ROBERT TOWNSEND: That's the starting point, but we're going to then move, in the latter half of the class, to what to make of it, what happens if you find that it's complete or incomplete, how to think about policy through the lens of this model, how to design insurance instruments, what damage they actually might do, what the distribution of benefits might be. There is a recurring theme today, which wouldn't surprise you, knowing me, which is theory is guiding the whole discussion, rather than starting with some randomized control trial and targeting, but we will get to those control trials in the discussion.

OK. So here's the outline. We'll look at risk and insurance and village India, move to a benchmark way of measuring the effectiveness of financial institutions on the formal side. We'll then couple that with work on the informal side and networks and so on. We'll get to this-- can well-intended interventions actually result in well welfare losses? And then issues of why insurance take up in many circumstances is actually quite low.

The idea here is that people in village India, and developing economies more generally, live in very high-risk environments. There are ways to cope with this risk-- ex-ante and ex-post. Ex-ante, you can diversify your crops, fields, activities, and occupations. Ex-post, you can have financial transactions, gifts, loans. I actually mentioned this the first lecture.

How well are households doing? So the idea in this paper is forget enumerating all the mechanisms. You're likely to miss one anyway and the measurement might be really bad. Let's just look at the outcomes. The null hypothesis-- If risk sharing were perfect, what would you expect to see in consumption and income data? I think this style of work, although nowadays it's not only accepted, it's highly utilized all over the place, but initially there was a confusion that perfect means complete as incomplete markets, and somehow you're saying you're going to have all these Arrow-Debreu securities being traded in a village economy, and no, that's not the idea. Informal institutions can be at work, as well as some spot market exchanges. You don't need to take a stand on any of that if you're just going to look at the outcomes.

So the idea is individual consumption should actually not depend on individual income, stunningly. Once you control for aggregate consumption, it's as if all these households lived in a big risk syndicate, and they were just pooling all their grain and so on in one big pile, and then depending on shares, and some people are more equal than others, they get a fraction of the grain. A pile of the grain can go up and down with village aggregate shocks, but controlling for those shocks, your contribution has nothing to do with what you take out.

So where to test? We're going to do these tests at the village level in this starting point of this lecture, but we'll move. It could be kinship networks within a village, family related people. It could be bigger than that. It could be a cross-village risk sharing. It could be that being a member of a financial institution, for example, is allowing you to be part of a larger risk pooling population. So it's a judgment call about where to apply it. Fortunately in the work today, I can show you it being applied at various distinct levels, which is a theme we'll come back to next Thursday as well.

Why did I choose villages? I had some idea that everybody knows each other, and the environment is relatively simple, and it just seemed like the idealized setting. Of course, in practice not all villages are alike, and they're pretty complicated, and you've got people interacting with one another, and they don't all love each other. But that was the motivation to begin with.

By the way, if there's no idiosyncratic risk there's nothing to test, because the benchmark says controlling for the aggregate is the idiosyncratic influencing consumption or not. And if there's no idiosyncratic, it's all aggregate, you can't do that test. This extends, of course, beyond households within a village to tests across villages and so on.

Now, many people have the idea that villages which they may never go into are all alike, and it's all a function of the rainfall, and it's all aggregate, and that's just not true at all. And I'll show you the front end of this paper, some statistics. Although it is true there may be a lot of diversity in a village, if they all completely diversified and did not specialize and they all held a common portfolio, then, again, there would be nothing to test as if they're doing it all ex-ante and not ex-post, because the starting point would be the income coming from all these diversified activities. But again, they don't do that.

So let's take a look, these so-called ICRISAT villages-- Institute for Crops Research in the Semi-Arid Tropic-- at one point that was the source of virtually the only panel data on developing countries in the world. And it was gathered monthly, and essentially for 10 years. They've started it up again, interestingly. And there are other data sets like that, where there are gaps of decades. So it's not true that ICRISAT was the only source of panel variation.

In Aurepalle, they're growing castor and sorghum. Each crop is quite risky if you look at the coefficient of variation. Those numbers are comparable to something I did in medieval England, where roughly every 12, 13 years, they would experience a famine. So these are high-risk environments.

Of course you could diversify across these crops. That helps. Depending on the pair of crops that you pick, the cross-crop correlations could be as low as 0.09 or something much higher. A low number is a good thing, because you can spread your risk. Soil isn't uniform either. You can look, for a given crop, at how the coefficient of variation varies from one type of soil to another. And again, the cross-soil correlation for a given crop is 0.37, and relatively low.

Nevertheless, as I said, it's not true that all households hold a diversified basket of all these things. Yes.

AUDIENCE: Are there any crops that are negatively correlated with each other?

ROBERT TOWNSEND: Not in these numbers, although I certainly wouldn't rule it out. Well, and some of those numbers are point estimates, relatively low. And you put the standard error bands, and yeah. OK.

This is by cruder categories, like labor income, income from trade and handicrafts, animal husbandry. And it's stratified, as ICRISAT did, by the amount of land that they hold-- none, small, medium, and large landholders. And these are sort of basically fractions of income that they get from such things. Large landholders have the bulk of their income from crops, and a substantial amount from animal husbandry. People without land virtually have no crop income. And the dominant source is from labor. And there's everything else in between, as you might anticipate. There were three villages involved. That pattern is pretty common.

And here, again, by those income sources-- (EXCITEDLY) hey, a negative number-- again, in the lower bound of the stair of the confidence interval. Well, actually this is-- even the point estimates went negative.

Sorry, no no. Yes, yes. This is hard to read because these items should also be listed on the rows. So if you're going down the diagonal, you're reading the coefficient of variation, and then above and below the diagonal, you're reading the correlation of, say, in this case, livestock income with profit income. It just throws you off, because it should be written down here. And then to avoid double counting, you don't fill in the bottom of the matrix.

So I've shown you this picture before, the Rocky Mountain picture, arguably the most famous graph I've ever produced. This is 10 years of data. And then the households are lined up-- which is why I keep showing it-- the households are lined up roughly by the size of landholdings. And you can see the levels kind of going up as they have more and more land.

And there's a 0 here. And so everything on this diagram is deviation from the village average at a point in time. Some people are below and some people are above. You're seeing that in the cross section, and you're seeing that over time.

The point of the diagram, I've said before, is that some of the valleys lie behind the peaks. And that means not everyone is experiencing good years and bad years at the same time. It's not like undulating waves of grain, as in Kansas. This is the consumption picture that is on the same scale as the income picture.

If I wanted to get bigger differences, I could rescale this graph. You know, it's all just suppressed down there. But the statistical work is really testing to see if there are common fixed effects, and also whether this diagram somehow, through the lens of the model, is related to the income diagram previously.

AUDIENCE: Have you observed this in terms of like greater weights on people? Because it looks like there's not only risk-sharing, but also the rich guys are-- if it's the same [INAUDIBLE].

ROBERT TOWNSEND: Yeah, exactly.

AUDIENCE: They're not just getting the average.

ROBERT TOWNSEND: That's right. So risk-sharing is sometimes misinterpreted to mean complete pooling, and everyone has an equal share. And that's not necessary. And when we look at the benchmark regression, essentially it has fixed effects. And those fixed effects are almost exactly the Pareto weights. And then we can test and see what those weights are related to.

So here's the optimization problem. You just basically-- there is a one-to-one relationship between all the Pareto-optimal allocations in this economy and solutions to this programming problem. And that's asking about the Pareto weights. Here they are-- lambda sub-k, different for different households, k. The household is the basic unit, although we will control for the demographics within the household.

So just to read across this, it's a weighted average, lambda, Pareto-weighted average of the discounted expected utility of, say, in this case, household k. OK, so it's over time, which could be finite or infinite. It's discounted at a common rate. I'll come back to that.

It's presumed everyone has equal probability over these states of the world. And these states of the world are not just the contemporaneous realization of all the shocks. It's the entire history, from the initial date out to that, like the branches of a tree. And then utility contemporaneously depends on consumption and these gender/age weights.

So and you just want to-- whatever sums of consumption is, it's this average and the aggregate. And then the idea is how to distribute that aggregate to maximize this subjective function.

You can put labor in here. We'll have a whole lecture on that. So what is the solution? Let's just pick consumption of household k at date t is entering in here, and it's entering in this resource constraint. This supplies, you know, for the whole history of states up to date, t, for every date, t. So we're going to get derivative here, and a Lagrange multiplier here. This is common resource constraints. So you pick up a common Lagrange multiplier. Yes.

AUDIENCE: Does the [? ht ?] include the use of [INAUDIBLE]

ROBERT TOWNSEND: Yes. Yep. A complete enumeration of everything, and not just crops, but babies, and deaths, and the whole demographic structure.

So again, taking that derivative, putting this thing over here, you'd get a beta-discounted probability, lambda k weighted marginal utility. And everyone's margin utility is being equated to this common Lagrange multiplier.

You know, so if you assumed a particular exponential utility function, for example, then you can actually get sort of closed-form analytic solutions for these first-order conditions. So this just says, in levels, that the consumption of household k depends on the Pareto-- the log Pareto weight of household k. It has something to do with the demographics and something to do with this common shadow price of consumption, if I dare call it that.

So here are the levels that we were just talking about. This is too small to read, but basically-- another nice thing about that exponential utility, if you're willing to make strong assumptions about the composition of the household, is it basically would dictate how to weight people. And little babies don't eat as much as 18-year-old males. And in turn, consumption is dropping for the elderly. So we actually had a dietary survey. And we assumed that strong functional form, and we made all those corrections. So you don't actually see people eating. So you don't see their individual consumptions. You see the consumption for the whole household, k, but you put it in per capita terms by dividing by the age/gender-weighted number of people.

Likewise, this thing over here is the sum in the whole population of per-capita consumption. And sort of reading out loud, this term is actually the risk tolerance, the inverse risk aversion, of household k. It's actually a j here. Sorry. Household j relative to the sum of the risk tolerances of all the other households. And I'll come back to this often.

If you assumed homogeneous risk aversion, which is standard, then this ought to have a coefficient of 1. So everybody's consumption is co-moving, one to one, with the per capita average.

If you have different risk aversions, there is a very nice interpretation here, which is this is aggregate risk. This is what's left over after all the other smoothing that they're doing intertemporally, blah, blah, blah, than what they actually eat. So this is like the macro risk. And then the higher this risk tolerance, the more they should be willing to bear that, quote, "helping" out their village neighbors who might be quite risk-averse. And we'll come back to that in terms of measuring risk aversion.

OK, so here's the standard regression equation. If the level of consumption depends on basically what is now a constant household-specific term, which are those Pareto weights, sum-- coefficients, time, village average consumption-- at date t, adjustment for demographics, and something else that ought to be zero, like income.

You're going to have a devil of a time running this. Even when you run one household at a time ignoring all the others, even when you use the leave out mean-- so the household is not included-- it will turn out, amazingly enough, that these coefficients, one for each household, when you average them up, have to take on the value of 1.

So there are sort of econometric issues here that you could mistakenly think that you've got a lot of co-movement, and it's just going to happen-- you know, why is this happening, essentially? Well, it's more obvious if you had the panel in front of us rather than one household at a time. You're looking at household consumption on the left, and the average of household consumption on the right. Now take the average of the whole thing, and the averages appearing on the left and the right, and it ought to take on a value of 1. So the average tendency of the dependent variable is kind of like an intercept. So that's kind of the intuition. This also comes up in many situations where you have this confounding problem. Yes.

AUDIENCE: [INAUDIBLE] talked about [? instead of their ?] [? compatibility? ?]

ROBERT TOWNSEND: No, because we're going to assume full commitment-- instead of compatibility? Did I hear you correctly?

AUDIENCE: Yes.

ROBERT TOWNSEND: Yeah. We can put in incentive constraints. And we will later. So we're going to build on this basic framework, and get to all the obstacles to trade, and how to estimate them. But you may think this thing would be rejected out of hand. And it is statistically rejected most of the time. But it turns out to be a very, very useful benchmark.

So this is the time series, one household at a time. I won't say much about this other than, does the data have any power. The problem is, you test one null, like full risk-sharing. And you can't reject it because the confidence intervals are so large. So you want to test other things like its autarky. And the weight of the evidence is it's not autarky. We can pretty much soundly reject-- And risk-sharing is doing quite well.

And this is with a panel. And again, I'm going to spare you all the details. But we did it three different ways-- standard, first difference, Jerry Hausman-- and estimated. Now, here we're picking up coefficients on various sources of income, none of which should be positive and significant. But some of them kind of are. For example, the highest number here is profits from trade and handicrafts in Aurepalle. They're climbing palm trees and tapping-- and then making liquor, toddy. And it's not insured, evidently.

Labor in general has a high number pretty much in all the villages. So now we're kind of easing into the policy thing, which is, if you wanted a target, wouldn't you want to run this thing first and see who's the most vulnerable, and then sort of guide your interventions. Based on that, I wasn't thinking along that line, I was just being honest about where we're rejecting and so on.

And here, finally, Matt, is where we took this coefficient and regressed it against other things like the area operated by the household, the value of their plow animals, their inheritance, and so on and so forth. And I'm not sure what you would expect. It looks good at first blush, in the sense that the amount of land and the value of the animals is correlated with those intercepts. The problem is things are more turbulent than that at the village level. And people do sell, or buy, or acquire land. And certainly those plow animals aren't just locked in place year after year after year.

So that's kind of some evidence of some kind of long-term impact that shows up in terms of the assets they end up holding. And that's very much along the lines of, is there a moral hazard problem or some kind of incentive constraint? And how should wealth levels behave if that were the case?

The biggest puzzle in these data, in some sense, is almost the opposite, which is although we see them over 10 years, we see the upper-caste, white-robed brahman-type guys in the village just consistently eating far less than their income, constantly. And so is this just a social transfer mechanism, or are they actually insuring themselves against future relabeling of-- castes, by the way do get relabeled, and histories get recreated. It it's a much more complicated history than you might-- at least that's what the anthropologists tell me.

OK, so that's risk and insurance in village India. So let's take that framework and start knocking off some policy questions. So first is about the effectiveness of financial institutions. Now, let me just say, the immediate policy implication coming out of village India was basically, miraculously, somehow, they're pooling all that risk. It's not a bunch of fragmented households on their own, in autarky. It seems to be more the opposite. So there was a huge firestorm people had a hard time accepting that something like that could happen. But villages, even if they were amazingly good not, are not necessarily interacting with each other. I'll come back to that. And there may be, as there is in Thailand, transactions with outside financial institutions.

So the question is, do outside formal financial institutions help, if not within the village, then across village, in the mitigation of risk? So roughly speaking, the treatment group is a set of households who are customers of financial institutions. A little typo there. And that's measured-- their use of the institution is actually measured in the data, which is reassuring. And I'll tell you about the instruments we use momentarily. So we're going to go back to that consumption equation, and also derive a corresponding investment equation, and then see whether being a customer or a member lowers the vulnerability in the sense of that coefficient on idiosyncratic risk. Yes.

AUDIENCE: So when we say those with financial access are compared to those without, so does that really mean those with access to the formal financial sector are compared without--

ROBERT TOWNSEND: That's what it means.

AUDIENCE: Yeah, because there's money-lending sectors et cetera.

ROBERT TOWNSEND: Exactly.

AUDIENCE: Yeah.

ROBERT TOWNSEND: So yeah, that's like one topic away. I'll come right back to that. But to answer your question, it's going to turn out that directly connected, at least in consumption, is helpful, but indirectly connected, through gifts and loans to someone in the village, even though you're not, is as effective in consumption-smoothing. So that's then consistent with the previous paper, although it's not quite so clean with investment.

So what are these institutions? Well, by now you probably know the names of them. They've been coming up in previous lectures. We find that this bank for agriculture and agricultural cooperatives is helpful in smoothing consumption and investment. It looks, from the transactions data, is if that has to do with its credit operation. They're basically embedding credit and insurance in their loan contracts. It's not, "you pay back the loan or default." There's a whole schedule for repayment. And it's adjusted by aggregate local/regional shocks, and so on. It's quite sophisticated.

In the financial crisis, the World Bank went in there determined to see this stereotype everywhere, since remember, finance companies and commercial bank failures triggered the Asia financial crisis. So they looked at the BAAC history of over arrears, and concluded that it was another disaster. But actually they were running an insurance operation. And that's what shows up in the consumption data.

Commercial banks are helping to smooth somewhat. It's largely, though, through formal savings accounts. Loans are pretty thin in the data.

So what are we going to do? You've seen this model before. This is Greenwood and Javanovic. By the way, we're sort of in the data micro part of the class. Don't forget about all that macro. So the Greenwood-Jovanovic model assumes that the benefit of being in a financial institution is better information and risk-sharing. So here we are looking at the data to see if that's a good assumption or not, actually institution by institution. So there is a direct, immediate link of this micro work with those macro models. Or even if we reject, then it should inform what we assume within those macro models. So remember, in Greenwood and Javanovic, there was a cost to joining the financial system. We're going to let this be a household-specific cost. And basically we're going to want something like instruments, like distance from the branch, which varies across households.

Just a little bit on the underlying model, just to remind you, because you've actually seen it before. You've got output, which is a function of the current capital stock, and then the sum of an aggregate idiosyncratic shock, aggregate at the whatever level, regional/national level, idiosyncratic at the level of the households. As I said, we're going to embed investment in this, adding that to the risk and insurance in village India. So there's going to be some adjustment costs to adjusting your capital stock. This is also buffeted around by IID shocks.

This is new. This is a preference shock. We're going to put something in the utility function which is autocorrelated.

I'm actually bending over backwards to create a bias here. What we're going to want to be able to do an instrument is to say that the instrument has to do with, does influence the selection of being a member of a financial institution, and conditioned on that, does not influence the outcome, in the future, at least. But you can imagine, with autocorrelated preference shocks, that part of the decision to join the financial institution today has to do with that urgency or patience. And if that carries on into the future, it will also affect the endogenous variables. So we want to make it hard, and then we want to get rid of the problem, to show that we have the ability to get rid of the problem.

You may remember how this works, though. You create a value function for being a participant in the financial system. If you join, you subtract off that fixed cost. And this is just the array of the other variables. You're joining after the realization of the adjustment cost shock, the idiosyncratic and aggregate shock, and the preference shock. You could choose not to join. You'd have another value function. I think we even used V and W in the earlier lecture. So but you don't subtract off the cost from the capital stock. If you saw the autarky problem, then you've kind of already seen the solution to the consumption part of the intermediated problem. You can just sort of list the equations for p equal 1, participant. They've got the V value function. You've got consumption as a function of your preference shock, and aggregate consumption, and your Pareto weight.

By the way, you haven't seen this today, but investment is also a function of the adjustment cost, the capital stock, and the same aggregate consumption. Because it's the same resource constraint. So the price of consumption today, actually, versus tomorrow, is not only influencing consumption, it's influencing investment intertemporally.

And we'll come back to that and repeat it in the lecture on Thursday. And this is just the autarky equation. You don't see aggregate consumption anywhere. People are on their own, and so on. There's no way you could see and totally understand all the individual variables at this level. I'm going too fast for that. But I just want you to think about what you see in the data. You're going to see kind of a weighted average of these things, weighted by participation. So actually you're either in this branch or you're in this one. And being an instrument for these Ps.

But anyway, it's convenient to write it down this way, and derive the sort of econometric equations you're going to take to the data. I actually take first differences, but the same principle applies. These instruments are the history of the person's institutional use. Like predetermined variables is an indicator for their current participation at the village level, actually, not at the household level. The time to the district center is the most compelling. That's like how long does it take to get to the bank. And then we have some GIS measure of access in the local area, which is supposed to pick up, like, their active credit officers in one spot and some not-so-active credit officers in other areas.

And then the data you know, this was from May of 1997, in this case, when we wrote this paper, through 2001. And I already gave you the conclusions basically. And in fact you're going to see the equations again when we get to this informal, because it's very similar-looking no point in repeating.

So this is what I've been doing with Cynthia and [? Whit. ?] So we have the measured chains of financial transactions, gifts and loans, from one family to another. So we've already talked about the fact that people with indirect connections who are not directly customers but are connected through other households are going to do quite well. There's a flip side to that, is we've been measuring the control group. The control group was an average of people who were indirect previously, indirectly connected or not connected at all. And when we separate that out, we have the "not connected at all" people. And it turns out they're actually the most vulnerable. So their vulnerability was being concealed. And there's some work in Mexico that is very, very similar in spirit. The data is a bit different, so the details are a bit different, Andgelucci, and Rangel, Rasul, and so on.

Investment's going to be a bit of a different story. Kin are going to matter, but it's almost like an out-of-equilibrium story. Let me come back to that. I won't dwell on this slide. It's just showing you that there is a lot of borrowing, and lending, and gift-giving. And when it happens, it happens typically, you know, depending on which line you're looking at, three to five times during a sample. And the amounts are large. The borrowing transaction is like 60% your monthly income. So you know, take your student stipend and-- it's a big number. No, I'm not saying your stipend. It's a big percent.

AUDIENCE: [CHUCKLES] We would never make that misunderstanding [INAUDIBLE].

ROBERT TOWNSEND: [CHUCKLES] So this is the equation I was alluding to that I didn't show you on the previous paper with Alem. But basically, first-differenced, we have the change in household consumption in village v at date t. This is the change in the income of household i in village v at date t. This is the vulnerability coefficient that we're interested in seeing how big or small it is.

And then we start interacting that change in income with basically a direct member or "indirectly-connected," quote, member of a financial institution. And one hopes that this might be not trivial, but this would be negative and subtract off for those who are members, and customers, and so on. K is kinship. That's also entering-- we're trying to distinguish sort of having family in the village as opposed to having measured transactions. It's a bit heroic. We don't have instruments for whether you're in a network or not. It's a problem for us and for a lot of this literature.

And then this is controlling for net worth. I mean, as you've seen from the previous slides, the relatively poor people, at least in India, were more vulnerable. So knowing that a priori, we want to take that effect out and control for it, at least in this additive way.

This is the investment equation. It's basically very similar. I is just investment. It's actually normalized by the capital stock. So it's investment per unit capital. And again, it allows for the direct effect of the income per unit of capital, as well as the direct and indirect connections. I should have highlighted that, in both of these equations, you've got these time-varying calm and fixed effects. This is what the theory is telling us to do. It's not as though consumption doesn't move with income. Consumption can move with this fixed effect.

I should have also said, previously, with my mention of the econometric problems of measuring average consumption and putting it on the right-hand side, that you can just replace that with an unobserved fixed effect. And the theory gives you the interpretation of what it is. Or you could subtract it all off from the left-hand side. And when the coefficient is 1, that's also a reasonable thing to do. Here we're just putting in these time-varying fixed effects.

So this is the coefficient for those people without any access, direct or indirect. This is how much you'd subtract if you look at direct bank-- yes, these are very similar numbers. I always get accused of biasing this stuff.

Anyway, so the sum of these things is significantly different from zero. And the point is you've achieved the benchmark standard. But equally interesting is, when you have the indirect connections, you're subtracting off virtually the same number. Kin didn't help in this case.

AUDIENCE: Sorry, is there a way to augment-- because we have direct connection and indirect connection. Would it help me more to have both?

ROBERT TOWNSEND: That's what indirect means. Sorry.

AUDIENCE: Oh, OK. All right.

ROBERT TOWNSEND: Yeah, yeah. It's mislabeled. It means both.

AUDIENCE: Oh, all right.

ROBERT TOWNSEND: Yeah, good point. And then this is with investment. Now, what are the differences? Well, two, basically. First of all, whatever we're subtracting off, it doesn't lower the total to zero. And if you tested, there's still residual vulnerability. So with investment, smoothing investment against cash flow, you are left with some idiosyncratic risk. Although kin are quite helpful, but measured transactions are not.

It's a little too early to get into this story. Cynthia and I are still working on it. But the idea is something like, just to anticipate when we get to obstacles to trade, you can't force someone to be in a risk-sharing group, but they kind of like being there because they can smooth their future income. So you can't force them, but they voluntarily choose to stay even though, in any period, they have the option to leave. It's, like, self-sustaining.

With investment-- and this is above median investment relative to the capital stock-- investment is big. The typical number is like 20% of your physical capital when it happens. So it's not little, incremental, teeny-weeny stuff.

OK, now imagine, say, someone comes to me and asks me for a loan-- I'm the community head, in case you didn't notice-- and you're going to put it in your project, and it's quite a good project. And then, next period, you got a lot more income than you used to, now it's much more tempting to renege. Because you're much better off in autarky than you would have been if you hadn't gotten the investment done. And so the theory is, A, that either doesn't happen, or B, all your relatives come and do their thing. And that's consistent with these patterns when we're working on numerical estimates. Yes.

AUDIENCE: I'm not sure if you can see this from this data, but if I'm talking indirectly, accessing the financial system, is that link between us, does that person typically have assets on stock, perhaps, because they know people will borrow from them, and they want to do it? Or are they actually just turning around and then borrowing the pharmacist, and then being more like an intermediary for a loan.

ROBERT TOWNSEND: I'm not sure exactly what you're asking. [? Whit ?] and I have this-- somebody will take a formal-sector loan. Then it's due. And they're not doing too well in that year. So they will turn around and get an informal short-term loan from a friend, or relative, or money lender, pay off the long-term loan, and then with the proceeds of the long-term loan, they pay off the short-term loan.

AUDIENCE: And that friend or relative, should I-- maybe the answer is it's both-- should I think of that person as having money lying around to be ready to lend to their friend who needs it, or did they intend to--

AUDIENCE: They can also take it-- there could be a lot of shame. And they could also get it from the financial institution themself if they're a normal guy. But there is some people who are acting as money lender, and they have money ready [INAUDIBLE]. So it depends on the person. If it's a relative, mostly they're not that rich. So they'd probably borrow.

AUDIENCE: I guess I'm just speaking of one of the things the portfolios and the [INAUDIBLE] book, is that people will simultaneously borrow and save at the same time. And maybe this fits in, if part of the reason I have that savings is because I might need to lend it out at short notice to somebody.

ROBERT TOWNSEND: Yeah. So there are all kinds of financial strategies. You know, savings and borrowing gets pretty complicated. Because even if you get a loan, then the way you get the loan proceeds is someone opens up a savings account. You don't necessarily withdraw it right away. So that's one mechanical reason to get that.

But loans can be long term. And then you have a liquidity problem. And you don't call in the loan. So you might want to keep savings to balance that risk. I mean, this is like the bread and butter of a lot of things people do in macro when they're looking at investment and trying to sort all that out. Yep.

AUDIENCE: Why [INAUDIBLE] smallest investment?

ROBERT TOWNSEND: They want to undertake the investment when they have an opportunity. And they may not be able to do it, or they may have to wait to do it until they have the liquidity. And that's costly.

AUDIENCE: But [INAUDIBLE] differentiate [INAUDIBLE] tp avoid two cases. The first one is that there is insufficient insurance, so I cannot borrow money from the bank. The second one is that I [? just ?] [? want ?] to invest because productivity [INAUDIBLE].

ROBERT TOWNSEND: Yeah, so we'll get to this structural model of village funds, which takes a stand on the process that generate shocks. And then we have to take the data, and kind of estimate that process. And then that model allows the-- and then you get these investment opportunities which are like random variables, large, chunky opportunities.

AUDIENCE: Here we assume that we don't have to differentiate [INAUDIBLE]

ROBERT TOWNSEND: This is reduced form-- I mean, here, I was talking about a model just then. But we're not quite ready in the class, because we haven't really explicitly written down these obstacles, neither the moral hazard or the limited commitment. So I was talking you through it like what next. But actually this is just reduced-form stuff. It doesn't really tell you. But we're telling a story which I think is reasonably accurate. But Cynthia and I are writing down a model, and trying to numerically compute it to see if it's a reasonable story.

This is what I said in words. And then this is the Mexican paper. So it has to do with Progresa. And they have treatment and control. And the details are a bit different. But I think you have these slides. And there's only like four or five of them. But I will just say, then, it's not just about Thailand, and we're not the only one doing this.

Networks is a really hot topic. Everybody wants to play around with networks. Esther and [? Arun ?] have pretty good network data in parts of India. So one example probably is enough.

Now, I said something about well-intended interventions causing losses. This is some joint work We first of all do estimate heterogeneity and risk preferences. And then the idea, splicing together these papers, it's like village India. So there's aggregate village level risk that's not necessarily shared very-- which you have to control for, and it's not "shared," quote, unquote-- let me say it again.

You've got all the idiosyncratic stuff going on within the village. So it's as if almost complete markets there. Then you village-level risk, which is aggregate risk. And then someone comes in ostensibly trying to insure that aggregate risk. So you would think, oh great, they're already covered on the idiosyncratic side. Now someone else can come in and cover them on the aggregate side. Must be a good thing.

And the answer is, well, no, not necessarily. I've already basically given this away. Because you have risk-tolerant people in the village who were willing to underwrite that aggregate risk. And it shows up in higher levels of consumption for them, on average. It's not as though people wrote a check, and you got the premium. But it's showing up, in whatever mechanism it is, in the give-and-take with informal transactions and so on.

So to see how we're estimating it and get some intuition for how we estimate the risk aversion, if two households had more or less the same coefficient of risk tolerance, then their consumption should more or less move together. But the opposite, if one were very almost risk-neutral, and the other very risk-averse, then, practically, consumption is almost smooth for the risk-averse guy by virtue of the risk-neutral guy underwriting that risk.

So you can imagine-- and people have done this-- looking at households pairwise, and just trying to see how many consumption reversals there are, and things. We don't quite do that. We don't have enough data. And I'm not sure it's very effective econometrically. We do take a household, j, take their measure of co-movement, essentially, compared to the co-movements of all the other households. So it's always j, less time, compared with all the other ones. Well, it's probably easier to do it with the notation.

So here is that basic first-order condition. With constant relative risk-aversion, by the way, it's in logs, not levels. But it's the same idea. You've got the Pareto weight. We're actually going to allow differences in discount rates across households. And so that sends sets up a trend, where some people are going up and some people are going down long term. We allow preference shocks-- actually monthly shocks, because there are holidays and things. So the idea is to put as many controls in, in some sense, that the theory would allow. And then we have this sort of common fixed effect.

And note in particular that actually everything, but in particular this thing, is normalized by basically the risk aversion. If risk aversion is really high, this coefficient is really low, and so that household's not moving around with the aggregate. So basically-- sorry-- almost what we're doing is taking consumption, regressing it against trends, and monthly effects, and household fixed effects, and getting a residual. The residual is capturing all of this-- call it v or eta, and then just sort of see. What the correlation is between a household i's residual, which is a stand-in for that aggregate that's moving consumption, and any other household's i-prime. So this is like a moment, basically. And we can use method of moments, in principle, for all the households, pairwise. But we didn't have enough data for that, even though we have a lot of households and a lot of months. So we just summed that we have household i's risk tolerance-- essentially what this thing turns out to be is household i's risk tolerance times the sum of the risk toleri of the other households in the village. But hopefully you're getting some intuition for where we're getting the risk aversion measures.

This, I think, you can ignore today. It has to do with whether or not the conventional regressions that ignored heterogeneity are biased in some way. And early work-- actually Sam's work-- suggested that it's even harder to reject full risk-sharing if you allow for the heterogeneity. Or it could seem that you're rejecting full risk-sharing when you assume homogeneity, because you haven't allowed this asymmetric adjustment that's possible on the part of households, depending on their risk tolerance. And that's all I want to say, probably even false intuition. Because it may not be true in general.

But then you kind of want to measure welfare gains. OK, so how much would a household be willing to pay living in a risky situation to remove that risk. In other words, what fraction of average consumption would they be willing to take a hit on and live in a world without the risk, as opposed to the world that they're currently in. Or how much would they be willing to pay for someone to come in as an outsider targeting insurance against, say, rainfall?

Lucas did this with business cycles. So this number, k, which is the fraction you're willing to give up, if positive means that you gain from eliminating aggregate risk, but k can actually turn out to be negative. Here is an expression for it. Suffice it to say that it has to do with the risk tolerance. And although this slide is a bit tiny, village by village, this is 0. So basically we're lining up the households by their degree of tolerance, from very risk-averse to not very. And these are downward-sloping lines. So basically the higher this risk tolerance, the more likely it is that this k or this welfare gain is negative.

So it's not trivial. And by the way, it's not correlated with characteristics either. You say, oh, it's the rich people. No, that's wrong. Risk aversion is not correlated with wealth. It's as if complete markets, remember?

AUDIENCE: Could it be due to [INAUDIBLE] within the village [INAUDIBLE] institutions. So two equally risk-averse people, but one [INAUDIBLE] so they [INAUDIBLE].

ROBERT TOWNSEND: We actually initially did this just on the network households, because I was worried about having the non-network households in there. Although when we actually run the full test now, for risk-sharing, taking into account the heterogeneity, we don't reject either.

I don't totally believe it. I still think there are relatively poor people who tend to be a little more vulnerable. But somehow, with all this heterogeneity, it's not showing up the way I would have thought. It's not soundly rejecting full insurance.

AUDIENCE: From the data, can you see, then, if the risk tolerance is greater amongst-- is there a distinction between risk tolerance and [INAUDIBLE] network [INAUDIBLE]?

ROBERT TOWNSEND: Oh, I'm not sure we have enough non-network households to do it. But the counter to that is that we certainly had enough in network with Cynthia to find something. So it's probably worth trying. I haven't done it yet.

Yeah, the worry here is that we've left something out. The model is just making a lot of strong assumptions that there's no alternatives. To try to say it in words, we're not looking at some indirect utility function. We're looking at the direct utility function. And indirect might be there if they have other ways of coping.

When we get to labor supply, we will be quite explicit about supplying labor as a way to cope with risk, and then revisit these equations. So people have-- I have, too-- offered insurance products like rainfall insurance in India. We haven't quite done it yet in Thailand. And I'll tell you why we have misgivings about it, and some of the puzzles. But basically with less of an exception in Ghana and Chris Udry, with Dean, most experiments in the world, the take-up of these insurance products is pretty low. Or in our case, in India, people will insure one plot, and not insure the other plots that they own.

So you know, it's not totally terrible. There is some price elasticity. But there are puzzles. I mean, if people are so poor and so vulnerable, then why aren't taking it up more? Now, one answer you may have already seen. They don't really need it. Because even if they're somewhat vulnerable, there are people in the village willing to assume some of that risk. And even if it isn't perfect, at the end of the day, given the way the policies are priced, maybe you're just not going to see that take-up.

Unfortunately-- and people aren't either taking the time to get baseline data or don't bother with it at all. They just go in there with an RCT. Well, we did. Yeah.

AUDIENCE: So is that just rainfall insurance? Or any kinds of insurance?

ROBERT TOWNSEND: Could be anything. I'm going to focus on rainfall as just a way to give it a label. Could be any kind of weather risk. It could actually be price risk. Case-Shiller index in the US is about price risk. It wasn't just to measure housing prices to get his name in *The Wall Street Journal* every month. It was originally designed because Bob Shiller had spent years creating a risk instrument which he was trying to promote on Wall Street.

So yeah, so it could be-- and in urban areas, maybe it is more likely to be vulnerability to price rather than rainfall. So rainfall here is kind of like, partly, a metaphor for the problem in general. But I'll come back with a vengeance to look at rainfall in these Thai villages.

AUDIENCE: [INAUDIBLE] even summer and winter, they say that rain affects [INAUDIBLE].

ROBERT TOWNSEND: Yeah. People don't come and buy when it's raining, and so on. But you came up with a situation where it was actually good. Just trying to remember what it was.

AUDIENCE: I think we said, umbrellas.

ROBERT TOWNSEND: Yeah, something like that. OK, so here's the vengeance paper with Kamilya. So this is all about rain. So we want to see the effect of rain on rice production.

Now again, the typical thing that you'll see in the development literature is rainfall is the risk. Let's just sort of instrument the relevant variable by rainfall, and then start looking at migration and all kinds of stuff. So here, the idea is, let's just look at rainfall in the village, start instrumenting income or rice production by rainfall, and then look at the vulnerability of consumption to instrumented income, which people do, a lot.

Now, it's going to turn out that things are more complicated than that within the village. The soil is different. The hydrology and water flow are different. They actually plant at different times. So a common rainfall shock, even if it is actually common, and shows up and all the rain gauges which we have, it's still the case that it has differential impact on the households, especially depending on when they plant, which is in turn a function of soil and so on.

So what are Kamilya and I doing? We're going to estimate a three-stage production function which takes into account the sequential nature. There's planting, for sure. And then you have sort of the germination period, and the planting of the rice, and so on, with maybe some fertilizing and so on. Then you've got sort of midseason, and all the stuff you need to do. And then you finally approach and include harvest.

You know, I'm going to jump, and just show you the notation. So here basically is stage I, with predecessor stages I minus 1. It's going to be looking forward to future stages if you're not at the last one. So if you're at the last one, this would be the harvest. That's the easiest thing to think about. If you're in the next-to-last one, the output is basically the condition of the crop prior to harvest. And what influences that? Well, if you've got land, labor, capital, and so on, as Cobb-Douglas type inputs in the production function-- but you've also got the state of the crop plot that you inherited from the previous stage. And then all of this gets hit, at the end of the period, with more shocks, including rain, and so on, and temperature.

So this is not a Cobb-Douglas. It's a CES that allows some substitutability to be estimated between the current state and contemporaneous inputs. It may not be as extreme as Leontief. So you're kind of doomed by the state of the crop plot at this time, based on everything that happened in the past. And then you just have to adjust inputs in proportion. Doesn't have to be that bad, but it turns out not to be like Cobb-Douglas. If you want to think about Cobb-Douglas, there's a whole lot of substitutability-- labor today, labor tomorrow. Oh, it's not that much of a difference. Why don't we wait and figure out what the rainfall is first, and then adjust the inputs later? So this allows us to estimate the ability to adjust.

And essentially the reason you can't use rain, aggregated up over the whole season, is because of this imperfect substitutability across the stages. Yep.

AUDIENCE: It says that it introduces-- [INAUDIBLE] that [INAUDIBLE] induces people to plant at different times, for example.

ROBERT TOWNSEND: No, that's not in there. There's a separate equation that has to do-- like a probit, basically, which determines planting times as we see it in the data.

So we can look at the effect of rainfall on rice production. And we can actually do something with climate change. Let me show you this thing. So this black thing here is the distribution of actual yields across households-- and actually it's households and time. This peak thing is just using a weather index. So you regress yields-- same data, our data, onto temperature and rainfall, and a whole suite of them, a big, long vector. It doesn't have to be, narrowly, rainfall. Try to be as realistic as possible. And then you run a regression. And then you get basically the predicted through the lens of that weather index model. And it's close. It's very, very peaked.

Well, sorry, that's the first thing. Then this thing is using the full model-- sorry, using the model, but only the observed rainfall variable, and setting everything else at their sample average-- common soil, common fertilizer. And then the big jump down here toward the data is when you start to allow the variation that the model has in it, certainly soil variation, and eventually full variation. So we're not one to one with the actual distribution, but we've made a substantial improvement from simple weather indices.

So people talk about basis risk. And what they mean typically is the reason there's not take-up of a rainfall insurance product is because the rain gauge is way out there, so distant that the correlation of the rainfall at the gauge with the correlation of the rainfall in the village is low. In fact, the Bank for Agriculture introduced-- I guess it's now three years ago-- an index product. And they already restricted it to be something like 20, 25 kilometers. And the take-up was terrible.

So then they said, oh, we know, it's basis risk. So then they went down to something like 15, lowered the radius. And take-up was still pretty bad. And I think they're going to have to withdraw the product.

So this is a kind of basis risk. What we're saying is, even though we're sitting on top of the rainfall gauge, within the village, the way that that rainfall is impacting these farmers is different across the farmers. So we're busy sort of thinking about designing, still, a rainfall product, but trying to take into account this heterogeneity. Yes.

AUDIENCE: Maybe this is kind of what you're getting at, [INAUDIBLE] how flexible is the sort of payout structure given? I'm just trying to-- because it seems like there's kind of like, a [INAUDIBLE]

ROBERT TOWNSEND: Yeah, so let me show you Kapphan's-- her dissertation. So here's her idea, which is basically-- and this is typical of the literature. She did better than most. You have a utility function over consumption. Consumption would be autarky income. But you have a payout-- could be positive or negative. And that's a function of this weather index, z. So this is supposed to be smoothing. And this is an optimal design problem. It's output conditioned on the weather index, and then the distribution of the weather index. By the way, where is she getting all of this? She's doing it in the lab. She's basically simulated millet production in Switzerland or something. But she's not out looking at the actual crops.

AUDIENCE: Is that [INAUDIBLE] soil?

ROBERT TOWNSEND: Hmm?

AUDIENCE: Is it [INAUDIBLE] soil?

ROBERT TOWNSEND: I don't remember exactly. But that is a big problem in practice when you take these sort of lab-type experiments. I forgot to say, we actually use a crop grow model as well. It's called DSSAT. And we carefully feed in soil depth and soil into that model I just showed you.

But to address your question, sort of optimized over this thing, this is the policy design. This is the optimal contract design problem. How cool. I mean, it really is cool. And you know, you can see, in her simulated data, the distribution of yields for any index. And then here's the optimized payout, which is basically positive, definitely non-linear. And basically it's like [WHOOSHING NOISE]. And many of the products that people already came up with are like these kind of linear, flat schedules. This has more wiggles, and it's more non-linear than that. But this is all negative territory. This is basically ex-post premia down here. Yes.

AUDIENCE: So these [INAUDIBLE] models, do they try to account for price changes that also result from aggregate weather shocks? [INAUDIBLE]

ROBERT TOWNSEND: No.

AUDIENCE: [INAUDIBLE] everybody [INAUDIBLE]. Because that could push it in the opposite direction.

AUDIENCE: Yes, absolutely. And I couldn't agree more. Now, at some level, you've got to either use data to decide how price is moving-- that's a good starting point-- and/or incorporate in your model like we saw with the macro development models, the movements of wages and interest rates and so on with interventions. So you don't have to do just one thing, but you're right-- you need to take it into account.

In Thailand, rice is the dominant crop. It's pretty hard to argue that like the price of rice in the village is a function of village or even necessarily regional rainfall.

AUDIENCE: It would have to be a pretty big aggregate [INAUDIBLE] price.

ROBERT TOWNSEND: Yeah. So we want to do something like that, that optimal design, to take into account what we've learned about within-village insurance, and their heterogeneity, and so on.

Just two comments-- first of all, is it true that within-village stuff is great and across-village stuff is not? Well, in Thailand it's not obvious at all. And in Hong's helped with this. We updated some earlier results. But basically this section of the slides, which I won't attempt to go over, is how to test not only within villages but across villages. And then you can aggregate it up, villages, and test one county versus another.

So you have you have layered sort of degrees of aggregation. And you can use [INAUDIBLE] method-- she did it in Kenya-- to see how good the insurance is across the villages. And in Thailand it's hard to say that it's worse than what's going on within the village. So it's just massive remittances. There's just huge financial flows coming in, which we pick up in the transactions data.

So I still have misgivings about pushing too hard on this. I'm sort of torn. I'm torn by the attractiveness of designing the intervention based on the data, and yet worried that they really don't need this rainfall product.

And the other thing I want to just point to so you can-- sorry, Hong, all your great stuff here-- is to say that Chris Udry, with Dean and co-authors, have been doing stuff in Ghana, and they have various treatments. They have rainfall insurance. They have credit. They have both. They have one and the other. They have controls. And then you see stuff. Like you see sometimes that credit seems to increase investment. But you also see that the combination of insurance and credit is actually not helping, although insurance alone actually can help.

So they wrote down a model, a simple two-period model, which is clear about the optimization problem. And it's similar in spirit to what we're talking about today, which is what's the shadow price of current consumption, what's the shadow price of future consumption. If there is full insurance-- which they assume in the village-- and there are no aggregate shocks-- which they assume in the model-- then basically you can get counterintuitive effects going on when you offer insurance. Because insurance moves tomorrow, it makes tomorrow kind of less valuable in some sense. It actually seems to undercut investment today.

And these other papers are again treatments, RCTs and so on. It's easy to get lost in the details of these things, and get caught up with the language of take-up, and behavioral responses, and some very real things that are out there. But it's important not to get lost, and to remember the basic sort of framework that you're taking in if you choose to design one of these randomized control trials.

Or let me put it another way. If Udry had not been thinking about those kinds of models, he would be left with some very strange puzzles that does beg for an interpretation. But even better yet, when you have the model and the structure, ideally, with enough data, you can actually design interventions which are informative about key parameters in the model and your guesses about basic obstacles.