[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: Well, greetings everyone. First of all, in terms of the reading list, today we're going to finish the two, three things I deliberately didn't do at the end of the lecture on firms and production sets, which was largely Kreps, but with plenty of stuff added. I should say the stuff that's added is largely anticipating an application we got to at the end, which is the impact of tariffs and so on, on economic welfare.

So you'll have the tools at the time we get there, although we will also provide some reminders. The second thing is, although it's not starred, there are these things about-- other things on the reading list, which are not starred in this case, about the great Japanese earthquake. And that's-- I'm going to talk about that momentarily, but I know everyone's quite busy, and there's plenty of required things to do.

But if you get a chance, I'm sure you actually would enjoy taking a look at that article. Today, the main thing I want to draw your attention to is decision making under uncertainty and also linear programs. And that material is largely coming from Nicholson and Schneider, although it does have some material that was moved relative to earlier years. As you will see, this Medville, which is short for the medieval village economy-- several of the slides today have material on that from that book, which you therefore might find helpful.

In terms of the study guide, we have some questions here. And let me try something a bit new-- but nevertheless, let me try. Armando, could you take a crack at this one-- define increasing, decreasing, and constant returns to scale and production?

AUDIENCE: Yeah, so first, constant return to scale means that if you double the inputs, you get double the outputs. Increasing means that efficiency increases as we scale it up, so if you double the inputs, you get more than double the outputs. And decreasing is that if you double the inputs, you would get less than double.

ROBERT TOWNSEND: Yeah, that's fair. That's a pretty good answer. We can define these things in various ways. They're all equivalent with one another. The way you did it is fine. And in some respects, the definitions are the reverse of what you would expect. Decreasing returns to scales means that if you pick any-- if you pick zero and any other point in the production possibilities set, and then pick alpha between 0 and 1, that that point is also in the production set.

I say it's weird, because decreasing-- you're actually scaling back as opposed to doubling and moving forward, but it's the same definition because you do get that curvature in the production set, and likewise for increasing returns. So thank you, that's good. All right, I marked a few of these. We don't have to do all of them. How would you prove that maximized profits are homogeneous of degree 1 in prices p? Caleb, if I'm saying your name correctly-- Aulins?

AUDIENCE: Yeah, sorry, what was the question?

ROBERT TOWNSEND: Prove that maximized profits are homogeneous of degree 1 and prices p. How would you go about proving that?

AUDIENCE: I'm not sure.

ROBERT TOWNSEND: Do you recall what homogeneous of degree 1 means?

AUDIENCE: No.

ROBERT TOWNSEND: OK, so homogeneous of degree 1 means that if you increase every element of the price factor by a constant that maximized profits, increased by that same constant.

AUDIENCE: OK.

ROBERT TOWNSEND: Do you feel comfortable taking a crack at the rest of this?

AUDIENCE: I'm not sure what the answer is.

ROBERT TOWNSEND: OK, all right, that's fine. Wang Ping-- and correct me if I'm saying your names wrong. I'm sure I am.

AUDIENCE: In the future, you can just call me Andy.

ROBERT TOWNSEND: Andy, OK.

AUDIENCE: Yeah.

ROBERT TOWNSEND: That's the A in your middle initial.

AUDIENCE: Yeah.

ROBERT TOWNSEND: OK, so let's see. Do you want to try the same question?

AUDIENCE: Sure, yeah, so homogeneity means basically, if we double the prices, then we also double the profit.

ROBERT TOWNSEND: Right.

AUDIENCE: Right, so I believe that profit function is equal to-- is p.z, where z is the production vector.

ROBERT TOWNSEND: Yes.

AUDIENCE: And so if we double the price p, if we double the-- if we double the production, so if we maintain the same production vector, z, then we can guarantee that we at least double the profits. But if we were able to do better, then we could then halve the prices again and obtain a better-- basically, a better production for the original price level. And so basically, the optimal profits at the doubled prices must exactly equal twice the value at the regular prices.

ROBERT TOWNSEND: Right, I think that's the correct answer. The way I like to think about it is if we-- graphically, if we were doubling the price of every component of the price vector, p, then we're not changing the slope of those isoprofit lines. And so the tangency with the production possibilities frontier would be the same point. And that's why the z, as in p.z is not moving. Great, thank you.

And let's try just one more. What is the definition of an isoquant? So Ke Qiu--

AUDIENCE: I think the definition of isoquant is like, there was a set of input boundaries which, given the productivity, they will output the same level of output y-- all input bounders which produce the same level of output.

ROBERT TOWNSEND: Yes that's correct. Great, thank you. OK, so I was too ambitious and I marked off too many of these. As I said, I sort of choose at random. In this case, this morning, I just marked a few of them. Take a look at the other ones. As always, these review questions are for your benefit, for the class. I'm only calling on you now instead of just taking volunteers because I like class participation, and it's just a reminder to try to keep up rather than having to make up lost ground later, when you're studying for the midterm or something.

If you keep up with the class, you'll be fine. OK, so let me go to lecture 4 and talk a little bit about three related examples of production I featured in the lecture, the example of trade. But here are some others that actually do relate to what we will get to at the end of lecture 5 as well. First off, I want to talk about Leontief input-output matrices.

So Leontief's view of the economy's production possibilities was interesting, albeit a bit simple. Suppose there are three sectors, for example, as in this matrix. You have a sector representing raw materials, which could be agriculture, or mining, and so on, and a sector representing services, which would be retailing, advertising, transportation, lawyers, real estate agents, and financial services-- could go on and on.

And manufacturing, which I guess is pretty self explanatory-- this entire page of lecture is just meant to be representative. Obviously, you can have many particular rows, depending on your willingness to disaggregate further and further into in more and more detail. So this matrix, what it's saying is the motivation for this is you have some final demand, and how much would you need to produce to satisfy the final demand?

You could have final demand for services, but the problem with that is that's the final demand, not all the intermediate demands. Services are in turn used in other sectors. And in fact, the column here represents that to produce $1 worth of services-- services are used in other sectors, as in this row. Services show up in producing raw materials, services show up in producing services because of that aggregation problem, and get used in manufacturing.

Likewise, if you go down the column, to produce $1 worth of services, you need $0.04 worth of raw materials, $0.03 worth of services, and $0.01 worth of manufacturing. So you can see, its input-output-- these are both in some sense direct outputs, but also indirect inputs into other sectors. This is a version, by the way, of the question having to do with aggregation.

Last time, we had a slide showing two goods, x and y, I think-- or was it Y1 and Y2. And one used Y1 as an input and used Y2, and the other industry used Y2 as an input to produce Y1. That was decreasing returns. This is linear constant returns, because everything here is linear, but it is the same idea of inputs and outputs. So we take this information in the input-output matrix and summarize it in a matrix A, and here we're just reproducing these row column elements exactly.

OK, so the problem is we have final demands, which are given, which is a three dimensional column vector. And it tells us how much we're going to need to end up delivering either to consumers or for exports, of raw material services and manufactured goods, in that order. So if we think about these equations here, we want to produce x.

And then we want to figure out how much to produce if we want to meet the final demands of 400, 200, and 600. But in addition to that, some of that output gets used in turn as inputs, and is not available for consumer and export demands. So the key equation here would be, if we need to produce x, but we lose a times x because some of that output gets used as inputs, then we have to produce i minus a times X. I is the identity matrix here.

So it just picks up the same-- it is 1 on the diagonal and 0 on the off diagonal. So for example, ia would just give you x back, as if you didn't have to use things as inputs. Ax, that we've been going through here, gives you-- if you're going to produce x, how much of that x gets used up in producing services, raw materials, and manufacturing? And what's left, I minus a times x, should just equal demand.

So this is a linear equation in the control variables, namely x. And also, demands are given. Although to anticipate a bit, you might think there's some trade off in demands, or maybe some linear combination of demands is fine, not just specific numbers. And then you could imagine maximizing demand, but in any event subject to this constraint.

Here, it's simpler. We're just given the demands, and we have to produce x in such a way as to meet the demand. So this is a linear system of linear equations. And if you remember a little matrix algebra, all we need to do is to invert this matrix, in minus a, which eliminates it on both sides-- on the left hand side, and puts it on to the right hand side. So the x solution we want is i minus a times dx.

So I going to say in this case, you can go back up and realize that this was 0.2, so i-- 1 minus 0.2 is 0.98, and so on. Nothing gets subtracted on the off diagonal, because i is 0 there. So as long as this is what we would call a non singular matrix, that the rows and columns are independent of each other, one is not a multiplicative factor of any of the other two, then this inverse exists and we can solve the problem.

And you can see although it should be intuitive by now that they need to produce more than the 400 of raw materials. They need to produce 449. Actually, these were denoted in billions, and likewise, how much more of each relative to the exogenously given demand actually depends in a pretty complicated way on that matrix A, in the following sense. To produce a unit of services, one has to take into account the inputs of manufacturing and raw materials.

But to produce manufacturing and raw materials, one has to take into account that they too utilize services, and so on. And actually, you can kind of keep going back into the matrix. Another way to write the system of equations, actually, to get the inverse, is to successively plug-in first approximation, second approximation, third approximation as you take into account more and more of the layers of the interaction.

So you could have a sector which doesn't contribute much to final demand, but is crucial as an input into the other two sectors. So that sector could have a high importance, even though it's not directly satisfying the demand of itself. So then we come to Google, and in particular their algorithm. When you type in a keyword, it pops up in a rank order the most likely things that are related to what you just searched for.

This was invented by Sergey Brin and Larry Page. I sometimes wonder if people's last names are not predictors of their eventually careers, Larry Page being one of the fathers of PageRank, but anyway, that algorithm is really based on the same ideas that underlie Leontief's ranking of industries, which is you Google up something, and it may be that there's an article that you might not have thought to be the most important one, because it may not seem to answer the question directly.

But in fact, it's the key article that started it all, that has been cited and used by other things that are more directly linked to the Google item that you're searching for. And this display is not by sectors of industries, but just by a network diagram. You can see these things directly related via pathways to things in the center, these outliers or clusters. And then it gets very thin, but they could end up being important because they're related to things in the interior.

These things are ranked by distance, actually, reordered in order for the diagram to display the degree of closeness. And they-- Page and those guys actually cited on Leontief and output matrix. And then as a final example, the Great East Japanese earthquake-- to remind you, there was an earthquake, a big one, and then a subsequence tidal wave. And actually, the coastline I think subsided by a meter, or some outrageous amount.

And so in came the tidal wave, doing a lot of damage. And this is highlighting the damaged structures. And this is highlighting the headquarters of firms whose headquarters' location was destroyed in the immediate aftermath of the earthquake and tsunami. By the way, the Fukushima nuclear power plant was right there, adding to the mix. So then we move over here. This is all of the main island of Japan. And this is the area that's highlighted, the small area here.

And then we ask, as an input-output sense, which industries had been supplying to headquarters which is not per se just where all the management is, but like the primary production facility. You have inputs supplying to General Motors, for example, in Detroit. You have tire manufacturers that are linked-- this is Japan. So we have input suppliers linked to headquarters, or intermediate. And likewise, in principle, it's not just inputs from upstream, but outputs downstream.

The disruption here means you can't produce something which is in turn supplied downstream, quote unquote, to others. So the first order connection, as in the PageRank, are these areas in red. And now, we've covered all of Japan, because there's either an immediate input supplier or output purchaser indirectly damaged in the sense of the first degree, again, that's the primary area of damage.

This is second degree damage, A linked to B, who was hit with C, and so on. So eventually, as you go in this case down to the fifth order, most of Japan was covered. And in fact, this Japanese earthquake caused a drop in national level GDP of about 1.5%. It is, though, just another illustration of the same mathematics.

So now, with some trepidation, let me go to lecture 5 and verify that you can see this on your screens as I move through it. Yes, you can see it right? So this is decision making under uncertainty, and then we'll get back to some of our techniques, in particular programming, linear programming. So let's suppose we have a lottery, and there are little n distinct prizes, possibly zero, possibly negative if it takes money away.

So these X's are the amount of the winnings or losings, if that even turns up in the lottery. There are n possible things that can happen, and they happen with these probabilities pi, pi1 being the probability that the gains or losses is in the order of x1, pi2 for x2, and so on. Because it's a lottery, the probabilities themselves are not negative. They're 0 or positive. And it's a comprehensive lottery, choosing over this vector of possible outcomes, so the probabilities add up to 1.

In other words, the lottery is a well-defined random variable, with these n distinct realizations. And we can call the expectation of the lottery, the expected outcome of the lottery of the random variable capital X, just to be the weighted average of the outcomes where the weights are these probabilities. So part of this-- these next few slides have to do with the ideas of lotteries over discrete items.

Part of the next few slides are quick reviews of some basic ideas and statistics. Here's a basic idea from statistics. If we have a random variable, capital X, when we want to talk about its variability or its variance, denoted is sigma squared X, that is just the dispersion of X around the mean, mu X. Mu X was defined here in the previous slide as the mean squared. So we square the differences. If X were equal to its mean, this would be 0.

The further away it is, the higher the squared value would be. Expectation again means that we're just taking weighted averages, with the weights being these pi I probabilities. So this would be the formal expression. The variance is the proxy for risk. The square root of the variance is called the standard deviation. And if we want to normalize it so we can interpret the units of it without knowing the magnitudes of the underlying variable, we divide by its mean.

So the standard deviation divided by the mean is called the coefficient of variation. What I mean by normalizing isn't just that we're dividing, but the reason for dividing-- you could well imagine if we have two random variables and one is twice the other one, then all the values are doubled, and the measure of the dispersion will go up too. So if we wanted to measure variability, that's kind of independent of the scale of the underlying random variable itself.

One more slide on statistics-- so suppose the prize has two dimensions, X and Y, not just the winning and losings of dollars, but maybe two different commodities. So a particular realization or outcome would be a specification of both the value of X and the value of Y Xi, Yi for realization i. Realization i happens with probability pi i, which is the same as before.

Now, though, we can define the covariance between two random variables, x and y, to be the expectation of the product subtracting off the means. So X plus its mean uX, times Y minus its mean uY. And I didn't write it out, but you could just sum it up over the same pi is instead of writing the expectation operator. It tells you whether the variables x and y are moving together or moving in opposite directions.

So for example, when a high value of X corresponds with a high value of Y, thinking collectively about these realizations Xi or Yi, then this covariance is positive. And likewise, if a high x means a low y, the covariance would be negative. You could think about plotting the Xi, Yi variables as in a cloud and trying to figure out if the cloud is running on the 45 degree line or running in the reverse direction, from northwest to southeast.

Again, it's useful to normalize, so we take that covariance and just divide it by the product of the variance, called the correlation. And correlations are by construction going to be between minus 1 and 1. So again, 0 means no correlation, 1 means perfect correlation, positive direction, n minus 1, means negative correlation. So those are the three statistical review slides, and we're going to come back to them momentarily.

They're all defined for random variables, but we defined a random variable by first thinking about a lottery. So let's come back to the lottery. Here, in fact, the underlying ingredients are x1 and x2, but they can only take on a finite number of values. These dots represent a discrete point, and this commodity space here, the values of x1 and x2, which are feasible, this is not a convex space, because you can easily find lines between any two of these dots.

And intermediate points on the line other than the end point are not in this consumption set. So this is a non-convex set. Back to the lottery interpretation, lets call, say, Cs a particular realization or state or point in the commodity space-- and capital C. S denotes the set of all possible states. It's just counting up the number of dots you're seeing on this picture.

So for example, if there are capital S dots or states, each state corresponds to a particular consumption bundle in R2, and we can call pi S the probability of being in state S. So we could have a lottery over this space. If we define those lotteries, a particular lottery would put a non-negative probability, perhaps 0, on any particular state. And the sum of the probabilities over all the states would have to add up to 1.

So that's a particular lottery. The set of all possible lottery is the set of all probabilities pi that satisfy these two characters called, call it capital Pi. It's also called a simplex. By definition, a simplex is a set where each element is non-negative, and all the elements add up to 1. So a lottery satisfies the definition of a simplex. Other things do, too. OK, so this becomes the new commodity space, not the set of underlying points, but the set of lotteries on those points.

And likewise, we now can define preferences with this squiggly script inequality over the set of possible lotteries. And what we're going to use is von Neumann Morgenstern expected utility, which is quite natural. And it's saying, take a particular point Cs in that space, take the underlying utility of it-- which we've been using before. Take the probability of that state and then sum it up over all possible states and define this expected utility.

So these are the underlying real value numbers of deterministic utility. This object over here on the left is the expected utility for a particular lottery pi. And the idea now is that all possible lotteries can be ranked by this expected utility criterion. Let's just note in passing that this expected utility criterion is linear in the choice objects, which are the lotteries. These u of Cs things become coefficients, real numbers.

But we're not varying the underlying commodity points. So these become fixed numbers. And the thing that we're interested in varying would be the probabilities that constitute the lottery. So the objective function is linear in the probabilities. There are some criteria we need to satisfy in order to rank using expected utility. What we want to get at the end of the day is that a lottery pi is strictly preferred to a lottery pi prime, if the expected utility under pi is strictly greater than the expected utility under pi prime.

In order to do that, we need three axioms, these three. The first combines two-- cheating a bit. Completeness says that all lotteries can be ranked and give us a number. And transitivity, I think you can see, is going to be satisfied, because if A is preferred to B, lottery pi is preferred to lottery pi prime, and pi prime is preferred to pi double prime. Under the preference ordering, we can map those preference orderings into these expected utility objects, which are just real numbers that satisfy transitivity.

So those are intuitive properties, completeness and transitivity. Here's something called independence. Take two lotteries of focus, denoted pi and pi prime, and say that pi is weakly preferred under that expected utility ordering to pi prime. Then it has to be true that some alpha weighted lottery-- this is a lottery, pi is a lottery. And pi prime is a lottery. And pi triple prime is also a lottery. So now we're taking compound probabilities, if you like, when we take the alpha weighted sum of the lottery pi and pi prime.

If alpha were, say, equal to 1, then we're already given that pi is weakly preferred to pi prime, and that would be an equivalent statement on the right hand side. If alpha is less than 1, however, then we're putting some weight on this extraneous third lottery, pi double prime, and the notion of independence means that whatever this other lottery looks like, it doesn't matter for fixed alpha when it comes to weighting pi and pi double-- weighting the composite lottery.

Because each one of these, The second term in each one of these expressions is 1 minus alpha times pi double prime, which is the same on the left hand side as it is on the right hand side. So the ranking is independent of whatever that term is. It seems like it ought to be true, but it's actually an axiom. Continuity is similar, although not identical.

It says if we have three lotteries pi, pi prime, pi triple prime, and they're rank ordered-- pi, preferred to pi prime, preferred to pi double prime, then there exist numbers alpha and beta such that this weighted sum of pi and pi double prime, weighted by alpha, is strictly preferred to pi prime, which is strictly preferred by the beta weighted sum of pi and pi double prime.

Now the way to see this is that it's not required to be true for all alpha and beta, it's only required to be true for some of them. So if we were allowed to set alpha equal to 1, we would have alpha pi on the left hand side, and pi prime would be over here. We'd have pi preferred to pi prime, but that was a given to begin with, so no new information. This is saying if alpha is arbitrarily close to 1, then that ought to be true, and hence the notion of continuity.

Likewise, on the right hand side, if beta were zero, which it's not allowed to be, then we'd be just getting pi double prime here. And we're already given that pi prime is preferred to pi double prime. The content of it is if beta is not 0 but very close to 0, then we should preserve the rank order. These are intuitive properties. They're here only to say that expected utility may not be, quote unquote, as general as you might think when facing problems of choice under uncertainty, although it is very natural.

And these properties are similar in some respects to what we did before with ordinary preference orders. And with regard to what we did before, there was already a question from one of you about this, what about choice under uncertainty? Now, our measure of utility becomes cardinal, not just ordinal. We are really going to care about the degree of curvature in this orange line, which is mapping the expected utility outcome as we vary the input x.

So what's going on in this setup here? We have income, say, moving around, and it has a mean expected value of x, and it has a couple of endpoints-- the lowest value x0, the highest value x1. Let's say for simplicity that's all that's going on, that the lottery that creates the random variable is nothing other than some weighted average of x0 and x1. And those weights are such as to produce exactly this expected value.

And that's what's being drawn with this line, namely expectation of x is a weighted combination of x0 and x1. Given the expectation of x, we could ask what the utility would be for that deterministic value, and it would be-- we would have to go up here if they were getting it for sure, and over here, and have the utility of the expectation of x. That may be a funny place to start, because there's a genuine lottery here, and you can't get the mean for sure.

You have to engage in the lottery and suffer the randomness around the mean. So what is the actual utility of the lottery? Well, that can be achieved in this diagram by taking a weighted average of the utility of the extreme points. Here's the utility you would get for the high value, x1. Here's the utility you get for the low value, x0. If you take a weighted average, you'll be moving along this line.

You'll actually take this weighted average here that also determines the mean, and go over here and plot the utility of the outcome. So this is the expected utility under the actual lottery, and it's just the appropriately weighted average of the utilities of the two extremes. You're with me so far? So as you can see, if instead of suffering the lottery, you gave this decision maker a deterministic value with the same utility level as the utility has under the lottery, what would that value be? That would be right here.

If you started with this value of x for sure, and went up to the orange line, you would get the utility there. But that value of x is obviously a lot lower than the mean. This is called the certainty equivalent value of the lottery. Let me put it a bit differently. If you told the household to choose between the lottery and a deterministic value with a mean less than the mean of the lottery, they might choose the deterministic value as long as that value is not less than the risk premium.

Anything here, they would choose. Anything over here, they would not. Right here, they would be indifferent. Now, as I think you can imagine, if you increase the curvature of this function, you're going to actually be lowering the expected utility of the lottery, because there's more risk associated with it, in the sense that you care more about the variance because your utility is more curved. And that's going to change the certainty equivalent amount as well.

So from now on, for much of the class, we're going to be considering cardinal rather than ordinal utility. What are standard measures of risk aversion? There's something called absolute risk aversion and another thing called relative risk aversion. Absolute risk aversion has to do with how your marginal utility is changing as you change the reward, the consumption. And relative risk aversion has to do with how much utility is changing as you change the percentage of wealth that you have.

I'll spare you the derivations of all of this, but I will show you some example utility functions that display either absolute risk aversion constant or relative risk aversion on the next slide. So absolute risk aversion is your willingness or not to take up gambles, like, starting with $100, add plus or minus 10. The lottery is from 110 to 90. The relative risk aversion has to do with proportional gambles. You start with, say, 100 and then have 10% gains or losses.

That would be the same, but you could start with 1,000 and have 10% gains and losses, in which case you'd be facing a lottery that moves you from 1,100 to 900. Under constant relative risk aversion, all you care about are those percentage differences from your initial wealth. Under absolute risk aversion, you actually don't like that new-- if you started with $1,000 as opposed to $100, your absolute risk aversion would actually be going down.

Because of the curved utility, you would be thinking about a lottery around a point that has less curvature. Attitudes toward risk do depend on these functions. So here are the examples. This one is constant absolute risk aversion. And it's an exponential function-- looks kind of weird at first because of all these negative signs. Say, well, exponential is positive-- yes, so as C goes up, utility is going down. What the heck's going on here?

But we put a negative sign in front of all of this to correct for that. So you could take the derivative and the second derivative and verify that this function has constant absolute risk aversion, and the order of magnitude is actually a. Or if you started down here, we'd have constant relative risk aversion, and you can verify if you did the math-- u prime u double prime times C that you would get a number R.

So R is the constant level of relative risk aversion under this function. And we often use these for different purposes. They'll show up in an application we'll get to in about, I guess about four lectures. Let's talk some more about lotteries and uncertainty. Actually, we'll now throw in not just uncertainty as in states of the world S, but also time. So time is evolving from day 1 to day 2 to day 3.

So from the perspective of your initial starting point, i0, you could experience a low or a high state at date one. And if you're at date one, you're still facing risk, so you could go low or high from there. If you went to medium, then you could go and turn all over again from low, medium, to high. These branches of the trees with the nodes reflect every possible realization and the path to get there.

From the standpoint of expected utility, we would start here, ex ante, before anything happens at all, and trace the probabilities at date one. You get this or this or this with certain probabilities, pi s1, given s0. But you'd also take into account day two, so you take the expectation of being here or here, or any of these intermediate branches. And finally, you'd have to add in the day three, where all these events at the thinnest branches of the tree could come to pass.

If there's anything counterintuitive about this diagram, it is that when some things happen, other things don't happen. So when you're over here and move left, you've eliminated the other two branches of the tree. If you took expected utility from day one onward, you would just be taking into account what's left that could potentially happen at random. But this criterion down here is taking expected utility as if nothing yet has happened at all.

So taking into account all the possible events at day 1, at day 2, and day 3. So we can immediately get to a building block that's crucial to understand finance and financial markets, and that's the notion of a state contingent commodity. We've already been doing it. The state contingent commodity with only one date would be the probability of getting a certain wealth payoff or certain xy combination of goods and so on.

So what you get depends on the state, commodity bundle you get depends on the state. We call it state contingent. You get wheat tomorrow if it rains, but nothing otherwise. And we imagine you can actually buy these kinds of contingent pay off securities in the financial markets. However, a much better way to put that is any financial security is a weighted average, probabilistic weighted average over all possible events.

For example, you're an investor and you bought in-- someone's lending to you. And let's imagine they are never going to default. Then, you're going to get that money back for sure, in which case, it's not risky anymore, but that's only because the event is the same. Or you could have something like a stock option, so that it doesn't come in to play unless the stock prices are particularly high and you execute the option.

So all possible securities can be represented as linear combinations of these primitive Arrow-Debreu securities-- primitive in the sense that they, building blocks, only pay off in the event of one well delineated event, and 0 otherwise. So when we get to competitive markets and how to handle uncertainty, we'll come back to these state contingent commodities. OK, so let's move to an application of various of these things, one lotteries, two, random variables, three, diversification of portfolios.

Except in this case, they're not going to be trading financial assets, they're going to be holding plots of land. So the next map I'm going to show you is a village of Elford in Stratfordshire, England, in 1719, before enclosures. And shaded in black are the holdings of Mr. Darlaston. So here's that village, Elford, Stratfordshire. And you have to look a little close, but you see these darker areas, they are shaded in and represent the holdings of Mr. Darlaston.

Mr. Darlaston has these long, narrow strips of land all over the village, in all the three fields. And you could well think about that as a portfolio of land holdings, and in some respects a portfolio of securities where in this case, the holding of land yields a dividend, which is the harvest. And you may well surmise there's some risk here, or they're weird people, because why else would they be doing something like this?

Here's another view of the village of Laxton. I should at least tell you what enclosures are. Given the way this land was spread out, it was impractical to put fences up that would mark the division of the boundary of all these various strips. It would have been very costly to put up the fences and very difficult to plow. See the little plows here? Eventually, enclosures meant they put-- they divided it up. They got rid of the area and divided up the land in a different way, just put fences around the consolidated pieces of land.

But they didn't do that for hundreds of years. They kept this, the scattered strips, for quite a long time. OK, so let's move to the data. The data we have from this period of time is from the Bishop of Winchester. The church was wealthy and they kept records. So this is what survives, the holdings of the Bishop of Winchester, all domains, the crops being wheat, barley, and oats, various types of grains. And these objects denote, say, the coefficient of variation, which I've alerted you is the variance divided by the mean, standard deviation divided by the mean.

So this number, say 0.35, is in the middle of the two others. Actually, it's a bit on the lower end. Even at 0.35, the coefficient of variation of that magnitude is a lot of risk. So let's look at the picture. If we have those scattered plots, those strips, we could plot the entire distribution of outcomes with their probabilities. We could denote the mean, Us, for scattered plots. And then we have disastrous events, so d represents minimal, subsistence consumption.

And if you have these scattered plots with that coefficient of variation of 0.35, then roughly one in every 12 years you get a very low yield, less than 2 times-- 2 standard deviations below the mean, and that's starvation. And they were starving, we know this, periodically, every 12 years or so. Not like clockwork, but on average every 12 years. Now, could they have done something different? They could have consolidated the land and that would have gotten rid of certain costs and problems.

That would have moved up the mean, but you'd be left with more mass in the tails, and hence a higher probability of having a yield that was disastrous. So the reason that they're spreading out the land like that despite costs has to do with trying to increase their expected utility by putting less mass in the tails. Also from the Bishop of Winchester estates, which you can hardly see here-- this, believe it or not, is Southern England.

This is the water down here. And each one of these dots represents an ecclesiastical estate. And we have the data of the yields, the harvest yields. In this case, the covariate, the correlation-- the statistical concept we were chatting about earlier-- is what is being plotted here for villages that are relatively near to one another, and villages that are relatively far away. So D means near-- sorry, it's hard to read the legend.

And so take two villages pairwise, look at the correlation across the two villages. It's 0.68 on average, picking villages that are near. And if you have villages that are far away one another, that correlation drops to something lower, so the yields are not moving as much together. So again, if you were the Bishop of Winchester, maybe you chose the location of your estates in order to try to ensure yourself of having somewhat regular grain for harvest, for eating, and to feed the horses, and so on.

Horses are serious, by the way. That's the cavalry. That's the main defense, so they need it to keep them alive. OK, then we come to a technique which we've already been using, but let's do it from scratch. And that's called linear programming. So this is a linear programming problem. X is the control variable, a vector. This-- say, column vector-- these constants, coefficient C, represent a row vector, so these are weights on the various x values.

This is a set of constraints, which is this matrix a times x, less than or equal to b. And the x's are non-negative. So this is notation for a standard linear program, namely the objective function. And the constraint sets just represent linear combinations of the control variables x. I had this in the back of my head when we were talking about the Leontief input-output matrix.

And I said in words, instead of taking final demand D as given, maybe it's a weighted average of demands that you're trying to maximize, and you actually have a choice of various production modes to maximize that weighted average. So we actually saw an early version of a linear programming problem, but the objective function was not made explicit. This would be the consumer choice problem in a competitive environment, so where there's a budget constraint.

So these discrete points, commodity points C, denote subscript S varying over error. S, states of the world was the expected utility criterion, so that's the objective function. By choice of lotteries, so the lotteries are non-negative and add up to 1. To qualify as a lottery, any given lottery has to satisfy those two criteria. The choice object, as I said-- to belabor it-- is not c of s. Those are fixed commodity points. The choice object are the pis.

And we'll talk about a budget problem, which is there are prices p depending on s-- that is to say, depending on the underlying discrete bundle, quite naturally. And you've got a budget in dollars that you can use to buy a lottery, but if the probability is non-trivial, it's a number less than 1 and greater than 0, then you scale the price of getting that bundle s back by pi.

In other words, you're kind of paying the expected price. If pi s were one, then you're buying cs for sure and the price would be ps, and that would look like the standard price times quantity cannot exceed the budget. But here, what you're buying is a lottery, and the lottery can have non degenerate elements. You have to decide how you would be paying for, what you would be paying if you were buying a lottery which were non degenerate.

And here, you pay as per the actuarial value of the payoff. Well, this is another linear program. The objective function is linear in the pis. The constraint set is in the pis, including the budget set, which is linear in the pis. So we can solve these things conceptually and numerically, recognizing that it's just a version of the general notation. And here's the third. Suppose we have a firm-- back to production again.

The firm can produce n possible goods. The jth good here has a price, pj. The firm has a fixed supply of inputs, which makes the problem special. They're not buying these, they already have them in inventory. And the technology tells us how much of input i is needed to produce good j. Those are the aijs. Now here, the natural objective function would be the revenue, price times quantity, summing over all the j goods.

And you're going to choose the qjs, but you can't over utilize your inputs. So if the firm is endowed with input i, then you would say, how much of input i is needed to produce good j times the proposed quantity of good j, summing over all goods j. This would be the total input requirement for input i, associated with producing this vector of goods q, indexed by j. And so this is a linear constraint. The qs are non-negative. This is another linear program.

And the last one is what's called a transportation problem-- maybe just say it in words. There are m warehouses, which are stocked with goods, and n retail outlets where goods are sold. So here, we have geography. And in particular, in order to ship the product from a warehouse to an outlet, you use up resources. So this cost, cij, is the amount of resources you used up, say, in dollars, to ship the good from warehouse i to outlet j.

And then the decision problem is you're given demands at all the outlets, and you're given the amount, initially, in the warehouses, and you want to then determine the shipments from all the different warehouses to all the different potential outlets to solve the problem, which minimizes the total cost-- cost from i to j, warehouses i to outlets j, and these are the constraints.

The constraint says that you've got so much in inventory at warehouse i, so whatever you're shipping out of i to all these outlets j, it cannot exceed that level of inventory, although it could be less. And likewise, you have to meet the demand at outlet j-- So the shipping that's coming from all possible outlets i to that particular j cannot be less than the demand.

So this is another geographic programming problem. So we did ordinary production. We did choice under uncertainty. We did a spatial geographic problem. They're all well known prototypes in the literature, and they all use this common tool, which is linear programming. OK, that's all I have for today.