ROBERT TOWNSEND: So this is about evaluating microcredit. And how to do it. It's also a methods lecture in some sense, because it's going to be an application where it was done both as reduced form, IV, as well as structural. So that gives us this great opportunity to think about pros and cons, pluses and minuses, kind of what you can get from one and not from the other.

And this is about a financial intervention. So we have been consistent in every single lecture in talking about financial constraints and financial modeling and talking about policy and policy implications of changing the financial system. So this is an illustrative to application of that general theme.

Very little review of the literature, but there are two, three applications-- actually the truth is there is very few RCT evaluations of microcredit. Abhijit, Esther, Cynthia, and Rachel did one in Hyderabad, Spandana. And in many ways, the results are strikingly similar to what we're going to find in the Thai villages and even some similarity to Morocco. What I want to emphasize, and they do too, is the heterogeneity.

So, for example, in Hyderabad, there's group-based lending, joint liability. They're small loans, female borrowers, subsidized low rates. On average, to be put in italics, there's no impact on total expenditures. But on average an increase in durables in the short run. And you might stop there, and that would be a mistake, because the really interesting stuff gets going when you start looking at the people who are running or might run businesses, business creation, business assets, self-employed hours and profits increase for those who already had existing businesses. Their durable consumption goes up. Non-durable does not.

Now, the consumption part is similar to the aggregate. But if you look, at say, the those who have as in a probit, a high propensity to start new businesses in the sample, they increase their durable goods spending. But they decrease their non-durable goods spending. So you are seeing something constant on the average and going down for one group.

And in fact, if you look at the flip side, which is not surprising given the average is what it is, the low propensity business owners actually increase their consumption of non-durables. So various heterogeneous impact based, by the way, on observables, namely current business owners, and unobservables that are backed out through these propensity scores, as in probits. So it's unobservable as well as observable heterogeneity that seems to matter.

Morocco similar, larger loans, lending to men. No effects on consumption. Maybe some small reduction in consumption for those doing agriculture and livestock. No effect on business creation, but within existing businesses, certain activities were discontinued less often, and the scale of other activities, as in agriculture and livestock, increase.

So again, another aspect, one, the consumption aspect is kind of similar in the sense that maybe for these guys that seem to have some business impact, consumption might be going down instead of on average staying constant. And then you have this kind of nebulous effect on businesses. A lot of people think of microcredit traditionally as an impetus to business creation. It's to allow talented, if poor, people to enter and expand and also enter into business. And, in fact, some of the early micro founded macro models that we went through in class very much feature that vehicle-- people going into business who couldn't previously, as their wealth expanded, you know, driving aggregate TFP.

Now, that's not to say those studies were wrong. These things can vary across countries. And they can vary within a given country, according to kind of the state of the financial system. And the overall development path. But we'll see, again, in the Thai study, this a bit of a struggle to figure out what's going on within existing businesses that's allowing them to be more profitable.

And that durable goods things, you know, again, as we've studied in the financial accounts, we make decisions about classification. But some durable goods are used not just as a service flow for household consumption, but also potentially used in the business. And if there are sort of increase in durables might suggest, they're actually expanding the business in some way, not just taking it out in terms of increased consumption.

So in Thailand, it's the Million Baht Fund. It was a government intervention. It was not a randomized control trial. And I guess I'll deal with that right away.

How are we doing an evaluation as if it were a randomized control trial? Well, it was basically a million baht per village. Villages are well-defined geopolitical entities. But they vary in terms of total population size. So all villages got the same amount of credit. The villages with fewer households got a higher per capita treatment. And there were 72,000 villages, by the way, in Thailand. And they all got something in the order of $24,000, $25,000. So this was approximately 1.5% percent of GDP. So you could easily claim it's the world's largest microcredit sort of expansion and evaluation.

Now, it's a quasi-experimental design. You'll see in the slides, but I'll say it now, you know, we carefully look at pre-intervention trends to see whether there's anything else going on systematically in a comparison between small and large villages. And we can't find much. We've even gotten the maps out and used this community development department data for all the villages in Thailand and looked at village size to see if, for example, larger villages are near urban areas or nearer to main roads and so on. They're not. We only find one or two variables that are correlated. And you kind of expect if you're searching over hundreds of variables to find a few things now and then.

So we're going to look at the effect of credit on consumption, investment, income. And we're also going to look at some general equilibrium effects. And again, just for worry that I'll forget to say it later, if you view villages as small open economies-- and we have seen papers on that-- then by intervening village by village, there's some sense in which you're intervening country by country. That's pretty rare opportunity.

These models we've been dealing with take a stand on what might happen with wages, for example, in general equilibrium. So we can actually see if there's a wage effect going on over and above the potential benefits or other impacts on households who got credit. Households who don't even get credit, but live in a village where business might be expanding would be villages given certain obstacles to total migration, would be villages where the wages are going up. And again, that's what we saw in those first few lectures to be the impact of the expansion of the financial system. Question?

AUDIENCE: So I guess, ex ante, did it surprise you that there was no correlation between things like village size and proximity to resources or productivity of the average business in the village or--

ROBERT TOWNSEND: One thing, it's a political decision. And they divide and subdivide villages. But they don't do it on a regular, persistent basis. So if a village were getting bigger and bigger for economic reasons, then it could get divided by the government into two separate villages.

AUDIENCE: So what is the measure that you looked down? The one decided by the government?

ROBERT TOWNSEND: Yeah. At a moment in time, we have the identity of each of the villages. And we count the number of households in the village.

AUDIENCE: So potentially something that we count as two villages is actually one big village?

ROBERT TOWNSEND: Well, so the real question is whether the funds could have leaked over to neighboring villages. The grant that established this village-level institution was for a village. You know, like Downtown Crossing gets a village fund. Beacon Hill doesn't get one. Sorry, close to home.

AUDIENCE: There's huge incentive not share it, right? I mean, it belongs to your village. Why would you want someone else--

AUDIENCE: No, no, what I mugging, is I think sort of agricultural or resource economics suggests that village size should be endogenous to resources in an area. Like there should be more villages near freshwater source, like in, you know, maybe a thousand years ago or something.

ROBERT TOWNSEND: Well, but then adjust for population. Actually, I think your argument goes the other way.

AUDIENCE: OK, sure, I mean--

ROBERT TOWNSEND: If villages are an arbitrary collection of people, then you might want to have the same number in each one of them. There could be more villages. But potentially there are same number of people per village if the government has some systematic rule. Now, the government is messing around, but not as systematically. That's one comment.

The other thing is this leaking, because if I could get credit from some other village and/or, as we've seen with the informal kinship networks, get more money from someone else who is connected to another village, then we wouldn't have any power. Now, we have looked at that too. And again, we've done the GIS. We've gotten all the maps out and looked at the impact on villages nearby, creating these sort of smooth neighborhood averages. And at least early on, the impact doesn't spill over. But that doesn't say in the long run it might spill over, which is another interesting potential thing to do. Yeah?

AUDIENCE: I have a question that's similar-- how essentially villages could be subdivided cities, how you can consider them small--

ROBERT TOWNSEND: Yeah, have I answered it?

AUDIENCE: I think.

ROBERT TOWNSEND: OK.

AUDIENCE: Well, I still don't get why you can, but if empirically it looks like there aren't spillover. I guess that's somewhat of a justification.

AUDIENCE: Yeah, I accept what you're saying is true. But it confuses me from like a priori point of view why is that just-- it seams--

ROBERT TOWNSEND: Well, we didn't take it for granted and get 3/4 of the way through the study and then think about it.

AUDIENCE: No, no, I know--

ROBERT TOWNSEND: We looked right away, and I would have thought--

AUDIENCE: I can tell from--

ROBERT TOWNSEND: My hunch was larger villages would be near the cities and the main roads.

AUDIENCE: Yeah, it might have a higher average TFP or something.

ROBERT TOWNSEND: There's a map in the paper. And, you know, I'd created these slides kind of too late to include it. But you'll be reassured by the map.

AUDIENCE: No, I believe you. I am just confused by--

ROBERT TOWNSEND: You can rub your hands on it and--

AUDIENCE: Some sort of that's true. But it's not true, like you have some sort of hubs all scattered around, so like the size, you know--

AUDIENCE: So the one thing is that if it's really urbanized-- so there's a limit. Once you reach a certain population, so this takes care of some of the heterogeneity. Once the population reaches a certain size, like the government will just designate it a town instead, so it won't be a village.

AUDIENCE: Sure, that would take care of some of it. But it shouldn't take care of all of it. There should still be some variation in the boundary. Well, one would there would be variation.

ROBERT TOWNSEND: It would be depend on how quickly they subdivide them and so on and so forth.

AUDIENCE: That's true.

ROBERT TOWNSEND: We just looked empirically and got reassured. Yes.

AUDIENCE: A related question is sort of who has the incentives to subdivide up a village? Like have there been other sort of village-level programs in the past which-- I'm not sure would got the incentive. But if predicted something like Million Baht Village was coming, you'd want to subdivide your village.

ROBERT TOWNSEND: Yeah, so on that, the program was announced pretty suddenly and implemented. Thaksin ran on this and implemented it very quickly after being elected. So arguably it couldn't have been anticipated.

Now, you know, the larger question is, why was if he was trying to maximize votes you'd want it per capita? Why give a lot of money to a village where there aren't very many people? But, again, in the Thai system, villages are thought of as part of the political system it's village, tambon, amphoe, and so on.

So you give it to a committee within the village to run a fund. And somehow it seemed, quote, "fair." It's great for us, right? I mean, there's wonderful mistakes in some respect. It's kind of a unique opportunity.

So we did that. Oops, I see. OK. So as I said, it's part of the GDP.

AUDIENCE: What is the 77,000?

ROBERT TOWNSEND: Yeah, I think that's supposed to be villages, not households. That's too low. This I've said. So if you don't mine, I'll just skip it.

Loans are about $500, short-term, as in one year, typically two guarantors. The interest rate was higher initially at about 7%, although the village committees vote on those things. And they allowed loans for both consumption and investment, unlike the stereotypical program where somehow consumption smoothing is supposed to be a bad thing.

So it's my data. But in this context, probably bears repeating because it's unusual that we didn't see this program coming either. We're out there gathering data year after year. And we have five years of the pre-intervention data. We're going to exploit that like crazy, because we're going to write down a model, which I'll tell you about, and estimate it, all based on the pre-intervention data, and then simulate the model in compared to what actually happened.

We actually have six years of post-program data too, which is, again, rare, and here, quite interesting. Again, if you're thinking about, which you may not, but you might be thinking about budget constraints to do RCTs and so on. It's quite expensive. And gathering the data is a big part of the expense. So running a survey for a long time before the intervention and running it for a long time afterwards is almost prohibitively expensive. Or at least you wouldn't be able to evaluate very many things.

But we didn't design the data gathering with this thing in mind. It just happened. I guess in that sense, again, we just got lucky.

We're going to look at four kinds of outcomes, short-term credit, overall, I might add, not just from the village funds, including, therefore, borrowing from the BAC, the Ag Bank, commercial banks, the reasons for borrowing-- although I don't have those slides for you here because money's potentially fungible. We do look at interest rates default and so on.

We have consumption broken down in these annual data, by the way, into these categories. Why these categories? Well, we had the socioeconomic survey. And we did analysis when we designed my survey to see what subset of consumption items would best predict the total in the socioeconomic survey. So that's kind of in the spirit of Blundell and Pistaferri and so on in terms of taking advantage, projecting to get something bigger than food. But actually there's some interesting stuff with alcohol. I'll show you, if I remember, when we get through these things.

Income and production decisions is another class of categories. We have assets and income growth, even in the annual data, divided into business, wage and salary, and agriculture and so on. And we see revenues and expenses. And again, we have some ability to look at female headed as opposed to male headed households.

OK, what's an outcome? y for household n at date t. This is like imagine running, oh, my god, a simple-minded OLS regression, the way people sometimes used to do. So Village Fund credit would include the Million Baht Fund. You want to see the impact causally on an outcome. Of course, we're going to instrument this. But the other things are these household characteristics, of various kinds, mainly demographics, age, education, number of kids, blah, blah, blah, gender. And then we have a time fixed effects and a household fixed effects, which are estimated in the data.

And here's the endogenous variable being instrumented. And this is pretty simple. And we also have those x covariates here and time and household fixed effects for household short-term credit, household and at date t. But we've also got some dummies that kick in the first year and second year of the intervention. Those are chis not x's. Those are 0, 1 variables. And they're interacted with the inverse number of households.

So inverse, that's just to get the sign right. Although it's a parabola, so it's not exactly not mattering. But if you wanted this thing to be positive, the larger the number of households, the smaller this number, and kind of the less you would expect the impact to be. And then we have the usual all the errors are uncorrelated with each other, et cetera, et cetera. Yep.

AUDIENCE: How many villages was this?

ROBERT TOWNSEND: This is 64, because it's my annual data.

AUDIENCE: And can you also look at outcomes in the CDD data? You know, 450,000?

ROBERT TOWNSEND: We have not amazingly done that.

AUDIENCE: CDD is at the village level. This a regression at the household level.

AUDIENCE: That's true.

ROBERT TOWNSEND:	But you could potentially do things across villages to see where the villages got more or less credit. So the spirit of it is possible. Someone actually looked at the socioeconomic survey and had enough connections to the Nash NSO to get it by village. But their timing was off, because their first year is the year of the intervention. And we have the pre-intervention data. But anyway, it's possible. No one has looked at that.
AUDIENCE:	What was year of the intervention?
ROBERT TOWNSEND:	The first year?
AUDIENCE:	Yeah.
ROBERT TOWNSEND:	2002.
AUDIENCE:	Why don't they have the ISS data? They just don't have the years? I thought the ISS go back 30 years.
ROBERT TOWNSEND:	They don't have a panel.
AUDIENCE:	Oh, OK. But they don't have panel afterward?
ROBERT TOWNSEND:	They do. They did. Yeah.
AUDIENCE:	But we don't have it.
ROBERT TOWNSEND:	We discovered it after. We have it. So we're going to look at new short-term credit, consumption, asset growth, and income. Tried to do it pretty carefully. We actually did it about four more ways than this. But reduced the number of that we show. OLS with all the problems that you know about. Here's an IV regression with all villages as opposed to fifth-- this is kind of outliers in village. If you think about that parabola, you know, it's kind of getting really steep when household size is really small-- number of households is small and not moving much. So we clipped off, took the big chunk, but we clipped off the tails of small and large villages and dropped some outliers from the household data too.

But every single way we do it, short-term village fund credit goes new, new, short-term village credit goes up. These are roughly the order of magnitude of $25,000 divided by the number of households. So, you know, credit goes up. You might be worried. And like Esther and Abhijit have papers where there's this substitution going on. You get cheap credit from one source. You borrow less from other sources. That doesn't happen.

In fact, arguably there seems to be almost an amplifier effect, but more weakly. Short-term credit goes up, for sure, one to one. And you can see it in the data. I mean, this is this like a huge spike in the financial system. And it didn't go down. That part doesn't go down subsequently. Yep. |
| **AUDIENCE:** | If we thought that going up variable was-- took that seriously, could that be something like people getting extra funds and re-lending and so you have to sort of multiplier? |
| **ROBERT TOWNSEND:** | It's true one. Would have to-- that's a good point. Once could be just-- not to say that it's a bad thing. |

AUDIENCE:	Sure.
ROBERT TOWNSEND:	But, you know, you have these sort of combinations of characteristics of loans. And if the village fund committee has this cap, which they're supposed to have on the loan size, then you get one leg up and think about doing something. Maybe you want to go the money lender or other sources and lever up even more when you have chunkiness.
AUDIENCE:	So this is because it's going more one to one, the credit?
ROBERT TOWNSEND:	Yeah. But the more conservative way to read it is just one for one anyway. So it's not like a substitute for other sources of money.
AUDIENCE:	So when I initially looked at it I thought this support for Esther and Abhijit's idea that if you can get over a hump, then you'll really want to expand or something. But I think what you were saying is could they be re-lending to each other? So I could I be--
ROBERT TOWNSEND:	Yeah, he did say that. And, yeah, that could be going on.
AUDIENCE:	OK. That's a convenient explanation.
ROBERT TOWNSEND:	And I don't think we tried to take that out. We're just looking at the household and at date t in some village j and adding up all the short-term loans that they have. So if you only counted the loan and not the lending, if it wasn't net, and I think it wasn't. I think it's just gross.
AUDIENCE:	Not all loans are one year. They're are loans that are one week.
AUDIENCE:	No, sure, yeah.
AUDIENCE:	They almost have different lending, really, really short term.
ROBERT TOWNSEND:	So that's why we did short term instead of long term. Consumption levels, other than the OLS-- I'll come back to that-- go up. Significantly, asset growth in one specification goes down. Now, I can't help it but say that we're going to have to choose a model. What I'm going to do now-- I'm giving you these stylized facts-- but what model comes to your head on why things like savings might go down? Like a buffer stock model.
	If you're on your own, you're saving for these future disasters. Now, in the future, you could borrow if you believe the village fund is persistent. So you don't need to have so many buffer stocks. So net financial savings goes down actually.
	And we have trouble finding stuff with the capital stock. I'll mention that. By the way, these consumption magnitudes are very similar to the credit magnitudes. So a bit naively, you might say they used it all for consumption. Now, again, what model do we have that would predict that consumption ought to go up one for one with an increase in credit?
AUDIENCE:	Life time permanent income model.
ROBERT TOWNSEND:	What would the permanent income model give you?

AUDIENCE: It would go up, wouldn't it?

ROBERT TOWNSEND: How much?

AUDIENCE: It has hazard shock is temporary or persistent.

AUDIENCE: You would think this persistent, right? So then it should jump to the new-- but I don't know how much.

ROBERT TOWNSEND: You'd put it in the bank basically and draw the interest off of it.

AUDIENCE: This is a temporary shock, right? It's just once.

AUDIENCE: Then you can get the funds back and you lend it again.

ROBERT TOWNSEND: Either way.

AUDIENCE: We're missing it. What's the--

ROBERT TOWNSEND: No, no, no, I'm just saying-- no, no, you got close. I mean, the point is the permanent income model isn't going to get you this. And you've got to have these credit constraints. You've got to be something that's generating hand to mouth like behavior.

AUDIENCE: So this is a shock to the credit constraint not the shock to income. So you should be facing this credit constraint-- you should be constrained.

ROBERT TOWNSEND: Yeah, that's right. And therefore, you're going to see that this overall average is masking a very big differential impact in the population depending on whether someone's at that constraint or not.

AUDIENCE: But even if you are not constrained, if you anticipate that in the future I'm constrained, endogenously you--

ROBERT TOWNSEND: It depends on the interpretation. Suppose you thought, you know, the government gave the money lump sum to the households indirectly through the village fund. And they just put that in the bank. Then their net worth has gone up. They don't have an obligation. Within the village, they are borrowing and lending against this fund. But the village, as a whole, just got a whole lot richer. And they would potentially, if you believed in the permanent income, raise their asset level, right? And then get the flow off of it. Yes?

AUDIENCE: How is it financed? Is it something like--

ROBERT TOWNSEND: Taxes.

AUDIENCE: But like the taxes mainly hitting cities instead of villages?

ROBERT TOWNSEND: Yeah.

AUDIENCE: I see.

ROBERT TOWNSEND:	Matt.
AUDIENCE:	So I was a bit confused between the OLS and the IV, because what is the expansion variable in OLS I was thinking it must be credit per person. But then I don't understand really what the instrument is, right? Because if its credit per person it's just a million divided by number of households. And that instrument itself is per household.
ROBERT TOWNSEND:	So there's a lot of variation in short-term village credit, both among households within a village and across villages. So we're identifying this through the cross-sectional variation. We're saying what part of the increase in observed increase in short-term village credit can be explained simply by inverse population size? And the OLS doesn't do that? It's just a flat out regression credit on the right-hand side and impacts on the left. All credit.
AUDIENCE:	OK.
ROBERT TOWNSEND:	And, you know, net income growth is positive in 3 to the 4, including the OLS, or 2 of the 3 if you don't. And so there is some impact on overall growth, not levels, growth. And this is by source of income. You know, we see a hint of profits, apropos the first two three slides of the lecture today. Profits may be going up. And wages and salary payments are going up pretty consistently.

So again, you know, what that's all about? But I've already given some of it away. The wage rate that's going up for villages that got a higher per capita treatment. And potentially employment goes up. Yep? |
AUDIENCE:	So this slide and the previous slide as well. I mean I understand we're talking about bahts, so the quantities are big. But like some of yours are huge.
ROBERT TOWNSEND:	But you can see the stars.
AUDIENCE:	No, no, no, I understand that. But even in the previous slides, on stuff that's like the growth rates, the standard errors are big.
ROBERT TOWNSEND:	Yeah, Yeah.
AUDIENCE:	You don't get clear zeros and stuff like that. So--
ROBERT TOWNSEND:	Wait, I'm not sure what you're saying. Yes, the standard errors can be big. That would make it tough to identify significant effects. But we on occasion do find significant effects.
AUDIENCE:	No, I understand that. I'm just wondering-- my question was more of as technical one. Like how-- is there a way to sort of make them smaller. I mean I guess it's an issue of ID and stuff like that. But--
ROBERT TOWNSEND:	Well, you know, this is with clustered standard errors and the whole thing. And we try to eliminate the outliers to see if that's kind of causing a lot of the noise. Sometimes it does. It doesn't do that consistently. OK.

We don't find much with ag or other sources of income or livestock. Now, as I said, we also have a lot of post data, post-intervention. Even when we wrote this thing, we're up to six years of data. But this is OLS. So you have to take it with a bit more of the grain of salt. For example, consumption isn't showing up, even initially because we're not instrumenting anymore. |

The probability of default's kind of moving around. And it doesn't go away. And that's a real issue. And let me show you the default rates. And, you know, potentially, by the way, there should be a lag. Depending on exactly where the 12 month falls, you borrow one year and you can't repay the next.

Net income goes up. But it's not significant after that. And short-term village credit remains high.

Well, what do I want to say? Back to models, you'll get one, don't worry. You know, the world isn't static. And we just shook things a little bit, because these villages arguably are wealthier than they were and they've also somehow used this to capitalize basically a savings and loan association within each of the village.

So if you were in a steady state, and someone smacked you in the head with something good, you might expect the adjustment isn't necessarily insta-- all the dynamic models we have with frictions, Paco's, for example, suggests there are slow transitions even to some ultimate new steady state. So this is rare, this look at these-- now, granted, the model's going to be partial equilibrium, even though we have some data on local wages.

And there's a guy named-- is it Furfine who has a study? Pandi and Burgess had this study in India where the government changed the branching rules and showed impacts on poverty and all of that stuff. But he went back after they published their paper and looked at the long-term impact. And he thinks something like this was going on in India too. So again, apropos, you know, RCTs, well, no one claimed the contrary, but if you don't know what happens afterwards, you're kind of left with the impression and hope that the program effects, if they were positive, persisted. Here, they're definitely mitigating over time. Yes.

AUDIENCE: Do you have the long-term impact on wages?

ROBERT TOWNSEND: I don't think we looked at that. That would be good to do. We struggle with the wage, because the annual data measure it very well. So this is the one thing that we use the monthly data for. And we only have 16 villages. So it's quite problematic. Most things we try don't show up because of the sample size is too low and the standard errors just kill us. But I don't think we actually looked at long-term wages. So that would be good to do. Yes.

AUDIENCE: Do you think the wage data in CDD is any good? Daily wage for laborers.

ROBERT TOWNSEND: I think when you sort of look at geographic averages of things, yeah, that it has-- there are beautiful maps showing how wages, not size, but wages move as you get near urban areas and get near Bangkok. And so arguably, something could potentially-- I wouldn't trust any one head man's response, but with 72,000 data points, you can do a lot.

AUDIENCE: So that paper, his name is Fulford.

ROBERT TOWNSEND: Fulford. OK, so this is a cautionary tale. I'm going to skip it. If we have time, I'll do it next time. That would be a better context for it.

So I've been trying to clue you into the puzzles all along in terms of-- the one thing I didn't say was about investment. And I had mentioned in other lectures that investment doesn't happen very often. And when it happens, it's big. So it's quite lumpy. And we're going to want to incorporate that.

But the flip side of that is we're going to have trouble finding effects of investment both in the IV and in the structural model, and the data generated from the structural model, if we limit that data to the same sample size that we have in reality. But when you see the structural model momentarily, you can imagine a 10-fold increase in the sample size because we can generate as much data as we want. And then, we definitely get these investment effects. So, you know, arguably one reason still on the table that we don't see microcredit directly in terms of investment activities is simply a sample size issue.

I mentioned that because, you know, people write reviews of what do we know about microcredit and whether or not it having an impact. And there seems to be a bit of a consensus that at best, impact is quite mixed. And maybe, you know, non-existent in many studies. And people come away with this impression that, at most, it's all about consumption. But I think that's a little bit of a rush to judgment. I'm not saying it has to be this too. I don't have more data in reality than we have.

So we're going to have this precautionary savings. We're going to have some limits because of those consumption numbers I showed you. We're going to allow default. It does happen in reality. And we're going to try to match the default rate. And we have income growth. That's actually arguably the hardest thing that we did. And I'll try to say something about why it was hard. But we actually allow persistent growth in the model.

OK, so what is the model? Liquid stuff is just your current income plus your return on previous year savings. Now, this reminds me to say, you know, so why aren't you doing moral hazard or costly state verification or full information? Hey, we just had a lecture about that. Same kind of data. Alex and I showed that in the rural data, the best approximation is something simple like savings only or limited borrowing.

Now, we got lucky, because Alex and I did that after Joe and I wrote this paper. But it is at least very comforting that the best micro underpinning that we now think exists out there was the micro underpinning that was being used in this paper. Likewise, I would not expect to get the same thing in the urban areas, because there, Alex and I found moral hazard was a better approximation.

So here, it's sort of this incomplete market's macro type literature or development literature, for that matter. Rotsio swears by this stuff. And it's in logs. So income looks like it's just a multiplicative transitory and permanent shock. That makes the log of income equal to the log of some permanent thing and the log is some transitory thing. The log of the permanent thing almost looks like a random walk, where log Pt plus 1 equals log Pt plus log N. So this is the shock to permanent income. This thing is the shock to transitory income, where income is measured in logs.

Now, it's not quite that because G is this drift, this damn drift. It's 4%, by the way. So you know, there's kind of sustained growth, not modeled. You get this boost to your permanent income. It's like TFP type stuff.

And then what's this? So investment, you know, why invest if it doesn't do anything? Well, on the contrary, investment gives you a boost of your permanent income, right? So this is this decision whether or not to invest, a dummy. And this is the amount of the investment. As I said, investment opportunities arrive stochastically from a distribution with a non-trivial mean. So on average, it's chunky, which means you may choose not to do it because it's too big to swallow. You don't get a choice about the size though. If you do it, you've got to do the size that arrived to you as an approximation.

Although again, I've said this, we see this in the data. People don't do sort of half chicken coops. A chicken coop without a roof is not a coop. OK.

AUDIENCE: I have a question about the drift term. I think we have in papers we have-- we basically assume some exogenous technical progress pushing wages up and then like papers where the focus on the structural model is to model that, because obviously like that's really important if Thailand over this period there was 4% drift where everyone or maybe--

ROBERT TOWNSEND: It is important. And we're not. A lot of the literature we've covered doesn't. Daron features that a lot. What is it? Is it innovation? What is this technological progress?

AUDIENCE: So would it be-- Are there any inherent reasons why it would be difficult to translate Daron's models to the Thai data or the--

ROBERT TOWNSEND: It should be done--

AUDIENCE: With regional variation?

ROBERT TOWNSEND: It should be done. It hasn't been done yet. No, I don't see any intrinsic difficulty. Actually, it's more like a real opportunity, because when you're on the ground like this, you should be able to see-- and in my data, you measure the capital stock and labor input and all of-- you know, if it's real technological improvement it's like, mm, getting rid of the water buffalo, and bringing in the walking tractor.

I showed you a picture of the walking tractor. I said it was land, labor, and capital. Remember that one? Well, it used to be water buffalo and the people in Bangkok still think that you know the farmers are out there with their water buffalo. Well, I mean, it's just not true. You know, they have pickup trucks and all of that.

So my point is that that's real technological progress. And so data like this potentially could tell you whether it's really that or just potentially something else. It could be better roads, you know, higher, easier access to markets that somehow showing up, everything else equal, but not modeled, and not at a deep level. OK.

So here's the max problem. Maximize discounted expected utility of this household. Rho is sort of the related to the coefficient of relative risk aversion. Beta is the discount rate. This is as of, quote, "initial period." Here's the liquidity. That, by the way, was the same as interest on savings, short-term liquid savings plus current income. And it can be spent on consumption or saving or investment. If you decide, D, to do the size investment that has arrived.

What are the state variables? In any given day-- this is true for every day, t. It would have been better, and I don't have it on the slides, to show you a traditional value function, where you'd have the utility today plus the value tomorrow. But this does tell you what the state variables are namely, permanent income today, liquidity today, and the size of the investment draw. So those are the key economic states. That's households are going to face.

This is just a parameter. What is s bar? Here it is. Now, it looks like savings is bounded from below, weird, weird. Well, actually savings can go negative. That's fine. That's just borrowing. So this is a credit limit. It says savings can't go too negative, credit can't be too big. And it's scale by permanent income. In fact, almost everything in the model is going to get scaled by permanent income. Part of that is the permanent income, we wanted growth. And the way to get the growth is to have this expansion exogenously in permanent income.

But the other thing is you look at the data, and you see, you know, like investments are large for small households relative to their assets. And you go to larger households, and their projects scale up. So it's not like you can sort of accumulate wealth and save your way out of these non-convexities.

So we scaled everything-- go to the other extreme-- we scale everything by permanent income. The arrival of project sizes and so on and so forth and the shocks are all scale. And what I'm not writing down, but we do in the paper, is to actually turn this into a balanced growth path. And the intuition is pretty much what I said. The math is trickier, where you just divide through by p everywhere. So the control variables are things like consumption per permanent income, per unit permanent income. And then we have to solve the value functions.

Now, these are the stochastic processes for the transitory shock. Again, everything is in logs. The permanent shock, so these are log normal, centered around zero with these standard deviations. Project side is not centered around zero. It has a non-trivial mean with the standard deviation. And you're going to see a table with parameters-- not to say you're really memorizing these. I am, because I want to go through the list when we get there-- mu sigma, U, N.

Here's a problem. We do have a borrowing limit. And we do not put in sort of a natural borrowing. We let households default. They default in the data, about 18%. It's not trivial. It's also hard to nail down exactly when a loan is not performing, because they can stretch out the payments. So it's not nailed at exactly 18%.

And, oh, yeah, so here's another story with a bonus. *Econometrica* requires the codes for any published paper. So you've got to give them the codes and the data, which is good. It's a good thing.

Of course, when you're doing the work, you don't really annotate as thoroughly as you might do. And then when you're ready to publish this stuff, you're going to like go back over it, and we found a mistake. And the mistake was around this default rate. So we actually know, for bad reasons, that the default rate matters. Yes.

AUDIENCE: So I saw two versions of a paper and they had quite different welfare coverages. Was that--

ROBERT TOWNSEND: That's probably part of it, yeah. Well, look, everyone makes mistakes in research. And it's in some respect doing it over and over is actually a good thing, because you get this robustness check. But, yeah--

AUDIENCE: So in the data, when we say a household defaults, what do we mean exactly by that?

ROBERT TOWNSEND: Never paid off the loan.

AUDIENCE: OK. Because in most microcredit literature a default is what we would call a delinquency in the sort of developed country finance literatures.

ROBERT TOWNSEND: Yeah.

AUDIENCE: So you're not counting deliquincies as default. You're counting actual defaults, like we just never pay the money back.

ROBERT TOWNSEND: You know, I have to go back and make sure. We could have put something like a two-year threshold on this. But, yeah, I mean, I always actually go the other way. I tend to think of defaults as a contingency in the loan contract. The Bank for Agriculture for sure runs an insurance company this way, where they extend loans and actually even more than that don't charge you interest on the renewed principle and so on. So, you know, people are risk averse, having these contingencies are a good thing. There is a literature on this, other literatures.

So anyway, there may be other ways to do it. Measurement is tricky. But we decided to put default in. And what is this thing saying? Well, it's a bit tricky to read. But we don't drive people down to zero consumption. We put in some minimum consumption. No matter what, you're not going to go below c lower bar and, well, scaled by your permanent income as everything else is.

Now, what is this other thing? Savings. Well, put savings times p on the right-hand side. It'll pick up a negative sign. That's credit. Minus savings is credit. So it's liquidity plus the maximum loan you could take out. And still, you can't cover that minimum consumption. So you're not respecting the budget constraint anymore. You get this, quote, "gift."

And then here are the rules that happen and when you default. You're not investing. You're borrowing up to the limit. And your consumption is the scaled version of permanent income.

So let's see if I've done my homework, or you have. R is the interest rate. This is the standard deviation of the permanent shock, standard deviation of the transitory shock. This is new. It's measurement error. We imagine that we see the data contaminated. It's a bit dirty centered around the mean. But this is the drift term. This is the lower bound on consumption under default. Beta is the discount rate. Rho is related to constant relative risk aversion. Mu is the project size. On average, sigma i is the standard deviation of project size. And this is the credit limit.

And R, big R, is the return on investment that augments permanent income. That's also quite problematic. What we did was use these return on assets numbers. But, Emily and Abhijit and I have been working subsequently. And we think not only does R vary in the population with some heterogeneity that we don't have here, but it actually is predictive of who's getting the money and what they're doing with it, which may be another reason that we didn't get as far as we thought we might with the investment thing. But anyway hopefully, we'll write that up real soon. Here, it's just a number, calibrated, you might say. And we're going to use method of moments.

Oh, now, a couple of questions for you guys, or one comment. Back to the codes, you know, Emily was about your cohort in a class like this. And very shortly thereafter, she got interested in the structural stuff. The codes are all there. And, you know, not only the published version, but all the other things, you're welcome to have them. Probably a little late to do it for the class, but it's a resource that's available to you. And, of course, in the end it wasn't the paper with Abhijit that's still not quite done. But she used structural modeling in her job market paper. And she was learning it by looking at those codes.

So the other thing I want to ask you, I hope you've seen at some point methods of simulated moments. One yes, three, four yes, OK. So we're going to--

AUDIENCE: For those of us who took Devop 771 in the Fall, Esther had a problem set on your papers.

ROBERT TOWNSEND: Have you learned have something new today?

AUDIENCE: Yeah, yeah. Like we did the methods of moments already.

ROBERT TOWNSEND: OK. All right, that's excellent. Yeah, that's excellent. I guess I knew that. I'd forgotten. But anyone, I'm glad you know it. So I won't try to reteach. There's not time today.

OK, another controversy, what to do with all the heterogeneity in the data that we don't have in the model? I'm emphasizing certain aspects of observed and unobserved heterogeneity in the data. That is in the model.

But we've also got all these demographic stuff. Now, granted not everyone would do it this way. We take it out. So we filter the data basically. We regress the household data on these observables and as well as time trends, potential initial business cycle, i.e. the initial data is one year after the financial crisis, the whole country is kind of-- so the controversy is, well, why don't you just put it in-- and I'm not going to have time this year. But I have lectured in previous years on Keane and Wolpin and stuff. So if you want to see a really, really long specification of what you can put into the utility function, I'm not debunking it. It's a perfect--

[AUDIO OUT]

ROBERT TOWNSEND: We kind of estimate it. But this is almost like calibration, because we have earned interest, which is interest time savings. And we have savings and interest in the data. So we are trying to minimize to get the mean right.

The other things that are more fun-- consumption, decision to invest, and so on is just basically looking at consumption in the data. And this is expected consumption through the lens of the model. It's conditioned on observables. Liquidity is observed. And income is observed.

Now, one caution, these are not one to one with the key state variables of the model itself. There are something like permanent income and project size relative to permanent income are the key state variables. And permanent income is unobserved. So there's some work to back out these expectations.

But anyway, you get an error, which is the difference between the observed and the predicted values. And you kind of like want to minimize the error term by choosing the parameter values here. Here's a la Blundell, Pistaferri and Preston. We put in a log moving average process for income. So if you're judicious about choosing log differences far enough apart, you can pick up the drift. It's really very similar to what BPP we're doing. And that's how we get that G basically. So these are time differences in growth rates.

Maybe I don't need to-- and then, of course, you can pick up other moments by the orthogonality conditions on error terms, assuming that they don't see something else. If they were to see something else we're not seeing in the model, then we're going to make a mistake with these guys, because there's going to be information contained in the error. But we assume not.

Simulating, what are we going to do? Well, first of all, how do we relax the borrowing limit? We allow that borrowing limit, that s bar, to move from one village to the next in such a way as to predict the increase in short-term village credit that we see in the data. So that's how we sort of calibrate the magnitude of the intervention.

And the rest is simulation. We're going to draw these shocks over and over again and get repeated samples. We're going to have basically 500 artificial data sets. Any one data set is a series of draws of the permanent transitory project size and so on.

And we have the pre-intervention years. You can stop there and use the data and estimate the parameters. You can keep going and see what the model predicts over and over again, not just one path, but averaging over these 500 different paths.

Why do that? Well, basically, you want to get average tendencies. You don't want to be so sensitive to the luck of the draw in the first year after the intervention and the second year after the intervention, because we don't know what-- the model says-- we know something about averages.

In fact, then we run-- this slide really deteriorated through the projector. But these are the economic variables in the model itself. Liquidity divided by permanent income, projects size divided by permanent income. And what's going on here, as you move in this direction, there's three black dots here. And they're kind of all fixing project size and varying liquidity.

So these guys down here are bankrupt. They're defaulting. They have almost no liquidity. They can't even really pay back their previous loans.

These guys, these are these hand-to-mouth guys. This thing's really curvy. They're constrained. They're liquidity constrained. That borrowing constraint is binding.

These guys have drifted into higher territory. They're not actually constrained anymore. But in the future, they know there are credit limits and so on.

If you go this way, you're basically fixing liquidity and varying project size. Now, the most interesting part of the diagram is the Grand Canyon here. And this sort of pre-intervention, this guy sitting close to the cliff, he's tempted to invest. But given limited resources, investment would drive his consumption really low. And he's not willing to do it.

But then you have this sort of increment in the borrowing resources. It's like another [AUDIO OUT] that effect on consumption, we've got these four. These guys don't move consumption basically. Oh, well, let me get at the welfare stuff right now.

So for them, this credit program is terrible. Why? Because we're not going to force them into default anymore. We're going to make them take out a loan at interest, whereas before they got exempt from repayment. So welfare goes down for these guys down here as a result of the village fund program. Welfare goes up for these credit constrained guys, because they're at the borrowing limit, and that's precisely what this new program is doing, at least through the lens of the model.

These guys, you'd say, well, they're not constrained. No, but they're these buffer stock guys. And they now have excess liquidity. Well, might as well spend some of it. So their consumption also goes up. And as I just said, these guys who drop off the cliff, so to speak, they actually have consumption drops. So this consumption average that we see, 1 to 1, is this mongrel weighted average of overall this heterogeneity, of course, that we only create sort of through the eye of the model. Yes.

AUDIENCE: Do you understand the guy who's like constantly bankrupt is worse off?

ROBERT TOWNSEND: Because before we would have forgiven the loans, sent him to a minimum consumption c bar and let him start over. And now, he's saddled with all that previous debt. And he can basically borrow more against it at interest rates and then have to pay off in the future.

AUDIENCE: Oh, I see. Because we expanded his ability to refinance, now he's worse off.

ROBERT TOWNSEND: Yeah, exactly.

AUDIENCE: That's different--

ROBERT TOWNSEND: If the government exempts your loan and you stay in the house. But--

AUDIENCE: So if would be different if you were borrowing from a money lending who is going to break his knee.

ROBERT TOWNSEND: Yeah, so you could imagine default doesn't work like this. But that's the way it works in the model. Anyway, and then you know after they climb out of the cliff, you get a repeat in terms of-- so if you looked at those welfare numbers, it's almost misleading. In fact, I thought it was wrong for a minute. This looks like it's just-- the bankruptcy region, there are these huge gains. No, the bankruptcy, these guys are losing. They're negative numbers down here.

What happens is very close, but not very adjacent to that bankruptcy reason, the welfare gains just you know rise like a tent. It's actually not vertical. It's just coming up real fast and then coming down. It's a very steep tent.

So these guys are really in the program. And then these other constrained people, they're kind of liking the program. All these people are actually relative to what? Well, basically, we use the model to look at an alternative program, where they just get a lump sum transfer. And we actually do this consumption equivalent calculation, which is how much of a transfer would you have to get to have a welfare gain equivalent to the one you get under a Million Baht Program.

And another way to say it, in this region, the government could have saved money. It could have gotten households the equivalent welfare gain without actually putting so much money in the village. But it's very heterogeneous. Some people are tempted, and the abstract on the paper actually talks about it not being a great program. But I don't know, strictly speaking, Pareto criterion, we can't say that. 24% of the population really love this program. For them, it's wonderful. But more than the majority would have preferred a different program. Yep.

AUDIENCE: So that type of comparisons is just looking in people's welfare in the model, not thinking about cushion being administer, for instance.

ROBERT TOWNSEND: Oh, yeah, we're extracting away from the people who are financing the programs. So there's all those tax distortions being created to generate the revenue to fund this program that aren't involved in this calculation.

AUDIENCE: Oh, I was just thinking organizational capacity. So maybe one reason for loaning to each village is that by head count, there's be a lot more variation--

ROBERT TOWNSEND: Maybe--

AUDIENCE: Than a simple version.

ROBERT TOWNSEND: And then I just want to say that we come back to those IV regressions, two comments. First of all, what do these coefficients mean? Well, you can think about them as impact of the program. But that's kind of like a funny number, because it's an average over all these different people who got treated differently and had different views about the program. So we're lucky in a way that we saw this big salient number. And it helped us figure out a model where something like that could happen. But then with that model, we realize this number is not a uniform benefit for all the households. So in other words, a big advantage of a structural model conditioned on believing that you're getting the structure somewhat right is that it allows you to really look at the distribution of welfare gains and losses.

The other comment is more favorable. Since we did those IV regressions and I reported them, we can run-- we can generate-- we did generate-- the data from the model over and over again, as I said, and then do exactly the same IV type regressions on the data from the model and take-- so we're comparing apples to apples. It's kind of a funny intermediate criterion, but it is at least a consistent criterion that's used on the model generated data as well as the actual data. And the fit is actually pretty good.

So I think that's it. Thank you.