have not shied away from telling you how to determine solutions. We did this with Lagrangian stuff, with a set of first-order conditions you could use to help predict and understand the solution.

And then we went to linear programs. We went to dynamic programs. And we'll do more as we go along, depending on what we need to do, depending on the lecture.

And then this list of economies-- we had a great discussion about this on the last slide of the dynamic problem. We shifted gears, and went to a modern US economy with some credit instruments. And I'll come back to that.

But this first economy, the medieval village economy, you've obviously seen depicted in about three different lectures so far. The Thai village economy has entered into the discussion, economies where they didn't try to diversify ex ante, and they need to potentially to share risk ex post. And we haven't done much with trade yet, although we haven't touched these other two yet.

The Lucas trees we did in that dynamic lecture because we talked about yield to seed ratios, how much harvest they're getting per unit of land planted in seed. And then I turned again, at the end of that lecture, to a dynamic problem in a different economy, where they had an asset that was like a borrowing or lending contract with its own return. So we went from farming to financial assets. So hopefully, that's an example.

This we just did-- what's an economy? Preferences, endowments, and technology. We've now got examples of it and we developed notation for it. And throughout, I keep trying to emphasize this can be applied to economies with dynamics with geography and with uncertainty.

Hopefully, you're getting bored with my comments, in the sense of yeah, yeah, yeah, we know. I hope you do know, but my experience has been these things bear repeating. What are we going to do? We did that in the first lecture, then we did it. And now we're coming back to what did we do. So it is a review, after all.

Lecture two was about consumer choice. And we reviewed what's an economy-- preference, endowments, and technology-- as if to say now we're going to focus only on preferences. Motivationally, I'm very nervous about this part of the course and also the same slide for production, because then we kind of forget the whole general economy thing, which is the prize that we want to get to.

But we can't get to the big picture with all the ingredients combined until we understand the pieces. So this piece was about the consumer. And we applied it, though we did try to not just make things to memorize or theory. We tried to apply it at the end of this lecture and in the following two-- maximization of utility subject to budgets.

So we described the consumption set, various special versions that have time and geography. And we've talked about this. This is, again, one of the great, little questions we had at the beginning of the following class.

And we talked about consumer preferences, rank-ordering goods, defined rational choice as preferences which are complete, transitive, and continuous. And then we talked about a utility function representing preferences. A couple of things came up here-- it was ordinal and not cardinal, and we talked about that. But we will and did need cardinal representations when we got to uncertainty and dynamic decision problems.

Second thing is this was a chat on Piazza. This slide goes from having a continuous utility. Any preference that can be represented by a continuous utility function is rational in the sense of the definition of the previous slide.

And then someone initiated the obvious question, does it also go the other way-- that any rational preference function can be represented as a continuous utility? And the answer was, yes, although there was a good discussion about distinguishing continuity on the one hand from the other two ingredients-- completeness and transitivity. So that was a great discussion on Piazza.

Local nonsatiation-- this is kind of a definition of being close by and getting better. I talked about no bliss points. I think in words, we talked about utility functions with bliss points, but we never put any pictures in. You might be thinking, could I draw a picture with a bliss point, and what would it look like?

More is preferred to less. Largely I'm scrolling through this, and hopefully quickly, just to make sure if you have questions or not. Indifference curves, the whole family of indifference curves. Indifferent curves can't cross because then they wouldn't be transitive.

Margin rates of substitution were defined, typically diminishing, although not necessarily. And various displays of into families of indifference for different utility functions, running all the way from perfect substitutes to perfect complements. This one over here, the right-angled indifference curves, represent Leontiefian preferences where there's no substitutability.

We're about to get to maximizing utility subject to budgets. You might play around a little bit and see what happens with these particular preferences when you do that. What happens when you change prices? That would be an interesting experiment to conduct because you'll get really sharp predictions that are distinct from what you would get in this more typical smooth case.

And I'm anticipating reviewing the production lecture in two lectures away. So now, you're able to go back and forth between what do indifference curves look like and consumption versus isoquants or isoproduct lines in production. This one, the Leontief thing, we didn't talk about it much in consumer theory but we did talk about it in production, and you probably didn't make the connection. But I'll try to remember to point out where we actually used it in production.

These are examples of utility functions and definitions of convexity and quasi-concavity. There are a lot of different concepts here. They were all illustrated as holding or not holding, depending on how we draw these pictures.

Convexity turns out to be important, especially when we're drawing pictures of tangencies and optimization. Things can really go kind of crazy when you don't have convexity. Here, it's convexity of the weak upper contour set of it, weakly preferred bundles.

A little more technical derivation of the marginal rate of substitution, and something about homothetic preferences, where the slope of the indifference curves are determined entirely by the ratio of the two goods. We kept coming back to this picture in different ways. In the next lecture in fact, there's something about angles curves, and necessities, and Giffen goods, and so on.

Suffice it to say that if you're moving outwards along this ray-- because maybe your income is increased, holding prices constant-- the goods x_1 and x_2 are increasing proportionally with the distance along this ray. And likewise, since prices are fixed, expenditure is also moving proportionally along this ray.

The preferences were not homothetic. Then this would not be a line from the origin. It would start to bend. And that actually was what was going on when we talked about income effects and substitution effects.

That's coming up in the next lecture. We haven't used this utility function very much, although it will appear soon. And we did talk about utility functions which are completely linear.

So finally, we get to the utility maximization problem-- max utility is subject to a budget-- and defined homogeneity in this consumer theory context, and started to generate these pictures of income budget lines, and the set of all possible purchases in and under the budget line, and the utility-maximizing positions. Although again, this is kind of the classic smooth diagram that sometimes is the only diagram you see. I'm bending over backwards to show you different pictures like this one, where the maximum is not interior.

You do want to keep an eye out about whether they would be happy with this quantity of y , or whether they would actually like to go negative and just cannot do it because of the consumption set. And we formalize that as an example of the use of the Lagrangian, where you write out this expression with all the constraints. And this is the generalized routine for optimization, and actually applying it to the case where we have these non-negativity constraints on the two goods.

This is the most general way to deal with the problem. It doesn't mean all these constraints to are binding. They may be binding and may be not. Typically with non-satiation, a household with this constraint will want to spend all its income, so that's usually binding, although you could try to draw a picture where it's not.

Back, by the way, to the bliss point, but I've never shown you that. And just because it says "greater than or equal to," in other words, non-negative, doesn't mean they would want it to go negative. That depends on what this picture looks like.

And then we did the example utility functions. This is Cobb-Douglas. And you do the algebra, and you get this result that $p_X X$ is total expenditures on X . And it's a constant, α times income. So this is an example of budget shares being constant.

As income vary, X and Y will vary. The amount spent will vary, but it's going to vary in a linear way. It's one-to-one with income.

And that Cobb-Douglas, this guy here, that's homothetic preferences. I don't think I said that when I was lecturing on this. Maybe I did. So in this case, you can almost see it in the algebra-- the marginal rate of substitution, which is here, depends only on the ratio of Y to X because this α is a constant. So this is a homothetic utility function.

And finally, we got to labor supply, which is like an application, since I was worried that you'll lose the context of what we're trying to do. So we had this picture, this beautiful picture, of someone starting with leisure. This was the first time I think we talked about an endowment.

So they're endowed with 24 hours a day or less sleep time. I don't know if you sleep eight hours, but now we're down to 16. And they can take that time and enjoy it in leisure or run it to 0 and work.

This diagram is in the space of goods that consumers want. They're going to end up consuming positive amounts of leisure, positive amounts of consumption, like x_1 , x_2 , with both of them being positive. The slope of the budget line represents the wage.

You might play around a bit and ask if you could draw the picture where the choice is not over leisure but labor, in which case you'd be moving to the left here. And then you'd have to draw the budget line-- it's going to look similar-- and then draw preferences. I'll show you what I have in mind when we get back to that Robinson Crusoe economy.

Now let's go to lecture three. I'm going to spend a little bit less time on this in terms of going through each and every slide, but try to focus more on some of the key aspects. So we had this demand system. This would represent data, but we're going to assume right away it's coming from utility maximization.

So you can think about this as restrictions on data implied by utility maximization. This is in the spirit of Frisch and RCTs and various other things. So we're increasing income. What would happen?

This is the linear expansion path. But it doesn't have to look like that. It could look like this. So this actually has a negative slope as you increase income because the amount purchased of goods is going down.

It could be entirely vertical, actually. I think you could imagine how to draw it that way. That utility function I said we didn't do much, which was transferable utility, x_1 plus $\log x_2$, that would give you a vertical line like this, or horizontal. I don't remember which good was which, but now the quantities are going down or up and the prices are constant. Because all we're doing is moving income.

Hence, the expenditure shares are moving in this diagram, not just the quantities, because of the nature of the experiment. So we defined inferior goods and normal goods, and got to Engel. Engel's law being, in particular, proportion of total expenditure on food goes down as income goes up.

So someone asks me-- I think it was in the Piazza chat or was it in class-- no, it was in class. One of you asked me, are we talking about quantities or shares of expenses? So it's kind of related. In order for Engel's things to be true, some good must be, by definition, a necessity. But necessities were defined by this diagram where we just changed income and didn't change prices. Just trying to tie these loose ends together.

Then we showed you what's going on in the US. One of you made a great comment about this diagram, it's kind of being tautological. And my answer was yes, but it's also quantifying the differences.

And here we did the income and substitution effects from price changes, in particular emphasizing the income effect of price changes for increases and decreases in prices, and plotted out the demand curves. Some people just start with demand curves, like this. So where did it come from?

In particular with aggregate demand, is it down-sloping like this? What's going to happen in the economy as a function of the virus has to do with these elasticities and so on. But I prefer not to start with this as a primitive. I prefer to derive it and put restrictions on it from these underlying first principles.

Then we defined Giffen's paradox-- that in the famine, the Irish-- which destroyed the potatoes, they were eating more potatoes. We decided that couldn't literally be factually true, and that Giffen's thing, I thought, had never been tested. And then I discovered Jensen and Miller, and they shared their data. And of course, the paper was published and so on.

So this was an important part of the course for two reasons. It's meant to bring the concepts of Giffen goods, and hence the distinction between income and substitution effects, to life in real data. It was also, as I said earlier, an example of an actual RCT in practice.

So it's not just as if we were running an experiment. They ran the experiment. They actually did it. They had households that were treated, households that were control, looked at behavior, and they did find Giffen goods. So it's an important part to bring the theory to life.

The second thing is it was a big part of problem set one. And that came with the experience of reading a professional paper, having the data there ready for you to manipulate, run some code to replicate what these guys did, and to extend it. And many of you answered those questions. Some of you did not or didn't answer all of it. At a minimum for now, you should understand what it was that these guys did, and why they did it, and what they found, quite apart from answering that problem.

So uncompensated and compensated demand, also known as Hicksian versus Marshallian demand. This eliminates income effects. So you move along indifference curves when you change prices, so it's all driven by assumptions about the marginal rate of substitution.

We tended to draw these smooth indifference curves with declining marginal rates of substitution, which is the right way to get started. But this concept should work as well when you have linear indifference curves. So that links back to that other diagram that I was mentioning, the one with the almost-linear indifference curve and the price.

But in this case, we'd be doing that and also changing the price ratio. So we have compensated and uncompensated demand, leading to Slutsky's equation. And I don't know if I helped or hurt, but I actually jumped ahead to lecture 18 or something, and showed you eventually how we're going to use that. But I also told you to ignore that detour if it was confusing.

But this whole lecture is about making predictions as prices change or income change. And this Slutsky equation turns out to be something you can measure in data. Even though we never see the compensated demand, the compensated demand has properties because the expenditure function is concave.

So we can actually put some sign restrictions on the Marshallian demand in this way. It's all a bit surprising, but more on that later. Anyway, it's about making predictions.

And then we did duality, which I was just alluding to, minimizing expenditures subject to-- here's the picture. So this was the first time we talked about maximizing utilities subject to budgets, hand in hand with minimizing expenditures to achieve a certain level of utility. So we defined a dual problem relative to the primary problem, and got some useful characteristics out of the dual.

Although in fact, if it's a very teeny, weeny change, and it's all local, nothing happens. They're equivalent. At least, they're equivalent with these other underlying assumptions-- nice, smooth indifference curves, no bliss points, et cetera.

You could experiment with putting in local satiation. And then you will discover that these things don't have to be the same. Maximizing utility subject to a budget and minimizing expenditure to achieve a certain level of utility will not be the same when you violate the assumption of local non-satiation.

Anyway, we developed all this in the math, in the end of this lecture, and then ended just the way we did the previous one, with, hopefully, another attempt to bring this stuff to life, which is to give the household money incomes denoted by a first date or second date, and then talk about price changes. So it is about the distinction between income and substitution effects. But here, unlike income as exogenous and moving around, you have money incomes in the two dates.

All we're doing is changing prices. We're changing the price of the good at date two relative to the price of the good at date one. So there were two different pictures here depending on whether initially, the household, as here, is a lender, eating less than its initial endowment.

We talked about an income effect. You can see that as we move from the solid blue line to the pale blue line, in this upper quadrant at least, the new budget line lies further toward the northeast than the original one. And since they started here, locally the effect of the interest change is to give them more income.

But they don't have to end up with higher consumption of the second good. They could have ended up with lower consumption. It's really hard to focus in there. I have to put it right there or something.

So it would be a little bit lower than it was initially, and that's despite the increase in income. So that's the case where the substitution effect is getting washed out. So again, here, they start out as a borrower, and now when the interest rate goes up, they clearly have less income, which intuitively makes sense because they were already borrowing.

And an increase in the interest rate can only increase the amount of debt if they stay put. But they would not stay put. Then the two effects have to do with income and substitution effects. They have the same sign for consumption today, but opposing signs for consumption tomorrow.

So I think I went back the next class and made the comment that this is an important stepping stone toward modeling entire economies because we're giving the household an endowment. In the labor example, we gave the household an endowment of time. Here, we're giving the household an endowment of money income in each of two periods.

So the endowments are part of an economy. Preferences are clearly being described here. Production was suppressed because it's just a consumer problem. So that was lecture three.

Next up is lecture four. So I may speed up a little bit. That said, there's a ton of stuff in lecture four, and we did not get through it deliberately the first time. There's three or four slides at the end that we moved to the next lecture.

Because by topic, it's about production, but there are just too many important concepts. Repeat-- preferences, endowment, technology-- now zero in on technology. That question that we were discussing about describe medieval villages as an economy also made reference to Northern Thai villages as an economy. Here's the picture.

So what's the technology here? Land, labor, and capital are used as inputs to produce this crop. So this, in many respects, should not have been hard. But then again, I'm not asking you to memorize all these beautiful pictures, just saying that because I took this picture. Sorry.

So production-- we have inputs and negative, outputs, which are positive, and vectors, and indexed by particular goods and by firms, potentially. We have this production possibility set, where inputs being negative, like labor over here, and outputs being positive. This one displays a region of increasing and decreasing returns to scale, to be defined momentarily.

We went through these properties of production functions. And there was a good discussion coming up, namely decreasing constant and increasing returns to scale. This was the picture of decreasing returns to scale.

This is a picture of increasing returns. This is a "typo" because according to the definition, you should take a point in the production set and multiply it by a λ greater than 1, which I mentioned in class. We have not tried to correct the picture.

That's because what this displays is a violation of decreasing returns relative to them being satisfied on the left. They're not satisfied on the right. So the picture is not wrong, it's just not a direct use of the definition of increasing returns, and this was constant returns.

So this is aggregating production sets. This one plus this one equals this one plus the endowments. So there was a lovely discussion on Piazza about this thing, which is, have we generated a free lunch?

Evidently from the picture, the answer is no because there is that finite production possibilities frontier over here. Depending on where you are, you're using one as an input and the other as an output or vice versa. The possibility is created here, however, that you use x as the input, getting y as the output. Take y as the output, put it in as an input in the second technology. Take y as the input, get x as the output.

So if you kept going like that, and you got more in output than you used in input after going through one cycle, you could go to infinity. There'd be no end to it. It would not only be a free lunch, it'd be an infinite lunch.

You could just keep going, in which case, although it's not entirely obvious, you could never draw this picture because there'd be no boundary. You'd always be in the interior. You could always do better just by cloning.

And then we define the firm's problem, maximizing profits. We drew these pictures. Again, you get into the habit, if you're an economist- maximize subject to. Maximize linear function subject to convex constraints. You end up with these tangent cases.

In this case, these are isoprofit lines not isobudget lines. We want to move out as far as we can on the isoprofit line. But the tangency is a similar thing if we move across these applications from consumer theory to producer theory. So hopefully, it's starting to become second nature.

We defined properties of the production function, including convexity. And we proved it. And then almost out of order, this definition of homogeneity popped up. So one thing we didn't do was link these definitions of homogeneity of functions back to the discussion of increasing, decreasing, and constant returns to scale. But you should be able to do that.

This was a proof of Hotelling's lemma. And it was an illustration of the envelope theorem, which is another tool that's going to keep popping up. Special versions of profit maximization, cost minimization-- so this is like the dual, except this one's going to get used, not just to show equivalence.

What I mean by that is this-- so we have this isoproduct curve called isoquant, and we're going to try to find the cost-minimizing bundle that allows you to achieve that given level of output. Then you could do two experiments. Just like with the household, we varied income and varied prices.

Here, you can vary prices. And that's the Walt Disney thing. So the input mix would change.

What you could also do, although we decided not to overwhelm you, is pick a different level of output-- a higher level-- and find the cost-minimizing combination of inputs to support that higher output. And the definition of a cost curve was defined. Cost function takes the input prices as given, but also the output.

So this is kind of obvious when you say it. This is the cost of producing a given output. Yes, that's true. But we never showed the picture.

But we did show the costs. They're sitting here in this slide. Average cost, marginal costs, and what happens when you change prices with constant returns to scale? If the price were too high, the outcome would not be well defined because they would make infinite profit.

And then when I presented this, I said there was a typo, and there is no typo. So I apologize twice. It's fine.

What's going on is the dialogue in the text below is moving from here to the intermediate line to this line. So when price of the output is going up, this thing is going down. These lines are getting flatter.

And when they get sufficiently flat, then you go ballistic and you can make infinite profits, which is consistent with this diagram. Except that it's never defined at a higher price because there is no determinate outcome. Then this thing I will set aside, but again, I asked you to try to write down a little bit better proof that these input ratios are constant when you have constant returns to scale. Although this is a version of in consumer theory, the idea that with homothetic utility, the expenditure shares don't move, and so on. So it is related to that.

This was a picture of Robinson Crusoe. This was the first general equilibrium economy although it was only one person, with a production set and preferences. When I said you should be able to take that labor supply and convert it by making the labor supply negative, I had this in mind, except that unlike a production set, you just have a straight budget line emanating from the origin, and then you would maximize utility.

You're just basically flipping the diagram from right to left like an overlay. But here, there is production. And this was an exciting part because we got to trade.

So now we're talking about a real, live economy with two goods, agriculture and output, and looking at what the consumer would do if they were maximizing utility subject to that production set, as opposed to what they would do if they were able to trade in the world economy. They would produce more of the manufacturing good, and export it, and consume both goods in such a way that their utility is higher.

So this is-- I can't remember if I emphasized it, but this illustrates the gains to trade. It's the gains to trade for a single consumer like Robinson Crusoe, which you can see evidently has higher utility because they get to trade away from the origin. Of course, we should ask ourselves whether this is a realistic description of the economy or whether we should populate it with more people, and potentially have heterogeneous utility functions.

And for that, we're going to need to talk about whether or not we can aggregate up, which comes up in lecture 15 or 16, or something like that. And this is the US and England trading with one another. Questions?

So I'm going to be brief. Otherwise, we're going to go over. This is decision making under uncertainty, and also illustrative of linear programs. So I think on this one in particular, I want you to let me know if you have questions.

It is a bit unnatural. This part's all very natural. It just has to do with a household evaluating nontrivial lotteries over prizes, although we used it to define means and variances, coefficients of variation, and variances, and correlations. We did all of that.

But then we went back to the award or the prize being two-dimensional, in this case, goods X_1 and X_2 . So this is a discrete, non-convex consumption set. And miraculously, we turned it into a convex-looking set. And we had a good discussion at the beginning of the following class about how to do that.

But I want to make sure you all understand it, because we're going to build on it in subsequent lectures. So one of you asked me in my office hours, it's kind of weird-- in finance, you take as given the returns on the assets or their beliefs about the returns. We don't talk about choosing the returns.

But here you do, because you're choosing the lottery. You're choosing a random way, although it could be subject to a budget. I'm practically out of time.

We went to concave utility. We talked about behavior under risk, different types of utility functions, this tree that we keep coming back to that was at the end of the lecture last time. This is Arrow-Debreu security, so in this case, instead of having a pure borrowing and lending economy or borrowing subject to a credit limit, we could have had them trading in an asset which yields a return of 1 in one state and nothing in the other states. And you could be either making that promise as a borrower, so to speak, or buying it as an investor. So we can go from medieval villages to sophisticated securities.

And then we kind of talked about this at the beginning of class. And this class ended with linear programming problems and how they crop up in general in utility maximization and linear activity analysis. This linear activity analysis is the Leontief thing I was alluding to when we went over preferences.

These are constants. These A_{ij} is the amount of good i , input i , required to produce good j , and it's a constant. The rate of substitution isn't varying that delivers these L-shaped isoquants.

All right, so I'm going to have to stop because otherwise we're going to go over. So we almost got to the end. I don't think I have anything usually uncovered anyway. Because the other two lectures were the dynamics, which should be fresh in your mind-- we just covered it recently-- and the one that we did and reviewed at the beginning of the lecture today on Pareto-optimal allocations.