[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: OK, we'll get started on time. If you look at the reading list and scroll down to where we are, which is Dynamics and Programming, you'll see Lecture 6 there. There's a chapter on Varians, which is starred, called Time.

That's because the application today will be featuring dynamics and what happens over various dates. And we're going to do it in the context of a storage problem, where people were close to starving to death, as in the medieval villages. And that is Chapter 3 of "Medville." It's also starred, which fills out even more than I'm able to in one CLASS the material for the lecture today. So that's the reading list.

Then there's study guide. So I went back and earmarked what we did last time, which was Decision Making under Uncertainty, Linear Programming. Some of these questions are worded as if you were doing it at home. But nevertheless, the question asks about important concepts.

So let me read the first one. Draw a consumption set which consists of discrete number of points argue that th set is not convex and then show how lotteries which put mass on these discrete points, transforms the consumption set into something convex.

So there's two sentences there, effectively two questions. Could I get a volunteer, please? So you can't actually draw it. Can you describe a consumption set?

Maybe these words suffice. Try to describe, in words, a consumption set which contains a discrete number of points and tell me why it's not convex.

AUDIENCE: Yeah. So I think we discussed this earlier. But one consumption set was consuming lunch in New York at noon versus consuming lunch in Washington, DC at noon.

ROBERT TOWNSEND: That's true.

AUDIENCE: So that would be a discrete consumption set. And since you can't consume lunch in a mixture of those two cities, then the consumption is not convex. But if you had a sum probability of eating lunch in either location, then you could have a probability between zero and one.

ROBERT TOWNSEND: And how would you draw it?

AUDIENCE: So for the original discrete consumption set, you'd just be two points. I think one of those points would be one, zero and one of those points will be zero, one.

ROBERT TOWNSEND: Good. And can you imagine how-- I'll help in a second if it's not clear-- can you imagine how you would be able to convert that set into something that is convex, using the lottery? Is there something you could draw to depict it?

AUDIENCE: If you just sort of drew a line between those two points, like your location, instead of doing how many lunches you eat at DC and New York, you could do the probability that you eat at DC or New York.

ROBERT TOWNSEND: Exactly. So now we have a line running on the diagonal from one, zero to zero one. The end points are still there. Those are lotteries that are degenerate. They're putting probability one on each of the endpoints.

An intermediate lottery, a non-trivial lottery, say, halfway between, would put probability 1/2 on-- would be distance 1/2 to one, zero and distance 1/4 to zero, one. And that represents an equal weighted probability. And you can inch closer to one endpoint or the other. Those are still non-degenerate lotteries.

So a lottery, when you put a point in there, represents the weights. But yes, your answer was perfect. Thank you. Questions about that?

OK. What does risk aversion imply about the consumer's utility function? Again, I'll take volunteers.

AUDIENCE: If you're risk averse, then you have diminishing marginal utility of income. So the utility function should be concave.

ROBERT TOWNSEND: Good. Strictly or weakly?

AUDIENCE: Convex?

ROBERT TOWNSEND: Let me repeat. The question says what does risk aversion imply about the consumer's utility function. So your answer, concave, yes. And I'm asking back, you mean strictly concave, weakly concave.

AUDIENCE: I mean, I think it would be weakly concave, right? Because if you have a straight line-- or actually would that still be considered risk averse, if your utility from getting a certain value is like completely equal to the utility of having that as your expected value?

ROBERT TOWNSEND: All right. That's perfect. You figured it out.

So the answer is we typically think of risk aversion to mean strictly averse to the lottery. And hence, the linear case of the linear utility function, or weekly concave function, but not strictly concave, we call them risk neutral and not risk averse.

But I never said that in class. So I'm amplifying. But yes, risk aversion is being risk averse. Strictly speaking implies strictly concave utility. So perfect, thank you.

So one more here. This, again, is going to require you to describe something. Combining states and time of the world, time and states of the world, what is the complete commodity space for a consumer? And what I'm really looking for is for you to describe to me the tree of possibilities that we had on the slide last time. And, again, I'll take volunteers.

AUDIENCE: So this is a tree that's oriented upwards. And so at where time increases as you go up in this tree. As you make a specific choice at any given time, you limit your future options to some specific set, and you exclude some others that are no longer possible for you to consume.

ROBERT TOWNSEND: That's right. So the tricky part of it is to remember that the state of the world at any date is not just a set of possible realizations of the contemporary shock. It also includes the history of shocks, up to and including the contemporary date.

So that's what's going on with that tree. You pick a node in the tree. It's described not just by the set of all things that could have happened on the current branch, but the set of all things that could possibly have happened up to and including that point.

And actually, I made the comment in class that one has to consider whether you're already at the current date, or whether you're at date zero, imagining what could happen in the future. Because if you're at zero and imagining what could happen in the future, then you have to delineate all possible histories and all possible states that can happen in any given date, including up through and including the date in question.

Great. OK. So, again, I earmarked some of these. I only went through three of the four. That was very helpful. Thank you.

OK, so the lecture today is-- so this is on dynamics and dynamic programming. It's going to focus on this application for medieval villages on storage, or call it inventory, or carryover, as well as the seed that they put in the ground. And I'll come back to that momentarily.

And if you go back to the economic science lecture, this lecture is illustrative of something we call calibration. So you have a model of the way the world works and you find, somehow, parameter values that are appropriate for the model. And then you simulate the model and figure out how well it fits the data. And if it fits reasonably well, then it's a success, that you have a logical explanation of in an artificial economy of how the world actually works, or in this case, past tense, was working.

There should be another bullet point here. Because I like to feature and earmark that when it comes up, we go through various techniques. And the technique that's going to become apparent is dynamic programming.

So most of the lectures up to this point, not all of them were static. We made an exception with that interest rate and income and substitution effects and a couple of the consumption set diagrams. So time is playing a big role here. And so if we're going to solve a problem or program, it's going to be a dynamic program. And there's some useful techniques that fall right out of the end of this lecture.

So we're going to talk about actual and potential use of risk reduction arrangements, namely carrying over grain from one year to the next. This first line says alternative, begging the issue of what else have we looked at. Well, we looked at an individual putting his strips of land all over the village. And I referred to that as portfolio diversification at one point.

We looked at, and we are about to see, again, the picture of the estates of the Bishop of Winchester, which were spread out over space, arguably, again, to diversify. And now we're going to be stuck in one village, and effectively, with one person in one village, on his or her own, looking for alternative ways of reducing exposure to risk.

Now the stylized facts, if we go back to the same Bishop of Winchester accounts, we don't see a whole lot of carryover. There are one estate at the time, rare spikes where a carryover jumps up. And then the next year, it's down.

I should say, or this comes up later, "carry over" means carried over from one harvest to the next. Obviously, time passes by the day and by the month. They only produce one crop a year. They have to have something to eat during the interim.

So there is some storage going on. The issue is whether, by the time they get to the next harvest, there's any storage left at all. And if there is, we call it carryover.

Well, likewise, if you look at the same estate accounts, at a point in time, you can see that some of them have carryover. But most of them don't.

So this model is going to be similar to a macro model. That probably doesn't mean much to you right now, a neoclassical growth stabilization model. We're going to put in two kinds of storage technologies, grain in the bin for storage, but also seed in the ground, because that's another way to carry over the crop, by planting it. And we're going to choose the parameters to be consistent with the available data. And we're going to, again, look at the numerical predictions of the model and see if we can match the fact that carryover was not common.

This goes more into how uncommon it was where, if they're not carrying anything over, and there's only one person in a village, like Robinson Crusoe as an abstraction, then they're eating all their output-- well, maybe-- over a 12-month time period. So variability in consumption is one to one with variability in output. And if output varies a lot, that's a lot of risk to bear.

We have data from that same lord. Treating him as the single villager, we can look at output as representative output in the population and see, in the cross section in 1236, only three of the 18 states had carryover.

In a good year, there was carryover. 1223 was good, for example, for one of those estates with 57% of the crop being stored till next year. And the following year for that same estate, 1224, zero carryover.

This is a review picture. These are the estates of the Bishop of Winchester. And as I said last time, it's hard to figure out what country you're in, let alone where you are in it.

But this is the English Channel down here. So all these dots are the Bishop's of Winchester estates. So that's the context that's generating the data.

OK, two ways to store, well, actually, four. We're going to adopt two of them and rule out the two others. Not model the two others, I should say.

In the artificial economy, first, grain can be stored after harvest, as I've been saying. Second, they could take some of that harvest and put the grain in the ground as seed planted for next year's crop. The other things we're not going to consider is taking the grain and feeding it to the animals. And eventually, you might kill the animal.

So in some respects, you're carrying over the calories from the grain. But it's not very efficient, since the animals are also surviving on the grain. So we don't have animals in the model.

Another more humorous possibility is beer. So they would take the grain and ferment it and make barley malt, for example, which English people still like to drink quite a bit, no doubt due to this history. But we're not going to do that either, no drinking in the model.

So let's take this storage possibility. The most obvious one is you take the grain and put it in the bin, and potentially have some left or not by next harvest. How the depreciation rate being delta, why does grain go bad?

Well, in England, it's a very wet climate. If it gets wet, and the grain gets wet if it's not secured, it spoils. Mildew, even if it doesn't directly rain on the grain, is not good for the grain.

And grain gets eaten by rats and other rodents. And you may think that's humorous. But it is real. I know in Thailand, we would go and look at what they call the rice bank, which was a shed where they had stored the rice. And they were showing us the shed. And as they opened the door, all these mice ran out.

So we're eating the grain. How big is delta? We have to calibrate this model.

So drawing from McCloskey, you can see prices increasing month after month after month, from harvest up through the next month. And from this fact, McCloskey gets a depreciation rate of 30%

So here's that table that he's using. Between September and October, grain prices inflated by 3 and 1/2%, between October and November, 2 and. 1/4%. So these are kind of grain price inflation rates through the year. And roughly, if you add them up and do a little compounding, you'll get 30%.

Now who cares? Well, let's imagine that whatever you store at one date, you get back 1 minus delta that number, say, the following month. It's linear. You don't get out what you put in 1 to 1.

If the depreciation were 30%, you put in 1, you get back 0.7. You put in 10, you get back 7. It's linear. The storage technology displays constant returns to scale. If you were a firm maximizing profits by the choice of how much to store over time, you would be thinking about this constant returns to scale technology and asking, how much would they be storing, depending on how much the price next month is relative to the price today.

So it says exercise. So I urge you to write that down. I did it just to make sure this morning. It's, like, maximize P1 times y minus i, where i is investment, plus P2 times Y2 plus i multiplied by 1 minus delta. It's kind of silly to be describing this in words.

Now what do you know about constant returns to scale? You already know, with constant returns to scale when we did the Production Lecture 4, that the solution is either infinite or indeterminate or zero. Well, it's not indeterminate. And it's not zero, either, because they were storing from month to month, at least.

So the premise is that, in order to compensate the economics is, in order to compensate for the loss of the crop, rather than selling it today at a certain price, they can sell it tomorrow they're going to have 70% left, a 30% loss. So the price has to accumulate at 30% over the year.

I'm going back and forth between month and years. Sorry.

So that's the economic intuition. And you already know the math. So that's one advantage for us having spent some time on production and, in particular, on profit maximization. So you know how McCloskey is calibrating parameter delta.

Second possibility is putting seed in the ground. Now it turns out that's a pretty healthy return. We have some data on that. And the yield to seed ratio is about 2.6 for one of these manners and 1.67 for the other one.

So that's a yield-to-seed ratio, meaning you get roughly, taking the average, two units of output for every units of input. So that's pretty productive. So if we stopped right there, you'd say, well, putting seed in the ground is a lot more productive than depreciation, even setting aside the risk.

Well, so there's a couple of things. One, there's a limit to how much seed you can put in the ground. And two, we're going to be considering the risk.

So here's the technology. If you planted Kt units-- put Kt seeds in the ground, at t minus 1, the output would be a function f of that amount of seed, the harvest. And there's some shock epsilon.

In fact, let's assume this is a linear technology. So whatever seed got put in the ground, alpha times it is the output. Alpha could be 2. And epsilon is a shock that's going to make it random and make it move around.

OK, first thing, you can't plant an unbounded amount of seed. In fact, the best way to think about this is the seed per unit land is constant. You have so much land, there's no point in putting down more seed.

So it's really a decision about how much land to plant. And we'll put an upper bound on that because there are finite land holdings.

So K bar is going to be the most land that you can plant. You plant all of your land. And if you were to try to put down more seed, it's like throwing it away. You're still going to only have the maximum amount, K bar in this case times epsilon alpha.

So we're about ready to get to a key equation here. Let's do it this way. If you focus on equation 1, you've got output f, harvest, which is a function of the seed planted last year plus an influence of epsilon. Now you could have had something in storage as well. Imagine they had done both things.

So they put some seed in the ground. And they had some storage. You're going to get 1 minus delta times that amount of storage back. And that would be the end of the story for consumption, except they have to plan for next year. So they still have these two options, storage planted now at t that will be available at t plus 1, and seed, which they can squirrel away now, which will be planted and influence the harvest at Kt plus 1.

It's kind of weird dating. Please be alert to that. In fact, the last slide for today will change the dating. Different people use different conventions.

What I'm complaining about, or alerting you, is this is consumption at date t. And these things, seed in the ground and inventory, have a date t on them. But they are predetermined from t minus 1. They are not decision variables at t. The decision variables at t are labeled t plus 1. The pro case for doing it this way is that we care about the state variable at a given day t, and that's determined by the seed that they planted before and the inventory that they have from before.

Anyway, just to alert you. I trust you will interrupt me when you have questions.

So this is the amount that's available. This is the amount that you squirrel away for next year. And the difference between what's available and what you store in one way or another is consumption.

Now the risk averse, we discussed that a minute ago, so this is a strictly concave utility function over consumption over a finite number of dates t. So this runs from 1 to capital T. Hopefully, capital T is really a long time away. And I'll come back to that.

Alternatively, I could have written cap T equal to infinity. They live forever. But we're not quite ready to do that. So I'm just going to show you finite time horizon first and then argue how to extend it at the end of the class when we get to dynamic programming.

There's a discounted utility. Beta is a discount rate, beta some number less than 1. So you care about the present more than you care about the future, present utility more than future utility. This is, with beta less than 1, geometric. You can see, as the horizon gets further and further away, you care less and less relative to today's consumption.

But these are strictly concave, so you don't like variability. We talked about concave utility when we introduced the concept of risk. Here, even if there were no risks, they would try to avoid variability in consumption.

They'd rather have the same consumption every day, rather than having it low one date and high another. Because the utility consequence of that would be less than having the same amount. So that's where the context drift concavity is coming in.

So we already calibrated the depreciation rate at 30%. The discount rate we're going to put at 0.95. Which is kind of standard, although we could vary it, we could make it lower. Once we have the program written, we can plug ir whatever we want. Ideally, though, the spirit of this should be that we're somewhat confident about these numbers we're about to choose.

Constant relative risk aversion, I defined that utility function last time. That's seed to the gamma divided by sometimes normalized. And we're going to set that at 0.5. This is kind of a low number. It's definitely a strictly concave utility, but they may have been even more risk averse than that.

However, you'll see this ex post justification that all three of these numbers kind of work in terms of explaining the paradox. Namely, they're about to starve to death every 12 years. And they weren't storing anything. That's the paradox.

Back to that epsilon, this is the way you put the uncertainty in. One unit of land planted today yields two tomorrow. So that's the yield-to-seed ratio.

But there is a little uncertainty. 10 out of 12 times, you get 2 tomorrow. But 1/12 of the time, you get 1/2 of it. And 1/12 of the time, you get 4, double it.

So the expected yield is a little bit above 2, since the weights, 1/12 are same on the low and the high end. But 4 is 2 times greater than 2. And the difference between 2 and 1 is just 1. So it's just the expected value is slightly higher than 2.

We're going to set the upper bound on storage at 1. And this number, 1/12, I actually just set y. And you may be reminded of that when you go back and look at, in the data, the coefficient of variation that we introduced last time of yields. It was 0.35, which meant roughly every 12 years, output fell below less than half of the mean value.

So we called that disaster last time for these scattered plots. So this mass over here under the solid density is 1/12. [INAUDIBLE] it's also true they could get lucky and have a high yield.

So you saw this picture last time. I'm really just using it to remind you of the context of where the risk is coming from. We wanted-- we just did calibrate the yield-to-seed ratio in such a way as to deliver disaster every 12 years, on average.

OK, so let's look at the dynamic decision problem. Let's call this the total available this year and denote, on the y-axis so to speak, the amount of each of the possible savings decisions, either of storage or planting, back to the key equation. The amount available is this, output plus storage. And the savings decisions are inventory or seed. So that's what's being plotted.

But we are varying the total amount available going from 0 to 6. And let's look at-- oh, where is this coming from? So we actually solved this problem.

We maximize 2 subject to 1, putting in all these parameters. And I didn't show you the code. You can solve it on Matlab or something. I'm not asking you to do that.

But anyway, it's, like, this picture could be anything. It's pinned down, because it's the solution to a max problem at the calibrated parameters. So this is our model version of what they were doing, in fact.

So if the total amount available was a little less than 1, then they were not planting all the land, despite that high yield. However, for anything else, they hit the upper bound, and they plant everything they can at the upper bound, which was 1.

Let's look at the storage decision. As a function of the total available, it's only when the total available is 3 or higher that positive storage kicks in. And when it does, it's moving linearly with the amount available this year.

It's actually not moving 1 to 1. If it were 1 to 1, higher available means higher storage, it would be parallel with this 45 degree line. But it's not.

So they never store 1 to 1. They may store less than the harvest. But they are storing some, something here.

All right. So those are the savings plans. Things to keep in the back of your mind when we go through some of the other slides is the fact that mostly they put all the seed in the ground, but not when the amount available is less and the amount of storage is zero, unless the amount available is above 3.

Now how much is consumption? Well, back to this thing, so consumption is the difference between the amount available and those two savings decisions. So if you look at this diagram, consumption is the difference between the amount available and the sum of these decisions.

This 45-degree line is now playing a role. Because the amount available, say, 2, go up to the 45-degree line. Then this amount is 2.

That's the definition of the 45-degree line. Right? It's a way to convert the units on the x-axis to the same amount on the y-axis. That's convenient here. Because we can look at, OK, here's the amount that would be available next year if they didn't save anything. But they are storing in this case.

So consumption is the difference between this 45-degree line and the sum of these two other lines. That means that as we get over here, the 45-degree line is really close to the only thing they're doing, which is planting seed, and very close. So when you have less than 1.1 available this year, they're still eating. And they're not planting all the land, but they're not eating much.

It's a really hard problem. Do we starve now or starve later? I mean, they kind of gamble on getting a reasonable return. If they ate the seed, they have nothing tomorrow. They're starving to death for sure.

In fact, I don't if you've heard this expression, don't eat the seed corn. Yeah, that came right out of this episode. Eating the seed corn is the worst thing you could do. Because you're kind of guaranteeing zero for tomorrow. But they actually do cut back. They eat some of it.

OK, so this is a related way to think about the dynamics. Let's take total available this year, implement those two savings decisions on the previous diagram, and then ask how much are you going to have next year. Sorry if I'm belaboring this.

This is the amount available this year. You can save it or eat it. If you save it in one or the other of these two devices, we're now talking about the first part of this equation with the two terms. But the date would be t plus 1.

So this would be the amount available next year, given the amount available this year and the endogenously determined-- oh, I used that word for the first time-- endogenously determined savings decisions. So here, we can't plot savings and the total amount available unless we're in some 3-dimensional space. So the savings decisions are suppressed, though they're still active.

Well, first of all, output is stochastic. Epsilon could be and likely to be medium, but it could be high or it could be low. So the amount that's available next year, even given the savings decisions, will vary with the epsilon.

So if you had three units available this year, which is right on the margin of storing, say zero, the most likely outcome next year is medium. And hence, there would be two units available next year. And so there would be no storage next year, either.

So what do those words correspond to? You're at three units available now. The most likely outcome next year is medium. The amount available next year is two. And, again, from the savings decisions, storage at three and at two is zero, so no storage.

So this is like what would happen over time. But let's just come back. Suppose we redid this experiment, and you had two units available this year, and then you got the medium output. You would have two units.

So what does that referring to? If we have two units available this year, and you had the medium output, which is the most likely thing, you get 2 units available for the following year. So two in and two out is the most likely outcome if you start with two.

So you could refer to that as a kind of steady state, meaning things aren't moving. Actually, they are moving. So technically, it's not a steady state, because we picked a particular shock.

But it is the most likely shock. So it's approximately true that if you have two today, you're going to have two tomorrow. But if you had two units available this year, and you got the low output, epsilon 1/12 being probability, the output available next year is about one. And at that point, you're not planting all the land.

So that's over here. You could have two today and get really unlucky and have less than two available, about one available next year. And, again, if you memorized a bit, those savings diagrams, with one available, you're not storing. And you're eating some of the seed.

So let's go the other way. Suppose you had two units available, which is, roughly speaking, a reasonable starting point, and you got lucky and got the high output. Then you would finally be storing.

Suppose you had two, and instead of getting the medium yield, you got the high yield. Then you'd have four units available next year. And at four units available, there's definitely storage going on.

So that's carryover. Rare, but it can happen. Does it persist? Well, suppose, just to run this experiment, that you get lucky and lucky and lucky.

So where would you end up? You'd end up at roughly a little above four. Four replicates itself, four in, four out.

It's on the 45-degree line. So that's the best case scenario, that you would converge to that. And you would be storing.

But I'm sure you can see this coming. If you were at that fortunate situation with storage, you got a medium or low output. Then you would climb back below three. And you wouldn't be storing at all.

So they had carryover once in a while, but not often. Questions?

AUDIENCE: I have a question.

ROBERT TOWNSEND: Yes?

AUDIENCE: If the depreciation rate goes down, then the inventory line will move, right?

ROBERT TOWNSEND: It will move up, yeah. It would shift up. Well, it would not only have a higher slope, but it would kick in earlier.

AUDIENCE: Yes. So it is also implied that in the medieval England maybe the depreciation rate was so high, so the carryover was so low, or we cannot observe the carryover.

ROBERT TOWNSEND: Say that again?

AUDIENCE: I mean, maybe it implies that in the medieval England, the depreciation rate is so high, that is why we cannot observe the carryover.

ROBERT TOWNSEND: Yeah. Well, if--

AUDIENCE: Yeah.

ROBERT TOWNSEND: I mean, we picked it to try to match the data. You could run the model with a higher or lower number. If we make the depreciation rate less severe, so it's like 20% or 30%, they would store more often.

And then we would have to look at that diagram and see whether carryover is still rare. I mean, I imagine if you go from 30% to 28% or even 25%, not too much might change. But it's clear if you ran the depreciation to zero, then they're going to be storing quite a lot. They would have a way of carrying the crop over one to one. It would still be uncertain, but that would be a big step toward keeping their consumption steady.

AUDIENCE: Thank you.

ROBERT TOWNSEND: Yeah. So no, it's a great question. Because that parameter and all the other ones are key. We could vary each, one at a time. We could make them more risk averse, for example.

AUDIENCE: Yeah.

ROBERT TOWNSEND: That would be an interesting experiment. Then they're going to be more worried about that low output. And that probably means they're going to be trying to maintain more seed and more inventory, which one I'm not quite sure.

Well, there's an upper bound on the amount of seed, the amount of land they can plant. So that's probably going to push inventory up. That's a version of what we saw last time in a static context. Namely, that you're willing to accept something less than the mean return in order to guarantee, to get rid of the risk.

So in this case, even with storage being somewhat costly, when they're risk averse, then we could make them even more risk averse. They're going to be more inclined to give up some of the consumption today, even though it's going to not survive the rats into tomorrow. Because when they have it, the consumption today and consumption tomorrow are more likely to be the same.

AUDIENCE: Thanks.

ROBERT TOWNSEND: Beta matters. If beta goes to zero, you don't care about the future. So obviously, these guys would be very myopic. And they would be behaving very differently, and so on.

So all these parameters matter. I guess I got lucky. I mean, some of these parameters were picked. And I was complaining they seem a little bit off, like beta and risk aversion. But that, along with the depreciation rate and the very carefully calibrated technology, did succeed in producing the data.

So I want to come back to this. We cheated a bit. Because we had a finite horizon and then I kind of told you to ignore it. That diagram was produced as if the horizon was infinite. And how would I know that?

Because if the horizon is finite, then unlike those figures, the amount stored is going to depend on how many years are left. Intuitively, if you're at the terminal date-- it sounds kind of foreboding, no? The terminal date? Well, you've got one more year to go.

That's it. Eat now. You're not going to be around tomorrow.

So there's not any savings in seed. There's not any storage. Now you could move back from there to the next to the last date, knowing that you're not going to save tomorrow. You're still facing a choice between today and tomorrow, or this year versus next year.

So you could solve that two period problem. We call this working back from the end, right? We start at the last date. We move to the next to the last date and solve the maximization problem. Substitute in the values, and then we have the maximized utility at the next to last date because we've already solved that problem.

And then you move to the third to last date. But it's still going to look like a two period problem because we've collapsed the last date and the next to the last date into one because we already solved that last two period problem.

I'm going to show you some notation. I'm not sure you could possibly be following me. But I'll say it in English first. And then we'll do it in the notation.

Anyway, there's a sequence of two period problems. As you let the horizon t get larger and larger, you're kind of farther and farther away as you go back toward t equal 1. So the larger is that distance between t equal 1 and cap T. Especially when beta is close to zero, it's as if you're solving an infinite horizon problem.

And that is the way numerically to solve it, literally. Start with the last day. Do the two period optimization. So the next to the last to the last, plug in the solution, move to the third to last, and so on.

And then keep iterating forever. Well, keep iterating until the solution isn't changing anymore. And then, amazingly, you will be guaranteed to have solved the infinite horizon problem.

So let me show you the math of that. And I apologize the notation changes a bit. This is from engineering or economics. It's called a control problem, the Bellman problem.

So x is the current state, which in our application was the current harvest plus the storage. And y are the control variables, namely, how much you can store and save for tomorrow. That you get to choose. Also, you choose consumption. You can choose two out of three, that determines the third.

And then you'll get beta times the utility for next year, which in this control problem language is this so-called value function which has in it why the endogenous variable chosen today becomes the state of the world for next period.

So this is the two period problem. Utility contemporary, utility today, and utility next period, it's really already maximized utility from tomorrow on. But think of it as a two period problem. That's the way I was describing [INAUDIBLE].

And what is the overall solution? It has to do with J, the guess about the next to last date, and so on.

But T now, here's my apology. T is not time. T is the mapping in this language of Bellman. And I realized only this morning that it might have been better to change the notation to be consistent.

But T does not refer to the time horizon. T refers to the mapping from the utility from next period on to the utility today. This is the maximized utility, where J is a given. You're surely going to have questions about this, so interrupt me.

Now there's something called a contraction mapping. And don't worry about the details here. But it turns out we have this mapping T here, and we have the underlying space being a metric m and we have some distance metric, some way to measure distance in that metric space.

Then the difference between two guesses about v times beta is an upper bound for the distance between the result. So, again, here we should think about v1 and v2 as guesses for the right-hand side. So this would be J1 and J2. And then we apply T to it, and this would be TJ instead of Tv. So that's probably the second thing to correct, is to change the v's to Js.

Anyway, the point is the distance over here on the left-hand side is not greater, typically less than the distance on the right-hand side. And so intuitively what that means is things are going to contract. Hence this is called a contraction mapping.

Or to put it more formally, if you start with an arbitrary J or an arbitrary v, and keep iterating, start here, get this, put it in. Do it again. Get this. Do it again.

As N gets large, and we keep successively applying that mapping to an initial guess value of v, we're going to converge to something called v*. And when T of v* is equal to v*, we are done. That's the infinite horizon solution.

That's like saying the J up here and the TJ is also J. So the time horizon has disappeared. So this is kind of tough stuff. I thought it would be better to show you the math rather than be mysterious about it. But let me pause for questions.

All right. We have everything we needed in that application to apply the contraction mapping. Our guess for v0 is just the solution to the last period problem. We get the optimized value, as if you were at the last date.

Substitute in the maximized utility number. And then go to the next to the last period and iterate again with that new guess for the second period utility. And just keep iterating, and we'll get to the infinite horizon solution.

So we can think about applying this technique more generally. So this is a related, but somewhat distinct problem. It's in a contemporary setting, a household trying to maximize the utility of consumption, but facing stochastic earnings, say, from year to year. And they're limited in the way that they can save.

But there is a savings account in the bank at some interest rate, or even a riskless bond, like treasuries. So it's like putting stuff in the ground today and getting stuff out tomorrow. That's like storage.

In fact, in the first lecture, I said something about these agrarian metaphors. And we just jump like that. We went from grain in the bin as kind of a storage technology to riskless bonds as another way to store money over time, hopefully with a positive yield. Actually, the real rates are negative right now. So anyway, that's a bit of a problem.

So we maximized utility. And now I have the courage to put infinite horizon here. So we're going to maximize discounted expected utility over the infinite horizon.

Beta is still less than one. This thing would converge. So this term is finite. And the choice will be over how much to eat and how much to save.

So the way to think about this is you're coming in with a realization of your stochastic income. You're having the amount of the bond that you bought last period. You get that back plus the interest rate.

That's the resources available today at T. And you can either eat that, or buy a bond and carry it over to the next period. So this looks really similar to the medieval storage problem.

A couple of comments about it, as I tried to alert you. The timing here is the date. The dating is different. So now, at t, this is no longer a state variable. This is a decision variable, bonds to becoming due next period.

Likewise, on the right-hand side, this doesn't have a t. It had a t minus 1. Because the savings decisions are dated at the time t that you make the decision, rather than the time t where the yield comes in.

And this is also standard. You just have to keep an eye out in terms of what the authors have in mind. And this starts at date zero.

So you need some initial conditions, which should have been k at t equals 0 minus 1 or k1. This is also a typo. This should have been k minus 1, not k0. And the y0 is fine. That's the contemporary yield.

So I apologize about these typos. Although it is kind of an interesting story. They're not mine. I took this slide from somebody else, and they had actually made the mistake, and it wasn't corrected in the publication. But we'll fix it.

The other thing about the typos-- I know you hate them, I do, too-- when we did that input-output matrix for Leontief last, when we did the last part of lecture before the first part of last lecture, the beginning of Lecture 5, we went back over Lecture 4, and I was saying how much of the service sector you need as an input. Produce one unit of output of services. How much do you need from manufacturing and raw materials?

It was right in the slide. The typo had been corrected. But that typo exists in Wikipedia. And it's still wrong there.

So what I'm saying is, I work very hard to eliminate the typos. But sometimes, they still persist. If you're reading carefully, though, you can spot them.

And please don't hesitate, if you find other things that you find questionable. That's the bad part. Then it becomes confusing about whether there's a typo or not. So if in doubt, ask us. And we'll try to answer questions.

AUDIENCE: Here, the kt could be [INAUDIBLE], right? in the latter constraint, and say that a kt just are larger than minus k?

ROBERT TOWNSEND: Oh, good question. So borrowing is allowed. So I referred to this as a riskless bond, as if you could save. But in fact, if savings is negative, you're borrowing.

So if you multiplied through by the negative sign, you'd have a minus kt less than or equal to k bar. So kt positive for savings minus kt would be now denoting borrowing. And borrowing has an upper bound, a lower bar.

So this formulation of the problem allows borrowing, but it's limited. So that's a great question. Thank you for asking. Is it clear?

AUDIENCE: Yes. Another question is, what is the difference between the self-insurance and full insurance? Is the self-insurance-- that means we cannot borrow, or we can borrow, but there were some constraints like this. We can borrow at a certain range, or--

ROBERT TOWNSEND: Yeah, so I actually took out slides that I used to have. Self-insurance is what we've been studying the whole lecture. Because it's, like, one village, or one person. And they're pretty much in isolation.

Here, it's still self-insurance, although there is this bond, or this borrowing possibility at a fixed interest rate, which I didn't tell you where that came from. But this is, like, partial equilibrium problem in which the agent can borrow or lend with the rest of the economy. And that's still referred to as self-insurance, although it's a bit misleading. Because one person's saving is another person's borrowing.

But anyway, we're not doing the general equilibrium here. We're taking-- the interest rate is given.

But the other thing would be, what if there were multiple people like this, and some people had high income at a given date, and other people had low income at a given date? Then we're back to thinking about whether they could make deals with each other. So we'll get there actually next time.

But the risk-sharing problem would be-- an extreme version of that is get rid of the storage altogether and the seed. Get rid of them. Not the seed, get rid of the storage. Get rid of the bond.

But imagine these guys can contract with each other, write financial contracts, or informal risk-sharing arrangements with each other so as to pool the risk that they're facing. Like in the village, they had scattered those strips. So they're trying to stabilize their starting point.

But on top of it, they could be borrowing and lending. They could be making a deal. If I'm high and you're low, I'll give you money, and likewise, vice versa. And that's not self-insurance. That's kind of dynamics of explicit risk-sharing insurance.

AUDIENCE: Thank you.

ROBERT TOWNSEND: Yeah. No, it's a great question. As I said, I used to have more slides here about that. And I thought it was a little too early in the class.

But there are clearly trade offs about-- how can I put this-- whether or not want to do self-insurance as in a dynamic problem versus whether or not want to do risk sharing, but do it repeatedly over many dates. In fact, the very first lecture, after we did economic science, we went through some example economies. And I gave you five of them.

And the very first two were one, medieval villages with these scattered strips that we've been studying now, last time and today, implicitly, or the Thai temple scheme. You remember the Thai temple scheme?

They had low land and high land. If the monsoon rains were really abundant, the low land flooded too much, and they only got crops off the high land. And when the monsoon rains were adequate or low, then the highlands failed.

And what they did was contribute their crops to the monks. Not for eating, that really threw me. Because I know monks don't-- I mean, you can feed the monks, but the monks don't feed you.

But nevertheless, they were putting their crops under the auspices of the temple with the monks and then handing it out to people that had either drought or flood. So that's an active risk-sharing scheme. And your question comes back to that, which is comparing and contrasting these two different ways to handle risk.

OK, so this is Lecture 6. There aren't quite as many slides here. But you can see there are some very important concepts. And we've now succeeded in having both time and uncertainty and dynamic programming techniques for solving that. So we'll come back and use those techniques in subsequent lectures.

Next Tuesday is a lecture on Pareto optimality, which is finally the full blown general equilibrium, although we've kind of been going back and forth and getting very close to that in the last couple of lectures. But in the lecture on Tuesday, we will fully embrace it and talk about efficiency. OK, that's all for today.