

ROBERT TOWNSEND: Welcome, everybody, today, for the next segment of the class. Actually, this is literally a continuation. It's not a sharp break between the last lecture and today. Fortunately, the lecture today has some repeat of the equations, or at least a version of those equations. Plus, I know you're all totally fresh from the exam, so you know what the last lecture was anyway. Just teasing you a little bit. Sorry.

I'll just say a few words about the calendar. So we are right here, today, Thursday, 10/1. Oh, my goodness. It is October. And we're going to do risk-sharing applications today and then also 9. Although this says risk-sharing applications, I want to feature that we'll combine it with a production even in a more direct way than we will today. So two lectures today, Thursday and Tuesday, on risk-sharing. So that's the calendar.

The second thing would be the reading list. And we are here at lecture 8. So the two readings today are risk and insurance in Village India, my paper in *Econometrica*. And we will cover that in a fair amount of detail today, although not every word, obviously. And also, chapter 2 in *Medieval Village Economy*. So that's the reading list. So today is lecture 8, and that is here.

So risk-sharing applications-- we're going to do two different economies using the same theory, although the questions to be asked, though related, are a little bit different. One is Village India using consumption and income data, which I refer to as an ex-post version of the implications of the risk-sharing model, and also, as anticipated, the medieval village economy, where we're going to talk about dividing up the land. And you've seen maps of this before, but we didn't have all of the tools we needed to really do a good job. But we do now.

OK. So for India, here's the motivation from my paper. A large part of the population in developing countries live in high-risk environments. That said, there are numerous ways to cope with it. One is to diversify the activities you're undertaking-- trying to have a portfolio, as it were-- or ex-post, after incomes get realized, to engage in financial transactions, including gifts and transfers among family and family networks. How well do any one of these ex-post or ex-ante mechanisms work?

Well, the main, quote, "insight" of the paper was that we really don't necessarily have to enumerate the mechanism. We just look at the outcome. I'll qualify that a little bit, and we can keep coming back to this. If they diversified perfectly and everyone had a balanced portfolio, then they would all have the same outcome, and there would be no further risk to share.

So in some sense, what we're going to go through today would not have much content in the sense that today, we're referring to risk sharing, but it's not done, per se, by sharing risk after the fact. It's done by ex-ante diversification. Anyway, those two strands to repeat are going to be in today's lecture. We're going to do the ex-post part first, which assumes partial but not complete diversification, and then look at how they could do it ex-ante with partial but potentially not complete diversification.

All right. So characteristics of village economies as in India-- very risky place, in large part because agriculture does not provide a stable source of income. There's a lot of risk, idiosyncratic and aggregate risk. Households diversify in many ways. One-- not planting all of the same crops.

And we are going to use this benchmark which assumes there's no information problem, no moral hazard problem, no contract enforcement problem. All of those words at this point in the class don't make a whole lot of sense because we haven't studied anything but these frictionless economies. But that was the premise.

And again, to anticipate subsequent lectures, today we'll see how well this theory does without any of these frictions. And then subsequent literature has started to assume things like contract enforcement and information asymmetries, but we are not going to do that today.

So how do they diversify? Well, in a village Aurepalle, which was under the auspices of the Institute for Crops in the Semi-Arid Tropics, ICRISAT, farmers are a poly-crop sorghum. They grow castor. Sorghum is a grain like wheat. Castor you may have heard of as in castor oil. It's both used in diet, but also, it turns out to be a lubricant for jet engines. Anyway, the coefficients of variation are 0.5 and about 1.

Now, you may not remember that in medieval villages, we talked about a very risky environment with the coefficient of variation of yields of about 0.35 or something like that, which delivered a disaster every 12 years. So it's more risky than that taking any one crop one at a time. But they can diversify over crops. So looking at the crop correlation, those numbers range from 0.09 to 0.81. At .09, it's almost the case that the yields on one type of crop are independent of the yields on the other one.

Again, you can go back and look at the lecture under Consumer Preferences where we were introducing risk and derived these kinds of statistics. Soil isn't uniform either, so even taking as given that they're going to do castor, they can plant it in one type of soil or another. Each type is risky with a coefficient of variation from 0.7 to, again, a little over 1. But you can diversify over soil as well with a correlation of only 0.37, so pretty good diversification possibilities.

That said, households don't hold completely balanced portfolios across soils and crops. One household will plant more of one crop than the other, or for castor, have more of one type of soil than the other despite all the diversification possibilities. And going outwards from crops to other income-generating activities, they can engage in wage labor. They can do trade in handicrafts or animal husbandry.

So for example, trade in handicrafts in Aurepalle consists of climbing palm trees to get the fruit from which they will make palm liquor, for example. And it's a pretty specialized activity. It looks like telephone repairmen climbing poles with pretty heavy and fancy equipment. Animal husbandry consists of grazing animals and so on.

Anyway, if you look at how much land they hold, it varies from no land at all to being a small holder or medium or large holder. And so for example, not surprisingly, the more land they have, the more is their income from crops, going from about 2% to 56% of the total income. Likewise, when they hold very little land, most of the income is from labor, wage labor, and that drops to 4% for the rich guys, large landholders in the village, and so on.

Those patterns are pretty much repeated in the other two villages, and Kanzara especially. Shirapur is much flatter in the sense that we don't see very sharp gradients in terms of percentage of income from various sources as we move from small. So villages are kind of a metaphor. Villages are not all alike. But the techniques we're going to use today do take into account the differences.

Now let's look at diversification over income sources the way we did for crops and soil types. This is a little bit difficult to read, and I'll tell you why. Going down the diagonal here is the coefficient of variation of each income source. So for example, crop production is quite risky with a coefficient of 0.42, which I almost said a minute ago. Of livestock income, the CV is 0.21, et cetera.

Now, these other off diagonal elements here represent the covariation. So what would have been better would have been to list profits, livestock, income, earned wages, and trade and handicrafts going down the rows, and then picking a row, move across the column to find the covariance. So for example, this guy here is the covariance between profits from crop production and livestock income. And it's actually negative. So huge diversification possibilities for a household if they were doing both crops and livestock.

These lower off-diagonal elements are not filled in because the matrix is symmetric. The covariation of activity i with activity j is the same as the variation of i with i . So it would be redundant to fill them all in, and they were deleted to make it easier to read. So taking into account different soils, different crops, different income sources, piling them all together, and plotting income over the 10 years for which we have the data, we look what happens over time from 1976 to 1985, essentially.

And what's getting plotted is not just the time series of one household. All of them are getting plotted here. In fact, the way the data were coded, low numbers correspond to little or no land, and high numbers correspond to lots of land. So as we move backwards along the household number, the households are getting more and more land and are actually-- net worth is going up. So that 1% of the population that owns 50% of the wealth, that's this guy back here.

Now, I refer to this graph as the Rocky Mountains because you can see all the peaks-- well, not quite. I was going to say you can see all the peaks and valleys. That's wrong. When you have a peak, you can't see the valley behind it. If all these incomes communed with each other, then everyone's peak would happen at the same time, and it would look like waves. Instead, we see this rather jagged, sharp picture.

And again, when you're able to see the valley behind, it's because there's not a peak in front of it. Now, the other thing is moving along this gradient-- I should have pointed this out earlier-- there is a 0 here. This 0 corresponds to the overall average of income both over households and over time. So these guys are below-average income. These guys behind it are higher-average income. And again, you can see that that's quite a salient gradient.

Here is the corresponding picture for grain. And I could have plotted all consumption items, but it would have looked similar. The paper actually does that. When I say corresponding, I mean the scale hasn't changed. The scale that was used to plot income, 4,000 at the peak, is the same as the scale that's being used to plot consumption.

So you can see everybody's consumption is compressed, so they're smoothing. They're smoothing over time. And to a degree there, a large degree, they're actually smoothing over households as well. Those very wealthy guys are consuming more than the poor guys but not nearly as much as income is higher for them relative to income of low guys.

The main point of the picture, however, is that here, you can almost see the waves. Grain consumption is kind of co-moving. That would be easier to see if we blew it up, but we will look at this statistically as well in a few subsequent slides. And I refer to this as Kansas. I don't if any of you have ever been in Kansas. I ask this every time I teach this diagram.

At one point, some student in the class said, I'm from Kansas. And I said, oh, I'm sorry. I hope I haven't insulted you. And he said, no, no. It is really boring there. So I like Kansas. I'm not saying that. But anyway, amber waves of grain, so consumption versus income. So let's think about a theory to try to explain the difference between those two diagrams.

And this is the bit that's a little bit of a review, but I think a useful one. So we're going to talk about individual k alive at date t that has experienced a history of shocks h_t . This history is the shock at date 1, the shock at date 2, and the shock at date 3. So again, this is the state space representation of the tree that you saw in Debreu, and I featured at least twice already.

What is the objective function? We're going to solve that programming problem to go back and forth between solutions to that problem in Pareto optimal allocations. So we're going to maximize a lambda-weighted sum of discounted expected utilities, weighted lambda k for weight k over all the individuals and their capital M of them.

The contemporary utility functions have as arguments the consumption of individual k at date t , c_{kt} , which we just did, and then this age-gender index. So not everyone is the same age. Not everyone is the same gender. 18-year-old males eat twice as much as anybody else, essentially. So we want to adjust the utility for these metabolic requirements. And they may not be treating the women very equitably either. But we are going to take these as metabolic weights that were measured in a dietary survey.

So the other objects here are the probability of that history actually happening. Note that for any day t , we sum over all possible histories that could have generated that node at date t . And then in addition, there are multiple dates, so we sum over those dates and discount them by β . And again, the lambdas are the Pareto weights that vary between 0 and 1, and they sum to 1, and we take the weighted sum.

There are many, many resource constraints. There is a particular constraint for every date t and all possible histories that could have led to t that simply says that the way you hand out consumption must not sum to something greater than the total that's available. So usually, we talk about income. Now we just went to total expenditure.

But obviously, whatever income is, if it were spent, it would lead to expenditure, and then you could reallocate it. And some people might be consuming more than others. But here we take the total aggregate consumption as fixed. We're not going to put labor in here. There's just consumption. And we get first-order conditions.

So if you pick a particular consumption for a household k at day t as a function of a particular history h_t , it would appear in the w contemporary utility function with the prime for the derivative weighted by lambda k . And that β to the t probability of h_t would also be over there, but we've multiplied through. So that now becomes a term in the denominator.

And the numerator term here is μ . So let me go back. So the μ is the Lagrange multiplier on this resource constraint. And it really should have a μ indexed by t and the history h of t because there is a constraint like this for every single possible history end date. So there are many, many constraints. That's why I caught myself when I said subject to one constraint. That's wrong. There's only one good, but there's many dates and many states.

The μ that's appearing here-- I didn't write an additional little t which ought to be there probably. But anyway, μ tilde is the normalized version after dividing through by the discount rate and the probability. And if you started with particular utility functions like one of our favorites, constant exponential utility, where σ is the degree of risk aversion, it is of this form.

Now, there is a step that's happening here which is imposing the way that these age and gender weights are entering into the utility function. Namely, it's consumption per unit index. So pick a 30-year-old male and then look at the other age groups and the other genders. And the question is whether they're consuming more or consuming less than that baseline case. So we, for short, can talk about this consumption per unit age. That's really the object that's entering into individual K 's utility function. And c tilde then replaces that normalized consumption c over a .

And when you do the math, you'll get the consumption formula, the risk-sharing formula, as it were. So this is the per-unit consumption of household individual K at date t . It has an intercept term, a term that depends on the demographics, and a term that depends on this Pareto, this shadow price μ . You've kind of seen this before, but let me remind you this intercept term, quite logically, has to do with the Pareto weight. The higher is the Pareto weight of individual K , the higher is the log of the Pareto weight, it's normalized by this degree of risk aversion.

If we move over here to the μ term, μ is the shadow price of consumption. It's how much the objective function would increase if you increase consumption by an infinitesimally small amount in the aggregate. So as μ goes up, consumption is going down. Why is that? μ is the shadow price.

So the lower is aggregate consumption. The higher will be the shadow price. So as μ goes up, there's less consumption to go around, so individual consumption is also dropping. And we will replace this by aggregate consumption and get a positive sign momentarily.

The other interesting thing about this equation is, again, this $1/\sigma$ which I feature in this term, although it's on all the other ones as well. And that is if individual risk aversion goes up, this coefficient is going down. So as risk aversion goes up, this individual K should not experience too much variation in his or her consumption when the aggregate is moving because other people would be in a better position to absorb the aggregate shock if they were less-risk averse than individual K . And again, you'll see that momentarily.

Questions?

STUDENT: My question is why there was a negative sign before the A , $\log A$, the age and sex index. It looks like in the utility function when A goes up, the utility will go up.

ROBERT TOWNSEND: When A is going up, though, consumption per unit age is going down.

STUDENT: But in the utility function--

**ROBERT
TOWNSEND:**

And they're becoming more urgent, so to speak, other things equal. So I'm sorry. I think I said it right the first time. Holding aggregate consumption fixed, for example, as their index is going up, the consumption per unit index is going down. And that's the negative sign. It's really just that simple.

All right. But thank you for asking the question. All right. Now, that was just one individual. We can put all the individuals together and allow explicitly for different risk aversion. And this formula appears in the paper, and I'm not going to try to derive it. It's somewhat tedious algebra. But the interpretation is easy. Namely, there's an intercept, something to do with demographics, and something to do with aggregate consumption.

In a little bit more detail, this is no longer individual I consumption. It's all the individuals in a household J . And unfortunately, that's all we see in the data. We don't see individual eating apart from that dietary survey. We just see aggregate household consumption.

But that functional form allowed us to sum up over all the consumptions, and that's what we see. And we also have the age and sex of all the members. So we also have the denominator. So this is household-level consumption per unit age. And that should be higher. The higher is household J 's lambda weight relative to the lambda weights of the other people in the village.

It's a bit more complicated because these lambda weights are weighted by the degree of risk aversion. So it's a risk-aversion-weighted version of the Lambda that we care about, comparing it to lambda K . This is the demographic. I may come back to that. This is the formula now in terms of aggregate consumption. This is the consumption of individual K and household J -- I , actually-- but then summing over all households I .

So this is aggregate consumption per unit age in the whole village. And again, you can see quite explicitly that J 's consumption will move with that when σ_J is going down, or the other way around-- when household J is very risk averse, this coefficient is going to 0, and individual consumption would not vary much with the aggregate. That's the aggregate allocation of risk bearing.

Because the theory has those age and gender weights in it, you can treat those as shocks. Even though they're quite not moving in a surprise way, they are moving over time. They're moving intertemporally, and maybe the whole village is getting older. Maybe one household is moving across metabolic categories faster than the other households are. And so the formula allows for that adjustment, comparing household j to the other ones.

So let's look at an empirical implication, a linear equation. The aggregate consumption, age weighted, of household J has an intercept. It depends on aggregate consumption. It depends on the index of household J and maybe on other things. This is just a version of this. Household J 's consumption depends on an intercept, its age, and aggregate consumption. It's just much easier to write this way.

And the theory has implications. Namely, if everyone had the same risk aversion, then this β_J on aggregate consumption would be 1. You can convince yourself this would be 1 because the sigmas are all alike, for example. Likewise, this demographic thing would have a coefficient of $1/\sigma$, the common sigma. And this other stuff-- I saved the best for last. This is any other variable, including household income. And the theory is saying that this coefficient should be 0.

Now, that's very counterintuitive. How could I possibly be claiming that household consumption doesn't depend on household income? What about marginal propensity to consume and all of that? Well, the answer is they're pooling their incomes together as if in a mutual fund and then redistributing it as a function of the Pareto weights.

So household J's income is kind of in there, but it's in the pooled version, and it shows up in aggregate consumption. And if you control for that, then there's nothing leftover. There's no idiosyncratic risk for household J to bear. Another way to say that is we're going to get perfect consumption insurance because, like a large mutual fund, all the individual shocks go to 0.

A version of this in China that I'm becoming aware of-- Alibaba runs a mutual fund over health shocks. There's 100 million households joining, and they pay ex-post premia. Typical medical expense must be as high as \$20,000. How much do they pay in? \$9 in yuan equivalent.

So when you spread that damage of an individual household, which is high, over all the hundreds of thousands of households, it almost goes to 0. And this is the extreme version of perfect pooling. That's contemporary China, not a village in India in 1985.

So we can test the theory by looking at each household one at a time over the 10 years, and we can test the theory by pooling households together into one big cross-sectional time-dependent regression. This is a little hard to see, but what's going on here statistically is testing these coefficients α , which should be 0, to see whether they actually are in practice. Under the null hypothesis that it's 0, can you reject in a statistically significant way that it's greater than 0 or less than 0?

And for most households, the bulk of the households, you can't reject much one way or the other. There's just a handful of households on either side of the null. This test has power. This thing is, oh, well, maybe it's just noisy data, and we can't infer much of anything.

Maybe they're literally hand-to-mouth, and they're eating entirely their income, in which case the coefficient would be 1. And we can't reject that for half of them. But for the other half, we can statistically reject that it's 1. So there is a lot of smoothing going on. And this is for the other two villages down here.

Now let's talk about pooling it together. Actually, let's just jump here, and then I'll go back. So here's the villages-- Aurepalle, Shirapur, Kandara. And we're going to somehow take each household over time and take households in the cross-section and come up with a coefficient on how much consumption is moving on average with household income. And the answer is roughly-- should be in rupees-- \$0.07 for the dollar. So for every dollar change in income, consumption is moving by only \$0.07.

It's positive. It's statistically significant. We now reject the model. But as an approximation, we're doing quite well because these coefficients are small. So most, although not all, of the household-specific idiosyncratic risk is being pooled. And you can peruse this table and see-- this is profits as opposed to all income, labor income, profits from trade and handicraft. You can look for the high numbers.

So there are certain activities in certain villages that look relatively under-insured. This is who's the most vulnerable. If we had the ability to come in externally and introduce a better insurance product, we would want to offer it to households that need it. Or if no one needed it, we would expect take-up to be zero. So labor incomes had a pretty high coefficient in Shirapur and so on. But a lot of these are 0-- that is to say, not statistically different from 0.

So in that sense, it's a good approximation. Remember, according to Lucas, models are abstractions. They're meant to be stylized versions of reality. Because we have a model, we have an exact prediction of what we should expect to see. And through the lens of the model, we can decide how reality is doing relative to the model. So I'm not overly alarmed that we're going to get statistical rejections. The question is, is it a useful starting point as an approximation? And the answer seems to be yes.

I will share with you that when I wrote this paper, it kind of created a firestorm of alarm among policymakers who were astounded that poor households in villages in the caste system, and so on, could ever possibly be achieving anything close to an optimum. It just went against the grain of their prior presumptions. But it doesn't fit perfectly.

So this slide which I deliberately skipped over just now is to say, look, maybe I'm getting into statistical stuff a bit, and we can follow up later. One reason why this coefficient might be low is because this variable x is very poorly measured. Say it just got a lot of measurement error that is not real. Then it's moving around a lot for spurious reasons, and you would not expect that to influence consumption. So that biases this coefficient to something close to 0.

Fortunately, we can correct for it with the panel. And it's putting explicitly this measurement error into, in this case, idiosyncratic income and then substituting that into those equations. We now see a correlation between the measurement error and the measured version of income, which means you can't run an OLS regression.

But on the other hand, another way to do this is just to take first differences. Differences over time over households gives you another version. And it turns out this within coefficient and this across-time coefficient allow you to essentially get rid of the measurement error.

So that's this IV Griliches-Hausman thing, which I'm not going to go over. I still don't know what the best strategy is for me. If I skip stuff, then it looks like it came out of nowhere, and you don't really understand it. Hence, I try in all these slides to show you the steps that are being taken. I'm deliberately not concealing something from you. But on the other hand, this particular sequence involves statistical material you may not have seen something similar to before. So if it's still quite opaque, don't worry too much about it. I'm just trying to be clearer.

OK. So that's summary, risk and insurance and village India. These villages are doing reasonably well, although not perfectly. They're sharing a lot of the risk. The reason that consumption looks so flat relative to income is because they're actively, somehow or another, engaged in borrowing and lending or transfers to bridge the gap between consumption and income.

Now, another way they could handle the situation-- so for this, we're going to jump to medieval villages-- is to potentially divide up the land. So this is the diversification part, whether land holdings of a typical village would be consistent with the optimal allocation of risk predicted in that model just now on the premise that what you get is what you eat. You have these strips. The strips are subject to risk. You add it all up across all your land holdings, then you eat it.

So this is the other extreme-- no expo smoothing, no credit markets. Even going back to lecture 1 or something where I presented all the different economies-- we talked about the medieval village economy. We're about to do it now in terms of dividing up the land. We also talked about Thai village economy with that temple where households did not diversify.

Some had low land. Some had high land. They made donations to the temple. And people that had a bad year, God gives back. So the second thing is what's going on in the first, some mechanism for risk sharing ex-post. This is the mechanism in the second, which is ex-ante division of land. I said that in reverse order, but anyway.

So you've seen this picture twice before. This is Elford Stratfordshire. Mr. Darlaston's land is shaded in black, and he's got something like 50 strips throughout the village. So why were they doing that-- not just Mr. Darlaston, but all of them to some degree?

So if you think about this map, for example, are these types of land differ in some way from these types of land in terms of soil or slope or elevation? Or alternatively, land is all alike. And then the question is, when the storm comes and destroys crops, it could destroy these crops and not these.

So there's really two different models, and we're only going to have time to do one, although the other one is covered in that missing medieval village chapter. We're going to do uniform shocks on non-uniform land, and the non-uniform shocks we'll save for another time. So there are two types of land. Start off keeping this simple. Say high-low or clay-sandy. There's only two households, another abstraction. And there's only one date.

And land type k has a yield vector which varies with the shock ϵ . And ϵ can take on capital S values. So this ϵ is a measure of rainfall, temperature, humidity, or other events which have beneficial or adverse consequences for the yield as captured by the way which the yield of type k land varies with those shocks. You've, again, seen this before a bit. When we did the dynamics, we planted land with seed and talked about how the yield could vary with low, medium, and high shock. So this is very similar notation.

Now, let's be serious even though we're trying to keep this simple. What numbers do we use for this vector? One thing we want to capture is how much that yield is varying as measured by the coefficient of variation. And in medieval villages, it was about 0.35, which we've covered before, and I'm not going to go back and dig up those slides for you.

But that was the number we used before. The other one is-- there's two types of land here, k equal 1 or 2. And we want to choose vectors of returns for both. So we want to capture the covariance in the two types of land, how different are they from one another.

And again, going back in the earlier lecture, the estates of the Bishop of Winchester, we had a cross-space correlation of 0.6. So we're going to take those numbers, .6 and .347, as givens. But here's a simple representation.

Let's take, again, only two types of land, 1 or 2, only three events, epsilon equal 1, 2, 3. Type 1 land has the highest return in event 1, then medium, then low. Type 2 land has the highest return at event 2 surrounded by the lower numbers, 3 and 2.

So if you aggregated up the total yield over all types, it would just be the sum of those columns. And it would go 12, 10, 7. Now suppose we divided up the types of land in a very extreme way. We let household 1 have all of type 1 land and household two all of type 2 land. And the question for you is whether that would be an optimal allocation of land. If we did it that way and they were forced to eat the grain off their strips, would the consumption that we see correspond to the implications of the model?

And the answer is going to be no because the aggregate is going down as epsilon is going up. And that's true of household 1's consumption, but it is not true of household 2's consumption. This proposal fails the monotonicity requirement of the theory that as aggregate consumption goes down everybody's consumption ought to go down, although some people's consumption go down more than others. So we reject this proposal for the allocation of land if the goal is to achieve an optimal allocation of risk bearing by land division alone.

Well, there are other ways to do it. Let's take the land type and give household J a proportional number of shares in the output of each type of land. So that means household J would get fraction α_J of the output over each type of land. So we're kind of re-dividing things here.

We're going to have household 1 getting some proportion of this column but also some proportion, not necessarily-- yes, in this case, the same of the other type of land. So this α_J doesn't have a K on it. The shares depend only on the household, not on the type of land. The shares can depend on the household.

All right. So it's linear in the aggregate. If we were over at school, you'd come up to my office, and you'll see my bookshelf filled with stacks of books devoted to medieval English villages. I had a happy time doing a lot of research. And you'd discover these amazing things.

Holdings were scattered over fields so that each household partook in equal proportion of fertilizer and poor soils. Dividing up the land in accordance with drainage and exposure, resources were divided into shares. One man might have more shares than another, but the shares themselves were equal.

Now, there are potential problems. For the theory to fit exactly, we got to worry a little bit about the intercepts because those divisions did not allow the intercepts to be other than 0. If you go back to where we first introduced risk sharing and looked at particular utility functions, I showed you, as example of the risk-sharing rules, two equations-- constant relative risk aversion, constant absolute risk aversion. And the intercept was 0 under these various types of conditions.

So it's not like they can't be 0, but it does require certain parameters to hold in order for that to be the case. So a little bit of a qualification. But all is not lost. We can even do more. And namely, we don't have to have a given household having the same number of shares in every type of land. And that's going to allow various kinds of intercepts and so on.

So here's the idea finally. Suppose there are at least as many land types with independent return vectors as there are states of nature. So what I showed you just now violated that. There were three states of nature and two types of land. But if we had three states of nature, we'd have to have three types of land with independent return streams. And likewise, s types of land, we need s dimensional return vectors. So one column cannot be derived as a linear combination of the other columns.

Now let's solve the risk-sharing problem and take this consumption to be the target, this consumption of household J in state ϵ . And we're going to do that, obviously, over all states ϵ and all households j , their little n of them. We get an equation of this form. So don't get lost in the notation. I'll help.

This is consumption of household J in state 1, consumption of household J in state 2, likewise in state capital S . That's a target. And just imagine it solved the Pareto problem. How can we achieve it? We have as controls the fraction of type k land held by household J . And now those fractions, α_j , vary with k . So we're allowing land-type-specific allocation of shares.

Anyway, if you do the linear algebra here, here's the yield of type 1 land in state 1, the yield of type 2 land in state 1, et cetera. These alphas would then multiply-- $e_1, \alpha_1, e_2, \alpha_2$, et cetera. So adding it up like the dot product, we would get consumption. Consumption is coming from a share-weighted version of the yield. Hopefully, the algebra is clear.

Writing it in matrix notation, we just want to get this quote endowment matrix e times the share vector for household J varying over type k land. It's a vector equal to the state-dependent consumptions. So e times α equals c_j . So then the question is, to achieve this target, can we find alphas that do it? And that's like solving this equation. So essentially, take an inverse of something.

If there are as many return vectors over types of land that are independent of each other as there are states, then this matrix has an inverse. It's non-singular, so you can solve this equation. So now we'll have solutions for the alphas. And it turns out if you do that not just for household j but over all of them, you'll get an implication that the sum of the alphas will add up to 1.

So they do look like shares. Type k land held in fraction α_j by the j -th household when summed over-- there should be a sum here-- there was one there-- when summed over j would equal 1. So all the land is divided. All the land of each type k is totally divided among all the households.

You could run into a little bit of problem if the alphas went negative because then that's like trying to go short. We're not going to allow that to happen. So there's a little bit of a fly in the ointment. But roughly speaking, there's a way to achieve an optimal allocation of risk sharing just by shares. So they don't need to transfer risk around ex post. They can do it all ex ante.

Now, in order to achieve that, we had to ignore the negative alphas. But more to the point, we had to assume there are as many types of land as there are states of the world. Maybe that's not true. I don't know. It could go both ways.

How many different types of land are there, and how many different states can a given land type take on? If there's incredible heterogeneity over land, it's entirely possible that the number of types was equal if not greater than the number of states. But I could also imagine the other way around. Like the Thai village, you've got low land and high land. You put them in two categories, and then you have multiple states. You'd have an incomplete version.

Let's see how bad it is, though. Maybe as an approximation, they were doing pretty well by dividing up the land even if they couldn't achieve the optimal allocation perfectly. So let's put in two different utility functions. For household 1, we have constant relative risk aversion, and for household 2, it's a quadratic function.

This quadratic thing is a pain and always scary when you see it. It's like there's a bliss point. We've talked about this. Consumption is less than the bliss point assumed so that the marginal utility of consumption is always positive.

Otherwise, you see a lot of negative signs, and it's kind of disturbing. Can this consumption be higher than the bliss point? No, because you'll get marginal utility being negative, moving negatively with consumption.

So anyway, just so you feel a bit more comfortable, let's put in α equal to 0.5, k equal 0.002, and b equal to 80. So we have completely parameterized this family of utility functions. And again, we have different preferences here but nice concave utility, strictly, for both. We've got two types of land, and we're back to, say, having three or more states of the world. So ex ante division into shares cannot work perfectly.

To see how far we get, we constrain the programming problem to nevertheless search over these shares, α_{jk} of household J share of type k land. And we maximize a λ -weighted sum of ex ante expected utilities. Oh, well there's only one date here, so ignore the dynamics. But there are multiple states.

And we're going to max this thing subject to shares can't go negative, and they have to add up to 1. So it's a constrained programming problem. We could call the solution to this constrained optimal. It's constrained by the premise that the mechanism only allows ex-ante division into shares and not ex-post consumption transfers, not gifts, not borrowing and lending, nothing.

Now, you know from our work on Pareto optimal allocations there's going to be a Pareto frontier. So it's not like there's one solution. It's going to vary with the λ s. So we'll do something special as an example, which is to maximize the utility of household 1 subject to the utility of household 2 at an arbitrary constant. So let me show you a picture, and then I'll go back.

We're in utility space, so we have the utilities of household 1 and household 2. And we have-- if we didn't have a constrained problem, we'd have this outer frontier of the set of all possible utilities on the boundary. And we'll call that the utility at the full risk sharing. There are many solutions. We got these through hyperplanes previously.

There is an equivalent way to do it, which is to fix the utility of one household, namely 2, and maximize the utility of the other one. So we would end up right here. But in the constrained problem, we're limited to dividing up land by shares. That constrained Pareto frontier is interior, still a bit concave. And we would move, therefore, from where that horizontal utility level for household 2 hits the outer frontier versus the inner frontier.

So let's focus on the inner frontier, the constrained problem, and then see how close they're going to get to the outer frontier. I solved that thing numerically. I was trying to give you all the parameters for utility functions. Some of it's missing here-- covariance, coefficient of variation. And you get a solution that household 1 has about 3/4, 72%, of land type 1 and about a quarter of land type 2. And household 2 would be on the other side of this at roughly 1/4, 3/4.

So one thing to note is that that generates a non-linear sharing rule but non-optimal rule. I think I missed-- when I presented this constraint problem, I neglected the little paragraph at the bottom. No, it's not there.

Well, let me remind you. Again, back to the preference lecture where we did risk, and I showed you two particular utility functions, constant exponential and constant relative risk aversion with linear schedules. But then I showed you another figure which was non-linear-- monotone increasing for both households, but nonlinear. That came from this specification of utilities.

So that's what I was looking for, seen in previous lecture 7. So we know the optimal sharing rules are nonlinear. Dividing up into shares suggests we're going to end up with linear rules. So we're not going to end up with the optimum. But we can still try to figure out how close the inner frontier is to the outer frontier.

So let's look at the transfers. Let me state this clearly. We're going to try to do it ex-ante by division of shares, get the constrained optimum, and then look at transfers across the two households, who's paying who depending on the state, and then finally, how those transfers would need potentially to be augmented. So here's the aggregate endowment. This is the sum of the yields over the two types of land. That's also aggregate consumption because this is a static problem.

Here is consumption of the individual agents, both of them. This is the consumption of person 1, constrained, linear dotted line. Here's the consumption of person 2. Constrained is this dot dashed line for person 2. Constrained is the dot-dot-dot line for person 1. They're both linear schedules.

Now, remember, they're not doing this with transfers-- I kind of misspoke a minute ago-- because they're doing it with ex-ante division into shares. But nevertheless, it's not the full-blown optimum. It's not the outer frontier in which you can not only have division but also ex-post transfers to get to the full optimum. And those schedules are listed here as unconstrained. And they are non-linear, as I was trying to say earlier.

So if we look at these states, say five states, and look at consumption under ex-ante division of land versus consumption under the full optimum, for household 1, you can see household 1 is monotonic. Household consumption is logically increasing with the aggregate, but it's a bit different, although monotonic, from what it would be under the full optimum.

And here you can see at the lowest state, household 1 would be consuming more under the full optimum but then less under the full optimum relative to the division, less, and then back to being more. So here it is over here. Why am I struggling so hard?

So this is the amount of the transfer from household 2 to household 1 that is necessary to add on to the solution from the ex-ante division of land shares only. If you wanted to look at it in a picture, the picture is literally the difference between these two lines.

So here, the amount that household 1 ought to be getting under the full optimum would be higher, and therefore, ought to get a positive transfer. Here, it would be lower, hence the negative transfer. And finally, at the end, positive again. So that's consistent with these numbers where these states, going from 1 to 5, is like moving across the different values of the aggregate endowment.

So in summary, I guess there is a couple of takeaways from the second half of this lecture. One, we looked at special mechanisms-- in this case, ex-ante division of land and whether or not we could get to an optimal allocation of risk bearing that way. And we saw some special cases where could. We had to make the shares different across the different households and maybe different across the different land types, especially so if we accommodated more and more diversity and preferences.

The second takeaway from the second part is we might not achieve a full optimum, so we would have expected some kind of borrowing and lending and gift-giving in that medieval village economy even if we can't find it in the historical records. Or we could just say, evidently, for institutional reasons we don't understand, they just weren't doing it.

But we don't lose our tools. Not all is lost. We are still able to solve for a constrained optimal allocation solving the so-called Pareto problem or programming problem, maximizing lambda-weighted sums of utility subject to constraints with the appropriate control variables, in this case shares.

And then the big picture of this lecture is we're looking at the application of the theory of the optimal allocation of risk bearing. We did that in villages in India, looking at consumption and income data. So from a science point of view, we have a model. We have data. The data can vary across the applications. You can still ask, even if we don't have the data on everything, whether we can see how well the theory fits.

And we could do that in-- in the medieval villages, we don't even have data on consumption of individual households, and all we see are pictures of land. But we do have the covariance and variance statistics. So we're kind of able to calibrate the model a bit the way we were calibrating the dynamic storage model, but in this case, not to try to explain why they didn't store very often, but rather, to explain this very salient pattern of dividing land into different strips.

OK. So that's all I have for today. All right. See you soon. Thank you so much.