

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: So let's get started today. First of all, are there any questions from last time? All right. Well, please feel free to ask questions along the way.

In terms of the reading list, we're up to lecture 9 today. We're sharing with production. It's our second lecture on applications.

There are two starred papers. And we're going to be going through the first one with Samphantharak in detail, and then a few selected slides on the Kinnon et al paper. But like all starred papers, I hope you will be reading them, because I know you're very well trained now.

And you will be able to read those articles. So I think you would get a lot out of it. But we will cover most of it today.

The other thing for the study guide-- so let me ask you guys some questions. Last time was our first applied lecture on risk sharing. We did it in both village India and medieval villages.

And today, I want to give a little bit more emphasis to the economics combined with the analysis-- more than I usually do. But it seems like the time to do that, since we're talking about applications. So the first one I picked here-- describe in words why villages in India are a risky place, how to quantify that risk and its sources, and how that risk might be mitigated or potentially mitigated via individual or social diversification.

So I have lately been taking volunteers. But let me ask. Dexin, do you want to take a stab at that?

AUDIENCE: Yeah, sure. So villages are a risky place, because it's mostly agriculture. And it's not a stable source of income.

And I think we can quantify that risk by just looking at the variation through the years and how it really changes, and how a lot of things are linked with each other in terms of their risk. And it can be mitigated by planting different kinds of plants, or planting far apart-- that's individually. And social diversification-- giving gifts or transfers. Yeah.

ROBERT TOWNSEND: Great. Very, very good answer. Thank you.

So we quantify the risk by the coefficient of variation of a given crop or a given crop in a certain soil type or a certain income source. Just to add a few more details-- your answer was fine. The covariance across those things actually helps mitigate the risk.

And that's what you're saying when you said different crops or different spots on the field. Those are ex-ante. So there's individual things that people can do. And the ex-post-- social diversification is the gifts. OK.

Next one. Comparing incomes over households and time with consumption over households and time, what can you infer from the data alone? So let me stop there and not read the second part.

Let's see. Who would like to? Charles, do you want to take a stab at that? Are you there?

AUDIENCE: Yeah. Let me think. If I remember correctly, this is something like you can compare the incomes, which is basically the crop harvest, and then what you end up receiving. And I remember, you can infer from-- you can see that it's definitely much-- you end up getting is much flatter than what's produced.

I guess I'm trying to answer the second part of the question now. But that is data that suggests that we're risk sharing. I don't think that's enough to prove causation, but it's definitely suggestive.

ROBERT TOWNSEND: Yeah, that's good. Maybe the second sentence helps. If average consumption of households increased when average income of households increased, would that be consistent with risk sharing or not?

AUDIENCE: Oh, for me?

ROBERT TOWNSEND: Yeah, for you.

AUDIENCE: So OK. So OK, average consumption has to increase, but-- I don't think it rolls out to some sort of big sharing, I don't think.

Because I think you have to parse a degree of risk sharing that wouldn't completely annihilate everything. But I don't think that would be enough. I think that would still be consistent with some level, right?

ROBERT TOWNSEND: Yeah, actually, the sentence is not terribly well worded. But it is worthwhile to clarify. So the first part that we looked at is just the levels of stuff over households and over time.

We referred to the Rocky Mountain picture as incomes going up concealing the valley behind, and very, very jagged rocky pictures as opposed to consumption, which was very flat. And we were even joking around a bit-- very boring, as in Kansas. So then, the issue would be with the theory that we're about to look at.

Apart from the data alone, would you think household consumption ought to move with household income? Actually, I think there's a question below. And what's ambiguous about that sentence is, household consumption is increased with average income of households.

Yeah, that's very treacherously worded. Average income meaning overall the dates of a given household-- that sets the levels of things. And maybe the household with higher average income over time has a higher average consumption over time.

And I'll ask for volunteers. If we have this regression where we run household consumption, on the left, as the dependent variable onto aggregate consumption and individual income, on the right, throwing in some intercepts-- so what restrictions from the theory are placed on the coefficient of individual income? And I'll take a volunteer for that.

AUDIENCE: I think the coefficient individual income should be 0, right?

ROBERT Yes.

TOWNSEND:

AUDIENCE: Because it's just related to the aggregate in average income of the household. And I have a question here. It's that I remember in the slides-- saying that when all the households have the same risk aversion, then the coefficient between the aggregate consumption should be 1.

ROBERT Yes.

TOWNSEND:

AUDIENCE: But why? I think the coefficient on the average consumption should be when not the aggregate consumption.

ROBERT Oh, well there's, say, a fixed number of households. So taking the aggregate and dividing by the number of
TOWNSEND: households gives you the average. It's just a renormalization of the units, that's all.

AUDIENCE: Oh.

ROBERT But you're right that it's meant to represent the aggregate. So you can think of it that way.

TOWNSEND:

AUDIENCE: OK.

ROBERT Because the economics makes more sense. The aggregate is going up, coming up over all the individuals. And if
TOWNSEND: there's constant absolute risk aversion and the same risk aversion for all the households, then whenever the aggregate goes up, every individual household consumption goes up 1 to 1.

AUDIENCE: Yeah.

ROBERT Which is what you said. OK.

TOWNSEND:

AUDIENCE: Thank you.

ROBERT Yeah. And the emphasis here is the household-specific income. It has a coefficient of 0. That's the salient
TOWNSEND: restriction that the theory places on the data as well.

And then finally, we looked at the division of land in these medieval village. And the question is, would scattering of an individual household's plots across the different fields be enough to achieve an optimal allocation of risk sharing? Paolo, do you want to take a crack at that?

AUDIENCE: Sure. I think the answer is no, just because it's not necessarily that each field is a random draw. You can have correlations like, it's just a drought here and there's no rain anywhere. That's a risk, which would be a shock to all of your crops, essentially.

ROBERT Yeah. So it's correct that it's not necessarily enough, unless you have very special characteristics of those plots--
TOWNSEND: how they move with various states of the world. In the lecture, we gave an example-- two examples, really-- one where there are two types of land and both are increasing as the states increase.

And then, it would be enough to have each household have a share, a constant share over the two types of land. That would result in a linear risk sharing schedule. So for certain utility functions, including an examination of the slope and the intercepts, that would be enough.

So in special cases, it might work. But in general, you have two different types of land. And the output goes up in one. It may go down in the other, depending on the state of the world, in which case, handing out shares of each type is not going to work, necessarily.

And we explore that further in the lecture. So that whole lecture culminated at the end with an example where we had more states of the world than we had land types and restricted them to having shares of each type as if they could not transfer stuff around. And it was clear that it was constrained optimal.

Taking that allocation of land with the harvest as given, we could do better by, add to welfare by having exposed redistribution of the grain via transfers. OK. So that's probably enough of a review. As always, there are many more questions in the study guide.

And probably, spend quite a lot of time just deciding which ones to ask in class. You should look at all of them. And let's get to lecture 9, is what we're doing today. And this is risk sharing with production.

Three main topics, really. Risk and return in production-- and that part is continuation of some of the themes, although we didn't really put all the pieces together. We had production in the dynamics lecture. We had risk as in yields coming from land plots.

But we never talked about inputs into land plots, and associated outputs, and so on. So we'll put those two pieces together. Another feature of this lecture is that we're going to use data on production but also on consumption.

So both things will be there. And this is illustrative of how powerful this risk sharing framework is. It extends to more and more complicated situations.

We're able to do more and more with it. And finally, I have some words at the end about social networks and sharing risk, and in particular, an example that will remind you for better or worse of COVID, and what's going on when people get sick with the economics. OK.

So this first starred paper on the reading list with Krislert Samphantharak. It'll tie, if you haven't guessed, extensions of the theory project selection. So we studied a model with risk and return in productive assets in developing economies-- namely, Thailand.

We're going to measure idiosyncratic and aggregate risk premia. Those words probably don't make a whole lot of sense, although we've been talking about related things just now with the risk sharing regression, idiosyncratic being household-specific income, and aggregate being aggregate consumption. And we're going to be able to quantify those two different premia and analyze what households are doing, how exposed they are to risk.

And finally, in principle, one could take not just the straight-out returns, but risk-adjusted returns and look for patterns across households, like whether poor households are more vulnerable or less vulnerable, does education matter, and so on. So here's the model. There are capital J households indexed by j , little j for an individual.

J is the number for the aggregate. There are I production activities, little i , big I . And you can think of those production activities as production sets or production technologies.

The inputs are capital. And I'll say some words a bit later about where labor and materials and seed and other things went. But for now, just think of there being a single input called capital.

And the outputs are in the same units of a single consumption good. So we're going to be looking at a household running a portfolio of assets. Now, actually, the way you think about this in finance is, you want to engage in a certain activity.

You acquire the assets necessary to do that activity. And a household could be doing one or multiple activities. This will generate a return. So we want to know how well the project is doing and whether it was selected properly.

So you want to compare the rate of return to the best alternative use of funds except that here, we're going to put risk in. So it makes it much more interesting. In particular, these set returns are stochastic.

And they can be correlated over activities and correlated over households. So again, you can think about the previous lecture with village India, although this happens to be village Thailand and its different kinds of crops and different soils, different income sources, and so on. So I mentioned early on in lecture 1 that we were going to go back and forth between actual real agrarian economies and associated metaphors.

So we can talk about trees bearing fruit literally in the data. But we can also talk about a tree as an asset in financial markets, households' whole assets. And they pay dividends. And there may be correlations in stock returns.

So exactly the same theory is being applied to two very different settings. And in fact, we're reverse engineering in this paper. We're taking measures of risk reduction and risk premia that have been applied to stocks in the New York Stock Exchange and applying it to all these activities across various households in a typical Thai village.

OK. So a little bit more in notation. We have F production function for activity I run by household J . And it has as a single input the level at which it's capitalized, the capital input in activity I run by household J .

So we can sum overall activities and sum overall households and get an aggregate production function. So again, early on in one of those lectures on production, we talked about aggregating two example production sets into one. So this is a version of that.

So the economy starts out with an aggregate amount of wealth, which is coming from two sources. One, the trees, and two, the fruit. Now, let me just say that as is standard in a lot-- not all, but a lot-- of economic models, we're going to have what we call a putty putty, meaning that you take the putty in terms of very malleable consumption, what could be consumption, and invest it into capital.

And then, that yields a return. You can certainly eat the fruit, but you can also eat the tree. So the putty putty part comes from undoing the tree.

Of course, in many actual situations, you can't eat the machine. But it's as if you could sell the machine and get the consumption equivalent. If you wanted to eat all of it, you could.

And you'll see that right now in the notation. We're going to have, quote unquote, "a social planner." But of course, you know now, that just means that we're going to have these programming problems for the determination of Pareto optimal allocation.

And we're going to act as if the planner is doing it. But it's really a math problem that we can solve to go back and forth between solutions to the math problem and various Pareto optimal allocations. So consumption of household J at a given moment-- say a year-- the sources of potential consumption would be the fruit, how much activity I is yielding a fruit run by household J as a function of previously determined capitalization K_{IJ} .

We sum overall I . And we have the tree itself, which as I've said now could be potentially eaten. But from these available resources, we can subtract off the capitalization for next period. Prime means tomorrow or next year.

And K_{IJ} is the level of capital carried out from this year into next year. So this is like storage in that dynamic medieval village economy. That's one activity. Seed planted in the ground is another type of activity.

This is a generalization. And finally, we have these transfers, τ . These are going to be, among other things, key control variables of the so-called planner, adding potentially to the resources of the household, or if negative, taking away. So then, we come to this programming problem for the determination of Pareto optimal allocations.

And despite some heavy looking notation, you'll probably understand this pretty readily. First of all, here's the utility function of household J . And all the slide does is stick into what would be consumption.

The consumption is defined at the bottom of the previous slide-- resources available minus resources utilized for next period year plus the transfer. These λ s are the Pareto weights for maximizing a λ -weighted sum of utilities today. And the control variables for the planner are these transfers, as I've already said, plus the level of capitalization of all the projects.

So we're trying to find an optimal allocation of capital as well as an optimal allocation of the transfers. Small apology-- this apostrophe-looking object shouldn't be there on top of the τ s. It's not a prime, because the transfers are decided today, not next period.

But there's definitely a prime on this capital, because the current state has already given the projects with their capitalizations. The thing to be decided upon is how much capital you want to invest in all the various possible projects for tomorrow.

AUDIENCE: Is it better to add a depreciation here? Because if instead of a tree, if you think of a machine? It will definitely depreciate.

ROBERT
TOWNSEND: Yeah, it's loaded into this output. And you'll see this in the data. You take net profits but you subtract off as an expense the depreciation. But it's true it's not there explicitly, but that's where it is. So yeah, it should be and it is included, OK?

AUDIENCE: OK.

ROBERT
TOWNSEND: Now the other thing that makes this thing complicated-- but again, you're already tooled up to get this right away-- these are these value functions. So what is the state of the world today? It's total wealth that's available at the bottom of the slide summing over all activities and households, all fruit and all trees.

This capital lambda is just the individual Pareto weights. And they're fixed for all time. And then, you decide on these control variables how much to eat and how much to save in the various technologies. That's going to deliver a state for tomorrow, which is wealth tomorrow.

It would be W prime. Just like W is this period, W prime would be for next period. And you would have a prime on all of these objects.

You'd have the fruit in the trees next period at prime. That's the W prime. This value function is as if they lived in an infinite horizon economy.

They're discounting next period by this fee. Usually, we use beta. But beta in this lecture is reserved for something else, like a regression coefficient. So fee is the common discount rate across all the households.

And you remember how we got these V s, right? We had a slide at the end of one of the-- at the dynamic lecture. Act like it's a finite period. So for the optimum, which would be, don't fund everything, eat everything in sight.

And then, go to the next to last period and third to last period. And then, keep iterating as far back into the present as possible. And this value function, which would have a T on it for the length of horizon, will go to something stable, as if we had solved the infinite horizon problem.

So that's all loaded up in here, too. Questions about this slide? Oh, one last thing.

You can't put negative capital into the ground. Capital could be 0, in which case you're not doing that activity. Or it's positive, in which case you are doing the activity.

This is true over all projects I in all households J . You can set that thing up as a Lagrangian with the shadow prices and solve. And you'll get a familiar condition.

If you differentiate with respect to consumptions or the transfers, for that matter, you will get that these lambda weighted margin utilities should be the same over all households and equal to a common Lagrange multiplier on the current resource constraint μ . So you've actually seen this before. It's exactly equal to what we had before in village India without consideration of production.

But we're going to get something else. Namely, when we look at the control variable the capital control variables, we're going to get something that links marginal utility today to effectively to marginal utility next period. And that is that the lambda weighted marginal utility of consumption today for household J -- if this were set equal-- should be equal to the future margin utility. Well, if you take a unit of resource today infinitesimally small, and instead of eating it you invest it, what you'll get tomorrow would be the derivative off the production function of that investment.

Plus you get the tree back, that one unit back. And that return stream is going to be evaluated at the marginal utility of wealth, which is the derivative of this value function at tomorrow's wealth. So again, we're just taking a derivative of this right-hand side term.

Questions? So it's written as a weak inequality. It could be less than 0. If it's less than 0, let's take the margin utility of consumption today, which is minus, and on the left-hand side, put it on the right-hand side. And that would say, the weighted margin utility of consumption today is strictly greater than the margin utility you would get tomorrow.

So they're going to hit a corner. They would like, actually, the capital level to go negative, because they're desperate to eat today. The marginal utility return today is higher than it will be next period.

But of course, we're not allowing them to go negative, hence the inequality. To remind you in other contexts, we write out the full Lagrangian program. We put an extra Lagrange multiplier on these non-negativity constraints.

Then this whole thing would be at an equality. But it would have as well the potentially positive Lagrange multiplier. Anyway, the conclusion is, not all activities need be funded.

You might actually want to unfund them and draw resources out of them that aren't there. You can't do that. So you set them equal to 0. Others get funded at a positive level.

We will be focusing a lot on these set of activities over households for which this is equal to 0. And those are funded projects. OK, so it looks like this is entirely closed economy.

But as I said, these techniques generalize. So we can think of a village as a small open economy. And maybe they can go to an outside bank and borrow money, in which case we'd have to change the wealth variable to subtract off the principal and interest that's coming due today, but add in the debt that they enter into today.

That will come due the next period. Nothing changes. These first order conditions don't change.

Oh, it's good for them in terms of smoothing. If the village wealth is low, the village as a whole might want to borrow. That's fine.

So the quantitative solution can change, but our characterization in terms of these first order conditions will not change. So we do not need to assume a closed economy. All right.

So what the heck is this? Well, this μ is the λ weighted marginal utility. I substituted that into this left-hand side of the second expression, set it equal to 0.

And then, you've got two terms. Dividing through by this, you would set everything equal to 1. So it's actually easier to look at the algebra than it is to say it in words.

We take μ here and divide this term by μ , which makes it equal to unity, and the second term by μ , which is what is being displayed here. And then, the rest is algebra, although it's revealing that we can take this expectation operator outside the whole expression, that these returns are random.

I'm not sure I've been emphasizing that enough. It looked like constants, because you just have the level of capital. But the return on a unit of capital is random.

Where did that first start to happen? I should have said it right here when we had that expectation operator. So anyway, bring the expectation operator out in front of everything.

And then, adopt some very cryptic notation, M -- actually, M' -- for this term, which is essentially the discounted Lagrange multiplier next period divided by the Lagrange multiplier today. So that's a discount factor.

Call it M' , and this R , capital R , is the total return on the project being run, the activity R run by household J . The return-- you get the tree back, plus you get the fruit. And this is a prime here.

It's the marginal return, obviously. Because when we looked at the first order condition, we were varying the level of capital. So we picked up a derivative in the production function.

So this is just notation. But it's very convenient notation. So it's saying that for all activities R run by households J , this product, this expected product, is equal to 1. Furthermore, looking at that expected product, there's a result in statistics that we didn't put into the lecture on uncertainty.

But it's easy enough to state, which is, expectation of two random variables is equal to the product of the expectations adjusted by the covariance of the two variables. If the covariance is 0, then these individual variables are independent. And it's just the product of the expectations.

But the covariance is an easy way to adjust if they're not independent. So now, we just substitute in that formula in the previous slide, which had this expression on the left. We now substitute in the expression on the right and manipulate by dividing through this expectation of M prime.

And we get this object here. So the expected return on activity I run by household J -- expected return next period-- has an intercept in some kind of product term. And I'm about to give you an interpretation of what these are. And in some respects, this beta is going to be a regression coefficient.

And the lambda is going to be a measure of social risk. In particular, finally, if we imagine those production functions are linear, then the return on activity I run by household J is linear in the capital. Constant returns to scale.

However, as I have been trying to say, this return-- this little r thing-- is a random variable. Anyway, the marginal return is that little r . So also, it's the average return, for that matter.

And we can talk about the growth rate of return equal 1 plus the net rate of return as a definition of capital R . Just so you might feel a bit more comfortable, if we have a constant returns to scale production function actually finally modeling other inputs like labor and materials and so on, if you optimize that and put the optimized inputs back into the function, you would get this derived function, which is linear in the capital stock.

So that's where it's coming from. But I didn't give you that algebra. Secondly, what about preferences?

Let's let the value function be quadratic. So you've seen these before. W star is a bliss point. And less than W star-- so that the marginal utility of wealth is always strictly positive. So if we go back to this guy here and start substituting in what we assumed about the production function and utility function, then we're actually going to get an expression which says, the average or expected return on household J 's activities less the risk free rate is equal to beta J for household J times the expected return in the whole village minus the risk free rate.

In fact, you can put this R prime F , the risk free rate, on the right-hand side. So that would be the constant. So it's constant plus theta times this guy.

And let me define these terms a little bit. R prime J no longer has an I on it, because it's the weighted return over all the activities run by household J . Likewise, the return for the whole village-- market, if you like-- M , is the weighted average return over all activities and all households.

So you take the capitalization of activity I run by household J , look at the return on that as a peer activity I run by household J , sum over for the household all activities I , sum over all households J , you're getting a total return over all activities in all households. And to get a rate of return, you divide by the level of capitalization. And the total level of capitalization is just the sum of the capitals over all activities and projects.

And what is this beta? This beta is just a covariance divided by a variance. As I've said, it's going to be a regression coefficient. If you're used to ordinary least squares, you know that a regression coefficient is $x'x^{-1}x'y$.

So x must be in the denominator here, which means that we're running on the right-hand side, the total market village return. And on the left-hand side, the dependent variable is going to be the return of household J . So we're going to regress the return of household J over all the periods of the sample onto the aggregated village level return and get a regression coefficient.

And that's going to be β_J . And that's going to capture how much the activities run by household J are covarying with what the village is doing as a whole. Now, as I think I explained to you last time, I could go from here to here. And it would all just look like some kind of claim.

Go see the paper. If you don't read the paper, it remains mysterious. You have an extra level of abstraction. You don't believe the results.

So I actually wrote in the algebra. I am not going to go through it today, because it is nothing other than algebra. But it's how to take those two assumptions about quadratic utility and production functions, and actually manipulate the equations we already had, substituting in these particular expressions and getting, at the bottom, the expression that corresponds to this, OK?

But I'm going to skip that. So again, mechanically, take the return of household J at year T . Regress it onto constant and onto the time series of the market return. And this is actually a prime here.

We're having trouble editing these latex files, I guess. So this is a prime. So we have prime.

This could be all today, today, or it could be next period, next period. Just be consistent about it. So the return of household J at day T regressed on to the village level return at dates T , running over all T . So the sample could-- I'll show you in a minute-- something like 10 years of data, a bit more.

Plus, an epsilon is the error term. After you run this regression using the time series for each household, one at a time, and having estimated these β_J s, we now run the next regression, which is entirely cross-sectional regression. We compute for each household the average return over the entire sample, which is a proxy for the household's expected return, the data analog to that.

And we regress that for each household J onto its β_J as derived in the first equation-- so 1. So the economics here is, the return on the project run by household J , the expected return, should equal, say roughly, the risk free rate. But it has a risk premium-- first time I've use that now with the notation. And the risk premium has to do with the comovement of household J 's activity with the village aggregate.

So we talked about diversification at the beginning of class. And the idea here is, if the things household J is doing are highly correlated with returns on what other people are doing, then household J isn't doing very much worthwhile. Household J is not diversifying the village level risk.

So in order to justify the set of activities that household J is doing, we have to bump up its expected rate of return. In some sense, we use the word risk premium, meaning household J's activities are pretty risky. Why?

Because they're correlated with the village average. So diversification is not-- and we need to compensate for that lack of diversification. The opposite is more revealing, perhaps, which is, if this correlation is low or even 0, it's a wonderful hedge. Even if not many households are doing those things, if household J does them and manages to produce a return stream, which is not so correlated with the return streams of the other households, that's great, in which case, we don't need to have a high risk premium.

Because actually, it's reducing risk. So that's the economics of it. It's all, also, entirely here in terms of the equations. All right.

AUDIENCE: I'm sorry. I have a question about a previous slide.

ROBERT Yeah?

TOWNSEND:

AUDIENCE: So it's more like econometrics stuff, because I think there should be some estimation error here, because the beta is estimated. And in step 2, you just regress. You regress return on some estimated beta, right? So you just regress some true data, or some estimated data? And I think the lambda here will have some error rate.

ROBERT Well, there are two things to adjust for. One, this is a regression coefficient. So it is measured with error.

TOWNSEND:

It has a standard error. And we do adjust the beta that we're using to allow for that uncertainty. That's part of what you're asking. Is that what you're asking about?

AUDIENCE: Yeah, probably. But yeah, OK. OK, yeah.

ROBERT The other thing is that here, we're regressing an individual J return on the village average. So you do have to be worried about regressing dependent variables on the left with their average on the right. And that tends to bias this beta toward 1.

TOWNSEND: That's another econometric problem, except that here, we're actually interested in dispersion across the different betas. So we're OK.

AUDIENCE: OK.

ROBERT And this is Fama, French, which I suspect you know.

TOWNSEND:

AUDIENCE: Yeah, I know that, yeah.

ROBERT

TOWNSEND:

Yeah. So we're just doing what they did with the New York stocks-- doing it here in the Thai villages, OK? We started our survey in 1998 with a big baseline. We update it every month thereafter. This said, at the time we were writing the paper, the survey was still ongoing, although even then, the later months were not ready.

So we ended up using monthly data from month 5 to month 160 of the survey. The other thing that's misleading-- it's not ongoing anymore. The project ended about 18 months ago. But we did end up with almost 20 years of monthly data.

So it's a vast warehouse of potential projects that can be done with it. And I'm happy to tell you more. But today, I want to focus on this paper.

So we only included the households that were actually in the village for the whole period. Sometimes, entire households picked up stakes and leave. We excluded them.

We also are looking only at households with production activities. Farm is included, but also fish, shrimp, livestock, trade and handicrafts. But not just wage earnings-- why? Because with wage earnings, they're not really producing anything.

They're earning something by working for others. So they're not entrepreneurs. So we did not include them.

And there are extensive networks of family in these villages. And I'm going to come back to that a little bit later toward the end. So just to give you a feeling of the data, we've got net income, which means income from household enterprise.

It's the difference between revenue and costs of all the inputs. Revenue is realized at the time of sale, or if it's a crop, it could be eaten. It could be given as a gift from others at all counts.

We add them all up. Revenue even comes from renting out assets. We include that.

Costs are all the inputs used in production, whether they're acquired in cash or in-kind with credit or gifts. We count them all. We put a valuation on them if it's not acquired in the market.

So we take revenue minus costs, include depreciation-- that came up earlier. Subtract it. And we get net income.

And these assets we're talking about-- we have agricultural, business, livestock assets. We've also got land and household assets. The agricultural assets, to give you an example, includes walking tractors-- you might remember in the production lecture, I showed you a picture of that guy walking around with the tractor-- or larger four wheel tractors, aerators for ponds, machines to put seeds in the ground, machines like sprinklers, et cetera.

So you're in the village now. You can imagine living there, I think. This is the level of detail that we have. We did put household assets in.

That's a decision, because some of the things that look like a durable good end up being used in farming, like a pickup truck. They use it.

So rather than-- but we've done robustness checks. But here, we include them all, OK? So let me just jump here.

If we look at that data, this is the return on assets, the expected return on household J. Each J corresponds to a dot here. And it's just the average return for household J.

On the y-axis and the beta, the regression coefficient on the village average is on the x-axis. So you can see that households that have higher returns are households that have higher betas for three of the four, potentially all four, although Buriram is not the greatest example. And I have shown this picture to guys working in finance with New York Stock Exchange or Swedish Data.

And they see this degree of correlation. And they're like, wow, that's really amazing. So the theory is working pretty well.

And this is the regression. Doesn't really add more. This is the regression from that step 1 with the corresponding-- sorry, step 2, it is, with the corresponding betas. And it's not significant in Buriram, but it is in the other three provinces.

Not sure if I've shown you this map. I don't think we did. Chachoengsao and Lopburi are relatively near to Bangkok. And Buriram and Srisaket are out in the agrarian regions of the northeast. So very different types of sample.

OK, so let's talk about idiosyncratic risk. The return on household J is determined by β_J that we've been talking about. X is to make the argument more general. So regression coefficient times the object, and then an error term.

So if you take the variance of both sides doing a variance decomposition, the variance of the return is the variance explained by the dependent variable with its coefficient plus the variance of the error term. So when we talk about idiosyncratic risk, we're talking about the error term here as opposed to how much of the variance is determined by the village aggregate. And so you can just imagine taking this variance term and dividing it through both sides.

And then, you'll have one like, 100% is equal to the contribution of the village return plus the contribution of the idiosyncratic variance to the total. So this is not theory. This is a way of quantifying how much aggregate village average risk there is and how much idiosyncratic risk there is in an individual J 's return.

Now, the next thing-- back to the regression-- as I have been saying, in the baseline, we're running the average return on household J onto a constant plus the β_J coefficient times β_J . And now, we do something extra. We add in σ_J , where σ_J here is the variance of this epsilon term.

How much is it moving-- oh, it's right here, except that I didn't call it σ_J here. The variance of epsilon J is σ_J . Stick that into this cross-sectional regression as an extra right-hand side variable. And get a coefficient on it.

So now, we can do a similar calculation. We can take the risk premium, which is the sum of both the return over above the risk free rate that's due to the aggregate return and also due to the idiosyncratic risk. And so we can decompose the aggregate risk as a proportion of the total, and likewise, the idiosyncratic risk as a proportion of the total. Question?

So doing that, we summarize the results-- the contribution of idiosyncratic risk to total risk and the contribution of idiosyncratic risk to the risk premium. The contribution of idiosyncratic risk to the total variance is huge. This row here running across the villages and looking at the quartiles including the median, you've got 98% of it, 97% of it, 94% of it.

65% of the risk is coming from the idiosyncratic component, coming from this guy. Let me back up. Coming from this guy. Second row of this table is a contribution to the risk premium of the idiosyncratic risk.

And it's pretty small. Two things to note-- it's almost always lower than the quantity. So the risk premium here is also small. Here, it's 21%, which isn't trivial. Or even here, 45%.

But the 45% is still less than the 99%. So the contribution to the risk premium in terms of the compensating returns required for the household to run those technologies is always-- the contribution to the total risk is always greater than the risk premium. Except Buriram screws up over here.

But as you've seen, Buriram screws up in other ways. And then, we do a robustness check. Question?

OK, back to the economics. And in medieval villages scattered around, random returns. Should they diversify? Do they do different activities, et cetera?

Well, that depends on the level of risk. And it turns out, there's a lot of risk, but a lot of it is plot-specific or activity-specific. Most of it is idiosyncratic and not aggregate.

It is not true that the returns are comoving much across all these activities. So it's a very risky place, as in India. But there are plenty of diversification possibilities by holding multiple activities, et cetera.

So we're back to a version of what we did in medieval villages, right? How much did they hold different plots? Here, we are asking, how much do they hold different activities and what can be accomplished by doing that?

And moreover, since it's idiosyncratic, you can hope that it might get-- the risk premium for the idiosyncratic stuff might get small. Anyway, all of this is from the production data. Let's look at the consumption data, and in particular, look at the gifts.

So now, we're going from medieval villages, so to speak, in terms of plots and production, to India in terms of ex-post consumption smoothing. The gifts and transfers are the ex-post part. It's how much they would give money to each other starting from-- do actually give money to each other-- starting from the yields on their production activities.

So if you look back here again, you would see these epsilons. So this could be high. This could be low.

If the variance is big it's moving the returns around. What happens when it's really big? Then, this household got lucky. It had a high return.

Just lucky. What does it do? It gives stuff away. Or another version-- they get unlucky.

The epsilon is low. What do they do? People give them gifts.

Well, I haven't actually shown your result. But that's exactly what is going on in the data. So if you regress those shocks, those idiosyncratic shocks on the right-hand side onto gifts and net lending, et cetera, you can see both. There is an ultimate impact on consumption, yes, but it's small relative to the way they're mitigating the risk.

Namely, when the shock is high, they're giving stuff away and lending it out. And conversely, when the shock is low, they're getting this stuff in. It depends on the sign.

This is net gift outflow when the shock is high. It's positive. It's statistically significant. You can add up both gifts and borrowing and lending and get a bigger coefficient here.

It's 40. So 40 relative to the impact on consumption, which is 4. In other words, for every unit of idiosyncratic shock, roughly 10% of it is making its way into consumption, which it shouldn't. It violates the risk sharing theory.

But 90% of it is being pooled away by this village gift-giving mechanism. Questions? In economics, we call this the smoking gun.

You theorize a lot. You think you know what happened. But what's the evidence? Well, they smoothed a lot. Must be the gifts.

OK, fine. If you think so. Well, here's the evidence. They're actually doing it.

And not only that, they openly tell us they're doing it. They are very aware that they help each other. I think you might have seen the video by now-- my movie, so to speak.

If not, take a look at it. Because it describes-- you'll actually see the setting of these Thai villages and how we've gathered the data, and so on. Michael is planning to show that at one point.

But it's available. It's a 32-minute movie. It's not so much of a burden. And you'll see how I spend 20 years of my life.

OK, so we've talked about in this. Now, let's come back to health shocks. No, it's not COVID.

But it is an illness experienced by the household. And it's not covered by insurance. And the expenditures go up.

So this is normalized so that the event happens here. And health spending is big, and other things like, they say they're sick. They're reporting symptoms to us.

We also measure that. So we're pretty much sure that we nailed the episode, et cetera. And sick days goes up and stays up for a while before it comes back down again. OK.

And what happens? There come the gifts. Again, centered around the episode, the event. As soon as the event happens, it's actually minus 1. As soon as the event happens, the gifts are positive and statistically different than for households that did not experience those health shocks.

Here, we're comparing placebo households, if you like, which is a dangerous metaphor given the illness part of it and the COVID. We're taking households pairwise conceptually-- one set have experienced these health shocks, and ones that have not, and looking at the difference in the gifts that they get. So if they're sick, they're much more likely to be getting gifts than the others.

This is the same setting as what I showed you in rates of return, except we've changed the subject from shocks to production to illness shocks. Now actually related, they're still producing or trying to. And the fact that they've doubled, say, their health expenditure is a big problem for them. Because they have only so much liquidity.

And they're supposed to be buying inputs and so on to fund their agriculture and so on. So this is a liquidity shock. It's not just a health shock.

In fact, sometimes, it's an elderly person in the household that is not even contributing to production. So we're looking at the economic impact of someone taking money away from their budget, essentially. OK, there are various kinds of networks in these villages.

This is the kinship network, who's related to whom. This is the financial network in terms of giving gifts to one another. This is the sales network, who they sell to in the village.

And this is who they hire from. So the way these pictures are produced with the software, roughly speaking, households in the middle are connected to many other households. Households on the fringe don't have that many connections, depending on the variable.

And the gift giving financial network is different from the sales and production network. And I'm going to show you what happens because of that. I don't have a slide to show you this, but if they're in this gift giving network measured by how many times they got gifts before the shell shock happened, then most of the impact of the shock is covered by incoming gifts.

But what do we have in the US? This is an issue. A household gets sick. Is the government going to be able to help them?

We have these PPP Paycheck Protection Programs and so on. Are relatives helping? Probably so. We don't know the extent of it.

It's not measured. We tend to think households are standalone. They're vulnerable. Households running businesses in particular get a liquidity shock.

And mainly, they don't have sales with COVID. Or their employees aren't able to work. So they're not generating revenue.

Well, if it were covered with gifts, as it is for some people in Thailand, then they pool that risk away. Great system. If they're not, there are other consequences.

So in particular, we can look at who's connected to whom in terms of the network. Households who experience shocks and were selling to other households-- their sales go down. It's a bit counterintuitive, because you would think they would want the money from the sales.

But the problem is, they had to reduce the inputs. They had to scale production down. And so they have less revenue as a result.

Sales are down. Inventories are up. Well, I should-- let me come back and clarify. These are households connected to the household that has experienced a shock.

These guys were selling-- I said it wrong. I'm sorry. These guys were selling revenue to the household that had the shock. Now, they sell less.

Likewise, because they sell less, they're accumulating inventory. And they're holding onto inventory longer. This is this inventory turnover ratio.

That's a measure of-- I'm going to [? tie it-- ?] It's a measure of inefficiency. It's also negative with an adverse effect. We do this for all the shocks and do it for the large shocks. So now, we see the COVID connection, which is, a household experiences expenses associated with the illness, even if it's not a working member. How do they cover those expenses?

They try to cut back. They cut back on spending. In these village economies, businesses are connected.

They're buying and selling inputs from one another, and selling output. So if they have to cut back and they don't have the revenue that they had before, the people who had been buying had been selling to them no longer have the sales. Just like one business is connected to another, they all go down the supply chain, because they're connected.

We had a lecture at the end on interconnections among households. We did Leontief. We did the input output matrix.

I showed you Fukujima and Shima, and the effect of the earthquake there. And those were all supply chains. One sector is connected to another.

One firm's connected to another. These Thai households are connected to another. And the point of this-- I guess I haven't said that-- is, yet, the closer they are to the adversely impacted household, the more they suffer.

So you think networks are good things. Networks are a good thing when it comes to financial gift giving. But if gift giving doesn't cover the adverse shock, we're actually getting an amplification in this economy that's showing up in terms of lost wages and less production, and less sales.

And finally, this distinction, which I've been pushing today, about aggregate and idiosyncratic shocks-- go to one of the areas, Chachoengsao, where they have fish ponds, shrimp ponds in particular. And it's pretty polluted.

And they didn't have green agriculture, so to speak. And the EU put a ban on the shrimp. By the way, if you go to Walmart-- you say, villages? Who cares?

Well anyway, they produce the shrimp that you buy in Walmart. But the EU put a ban on the shrimp, because they say it wasn't clean, which it wasn't. Now, the point is, most but not all households are doing those shrimp ponds.

And the ban is like an aggregate shock. So we can compare the aggregate shock, which is most of the households, to the idiosyncratic illness shocks. And you can see that although gifts go up-- I already showed you that-- for the idiosyncratic shocks, including in this village, there's not that clear picture when it comes to the shrimp shock.

Because it's an aggregate shock, there's not a lot of ways-- excuse me-- to reallocate that risk. And this is the associated decline in employment, which is bigger under the aggregate shock. OK, questions?

OK. That's what I have for today. And we'll continue on Thursday. Thank you for coming today. Talk to you soon.