PROFESSOR: So we had a big lecture last time on perfect insurance and implications for targeting, and policy, and evaluations. There is a series of lectures now that have to do with working off of that basic risk sharing framework.

Today we're going to talk about the implications for what we should see in rates of return. And then next class, we're going to talk about alternative models of what we might see in the consumption data. And then the lecture after that we're going to return to the perfect market model and think about labor supply. And then finally, we're going to get to explicit obstacles to trade.

I must say that rather than put off the obstacles, we've been noting and even talking about alternative models along the way. Because these benchmarks are sometimes rejected and you're going to see that a bit today as well. So this perfect market standard within villages that has to do with relatives and friends interacting with each other has implications for how you adjust for risk. And we're going to do that with a vengeance.

But it will turn out that it does-- that framework allows us to distinguish aggregate from idiosyncratic risk. And the idiosyncratic risk should not be showing up in rates of return, but it does at the end. And then I'll say a few words about what else we know about the tie data in terms of the decomposition of volatility linking back to a paper we saw the very first lecture on quantifying the growth rates by sector. But we're going to do it by wealth. At least I'll say a few words about it.

And then I'm going to toward the end of the class sort of return to this idea that either things aren't perfect and/or people are making mistakes. And we'll do that in some unusual contexts, one that has to do with data from Sweden, an unusually rich database from Sweden, and another that has to do with the US airlines.

But the basic frameworks that are getting used, whether in the developing context, or Sweden, or US airlines, are basically all the same. So there is a strong complementarity across these papers. This paper with Chris has been in the works for a while. But we are now able to use you know 13 years of monthly data. So we are doing better than before.

So we want to talk about risk and return of productive assets in developing countries. We already looked at some of the slides showing how rates of return differ a lot across different households, differ in large part by wealth and how poor people have relatively high rates of return. And they're saving. And they're accumulating assets and so on.

Another way people look at rates of return is to say, you know, well, some people are more talented than others. So if we look at rates of return, we get some notion of underlying productivity as in TFP and so on. We've had many lectures about TFP and coefficients in front multiplying in front of production functions and so on. But correct me if I'm wrong. Nowhere in any of that stuff did we ever adjust for risk.

So what's the intuition for adjusting for risk? Well, if you want to look at compensating differentials, the higher the risk of a project, the higher the rate of return ought to be. Other things equal to compensate for that risk. You know, like junk bonds have higher yields and so on and so forth. So maybe these guys with seemingly very high rates of return actually are just doing risky things.

And indeed, if enough of them do risky things and on average some of them succeed and maybe that would generate the high rate of return. So we really need to think about how to measure the risk. And we're going to rely on finance theory. But I'll be very careful to sort of build into it the economic models. You may learn a little bit of finance jargon along the way. But I'm not going to start with that. I'll get you up to speed. It's really not that hard when you look at the equations.

OK, so what do we find? We do find that higher risk is associated with higher rate of return. But the risk that we're going to be concentrating on in the model is not idiosyncratic risk. It's village level risk. So when you think about households together pooling resources, as in a risk sharing syndicate, they would care about say aggregate consumption should move individual consumption. And idiosyncratic income shock should not have an influence on consumption.

So when is something really valuable? A project is valuable when it has a return, which is negatively correlated with the village average return. So that the resources come in when you need it. So the quote market, the way the finance guys would use it, for us is the village. And indeed, the more a household has a return on its projects that are correlated with the village average return, the higher is going to be the measured expected or average return on that household's assets.

And we have all kinds of ways of adjusting for human capital and so on, which I'll get into. As I said though, it does in the end leave some idiosyncratic risk. And that's also generating higher expected returns. Aggregate and idiosyncratic risk is negatively correlated with age. So the older a household head basically, the quote, less risky things that they're doing.

But both in the idiosyncratic and aggregate sense they tend to do things which are less correlated with a village average. It's negative with wealth. Meaning the lower is wealth, the higher both the idiosyncratic and aggregate. So this gets us into poor people and how exposed are they to risk? We've seen a bit of that already in risk and return in village India that there are some landless laborers and labor income that tends to move consumption around more.

This is a bit different way of measuring it in the sense that we're not going to look so much at the consumption data. We're going to assume everyone is in this risk sharing syndicate. But we turn around and see what is reflected in rates of return. And in particular, relatively poor people are more exposed to risk. But it's not just their own idiosyncratic risk. It's the aggregate risk.

And there is a little bit of role for education. But it's kind of problematic. And there's some selection bias that I might talk about when we get there. And then this is maybe the punch line. If you want to think about talent, then you want to adjust for all sources of risk. And when we do that, we find who are seemingly the really productive people in the village.

And all of a sudden, it's the higher wealth people who have the higher risk adjusted rate of return and female headed households. So relatively wealthy females actually dominate in terms of getting a higher rate of return. So we have a basic model. They're going to be j households and l production activities. Each activity utilizes capital as an input. And let's simplify and say all these technologies produce the same thing, consumption goods. And we'll simplify further and assume all the technologies are linear. Yeah, Matt.

AUDIENCE: Can I just actually ask a couple of things about previous slides? So are we assuming there's no sharing between them?

PROFESSOR: So we are again, going to take a stand on our baseline. And it's actually going to be all four villages in a province, not just a single village. But then we go back and check for villages individually. And we actually go back further and look one network at a time. The basic results don't change.

AUDIENCE: OK. Another thing which I don't really understand is when we have this result that returns risk adjusted returns are lower for those with lower [INAUDIBLE]. How do we think about in terms of one story's that the people with low initial wealth, they're very productive people amongst them. But [INAUDIBLE] get sufficient size of their business. So their marginal returns might be really high. And the average returns as well, depending on--

PROFESSOR: Yeah, so this paper isn't going to help much with that. Because we're going to assume linear technologies. And margins are going to be equal to averages.

AUDIENCE: OK.

PROFESSOR: And more generally, you know, wealth is clearly an endogenous variable. So we're not saying high wealth is causing people to have-- I mean, it's more likely the other way around. But yeah, we've gone through a series of models where wealth helps overcome constraints. The first four lectures were like that.

So in a way, this is another fact finding mission. But it's not just data summaries. It's looking at the data through the lens of the model. And you have to make some assumptions to do that. OK, so we have j production activities. We're going to start the economy off with some aggregate wealth, w. And those are basically the projects, the assets, or the trees if you like. And those projects are going to have rates of return, the fruit.

So at the beginning of the period, the total sort of wealth available is the sum of the trees and the fruit. So this is a standard kind of solo reversible capital assumption that the assets can basically be sold. There's no irreversibility. You can convert them into consumption if you choose to. So we're going to have a social planner sort of directing the show, which is not to say this might not decentralize in some way.

That said, I want to emphasize it's not like finance in the sense that people participate on the New York Stock Exchange or the financial markets, and they're buying and selling stocks. We don't even have the data of all these capital goods. That's not the mechanism.

This planner's problem may be decentralized in some way. But it's just much more straightforward, just like when we did the risk sharing stuff not to necessarily be forced into some interpretation. We're just going to solve the planning problem directly. So somehow or other, a household is assigned. Household j is assigned capital and is operating project type i. And this is the net rate of return. Notice that we allow individuals to vary in their rates of return.

The planner could take advantage if it knows that some people are better than other people at doing things. So we're not just getting rid of talent. We're embedding it at some level in these rates of return. And then they have basically the tree and the fruit. That's their starting point. And then they get taxes and transfers, these Taus.

I mean, here it's written positive as if it's incoming and adds to the resources of the household. But in fact, like a risk sharing group you can give and get and contribute to the social fund. So the planning problem is to maximize discounted weighted expected utility. So these Lambda j weights are the pareto weights for households type j. They're not necessarily the same over all the households. And this whole term here is the utility function, which as noted, is just the return on the trees plus these transfers.

And then discounted expected return for tomorrow I usually write as Beta. But anyway, it's an intertemporal discount rate. And so you have the value function today given wealth today and expected value of wealth tomorrow w prime. And wealth today can be allocated into basically consumption. So this is just summing over not only over all projects for household j, but all households j. So this is again, just consumption.

So consumption plus capital allocated tomorrow equals total wealth, pretty standard kind of growth market structure. And what is total wealth of the economy? It's basically not only the fruit of the trees, but the trees themselves. So you can take this expression and put it on the right hand side if you want. Because they have pre-existing capital. They have the rate of return on that. You've got total resources available. You can allocate it to consumption or set some aside to continue the capital stock for tomorrow.

And so we just substitute in the expressions for tomorrow's wealth and today's wealth for that matter. And then this odd looking thing here is just a substitution from the previous equation. But there's a common term on both sides. That's why it looks kind of weird. Because this is really 1 plus r times k over here. But there's also a 1 plus r times k here. And the rk things are canceling out. So this is just straightforward algebra.

So we're going to maximize this thing by choosing the transfer as Tau j and the assignments of tomorrow's capital stock.

AUDIENCE: Was there supposed to be uncertainty? It's just their return on capital?

PROFESSOR: Yeah. Yeah, there's no preference shocks here.

AUDIENCE: What's the process [INAUDIBLE]? [INAUDIBLE]

PROFESSOR: Don't have to specify yet.

AUDIENCE: OK.

PROFESSOR: That's probably a good thing. OK, so if you looked, we've actually seen this before. But it's explicit here again. Consumption and capital are interrelated through this resource constraint. So when you take a derivative to maximize with respect to the transfers, you'll pick up a derivative in the utility function. And the sum of total transfers is down here in the resource constraint. And there's a Lagrange multiplier in front of it, right?

So you're going to get margin weighted, margin utility of consumption today equal to a Lagrange multiplier. And when you differentiate with respect to the choice of tomorrow's capital stock k prime, well, that's going to enter in tomorrow's value. So you'll have a derivative of the value function and then internally this sort of rate of return. And this prime thing is also entering into the same resource constraint. So you're going to get the same Lagrange multiplier as on consumption.

So here is a summary of those words. Weighted margin utilities are equated across all the households to a common contemporaneous Lagrange multiplier. And then this is kind of a-- it's an exactly Euler equation for capital accumulation. This is the price of capital today. It's how much consumption you're giving up. And this is expected discounted value of having capital tomorrow. It's the derivative of the value function times the rate of return, depending on which project i and J is being chosen.

Note that this is for every project, every capital stock of type i for household j. Now, so far I've talked about real assets. But it's not too hard. And we do actually allow external borrowing. You could imagine a village is able to borrow in land. Or the whole set of four villages have some relationship with the rest of the external economy. That still leaves this sub problem. I mean, throwing these things in might deliver some extra conditions that we're ignoring. But what we're doing is not inconsistent with having this equation.

All right, so if you just take this one and divide both sides through by Mu, you get this. And then you take the expectation operator in front of everything, including Mu, even though it's not random from the point of day t. And then you can rewrite this guy here as something called m or m prime, since in effect it's for tomorrow. And this r prime ij is basically capital R is little r. One plus r is sort of the gross rate of return.

So this is hopefully not mysterious. It's just basically a relabeling of variables. And it's like the Euler equation. But if you want to impress your friends and bewilder your enemies, you call this thing the stochastic discount factor. But all it is is the way society as a whole is discounting the future. It's tomorrow's sort of shadow price of resources relative to today's. So you're discounting tomorrow.

And it's random. Because these rates of return are random. And so the amount of resources available in the economy is likely random. So it's stochastic. There's a danger in talking to someone about stochastic discount factors. Because the conversation will keep going. And you may learn more or get bewildered. But anyway, as promised, we'll keep track of this.

So it's true for all households, all sectors. And the finance jargon is we just looked at something called the pricing equation in the consumption based asset pricing world. But for us it's just an Euler equation holding at say the village or let's call it township level.

Another thing you may or may not know or remember from finance is you can take assets and bundle them. So you can have a particular asset held by a particular household, a project. Or you could take a collection of assets. That Euler equation has to hold for any asset individually and all assets collectively. You can bundle, package, rebundle anyway you want. And anything that's actually held-- by the way, there's sort of out of equilibrium projects we're not seeing. They wouldn't satisfy the Euler equation. They have a low rate of return relative to the other ones.

So we're only seeing these equilibrium ones that the households actually hold. And of course, we're looking at the data at what they actually hold. I'm mentioning this bundling because-- and you'll see it again. But we might as well get our minds around it here. We are going to look at a household's rate of return. And typical households, as you know, you've already seen this, are doing more than one thing. Some of the incoming money is from labor supply. They have crops, maybe multiple crops. Some of them have both that and a business and so on.

We try very hard to assign these assets individually to particular activities. But it's kind of treacherous. And we're likely to make mistakes. What do you do with a pickup truck for example? Is that just used in farming? Or they're going to the bank. They're taking the kids to school. And so rather than face all of that we just aggregate up over all the activities any particular household is doing and call that a collection of assets. And the Euler equation should apply.

Now, remember a little bit of your sort of econometrics with means and variances or easy mistakes to make. The expectation of two random variables is not just the product of the expectations. You've got to adjust for the covariance, unless this happens to be zero. But this covariance is going to be crucial for us. Because you can see already it's the covariance between dare I say it, the stochastic discount factor, which reflects the community's relative valuation of resources and the particular return on project i held by household j.

So since we already had this equal to 1 as the Euler equation, right? So now we just substitute this stuff into here. And we get this. And you could basically divide through by the expectation of m prime. You're going to get 1 over that here, get rid of it here. Get the covariance divided by it here. And then you divide and multiply, which doesn't do any harm, by the variance of m prime and rearrange terms. So you get this thing.

So we're finally arriving at something close to what we want, which is the expected rate of return, which is going to be in the data, the average return for household j over all the time periods. Not quite. It sort of looks like a linear function, something with an intercept, and then Beta Lambda. Beta, and finance guys talk about betas all the time. Beta for us just means sort of the normalized covariance between the returns on ij and the stochastic discount factor, which I already pointed to above.

It's normalized by the overall variance. So that's kind of the risk factor, or the quantity of risk. And this other thing could be called the price of risk. But it's nothing other than this. It's the variance divided by the mean. It almost looks like a coefficient of variation, like Sigma squared divided by Mu or something. And there's a name for it. And I'll tell you what it is in a second. All right, so this is what I said already.

Quantities, prices-- oh, and one more thing. What was that intercept? Well basically, let's just define this thing, this Gamma, well, it was already defined to be 1 over the expectation of m prime. But let's call it the risk free return. First of all, why is there no risk? Because all the variability in the economy is in m prime or the one that we care about for that equation. And we've already taken the expectation. So there's nothing random left. So this is a number. It's not something stochastic.

It should be obvious. But I always do it myself, which is take that to be an asset, plug that back into the asset Euler equation that we already have and thus convince yourself that that equation applies for the risk free return. And you'll get this right back again, OK? So anyway, the point is risk free because the covariance of this risk free rate with m prime is 0 by construction.

So let's just look at this again. This is the rate of return of the mean, or average, rate of return on project ij. Subtract off the risk free rate. This is the return differential. This is how much the-- in finance, it would be like the junk bond premium over and above treasuries. And so that difference is sort of the premium that you get for holding a risky asset. And it's linearly related to the covariance of that asset with the market.

So again, I'm going back and forth a bit more with this finance language. But it's just all right off the Euler equation. This is just standard sort of manipulation. So we can even make it more obvious if you're wondering what m prime really is. Let's go with quadratic utility. So this is the form of it. The community's value function as a function of wealth is just basically the squared of the difference between wealth and some subsistence level with a coefficient in front of it.

So of course, if you take the derivative of it with respect to w the twos cancel out. And you get eta times this difference, which is linear, by the way. That's huge. So marginal value of wealth is linear in wealth. That's what you get at a quadratic. And then you just substitute our expression for w prime tomorrow into that expression and do a little relabeling.

So all of a sudden, this looks like the return on the village average and the village average capital stock. OK, the village average capital stock is the sum of the capital stocks allocated over all the projects and all the households, as it should be to be in total, right? And then what is the rate of return on it? Well basically, it's like a capital weighted return. So the total return is coming from a bunch of individual projects. So you just take the return on all the individual projects but weight them by the capital being allocated to those projects and renormalize by the total capital stock. So it's a portfolio weighted return. That's just a definition.

But of course, then it looks like this. And you take this up here and this here. And you realize that there's things canceling. And that's why it's really just equal to this. But we'll use this now. So m prime, which is the derivative tomorrow, is just this thing divided by Mu. Remember we have to take the current shadow price and divide through. And then this is all additive and linear. So we get this.

And now, m prime here looks like this plus this. This looks like a constant term that doesn't depend on the rate of return. And this is a coefficient premultiplying the rate of return. So it's starting to look like a pretty convenient equation for the stochastic discount factor. It's linear, linear in the village rate of return. And again, we have this thing we already derived, which was a manipulation of the Euler equation. And everywhere you see m prime, you start substituting a minus brm prime.

And then you got to remember certain rules about covariance operators and variance operators. Variance operators will square coefficients in front of them. Covariance operators will bring the coefficient outside. There's a lot of substitutions here. there's b's and b squares. And they're going to cancel with each other. And you end up with this thing.

So this says the expected return, or sample average return, should just be the covariance of the return with the market with the village average aggregate. By the way, we know how to construct that. Because we have it in the data. And we have the variance as well. And then this is just the sort of coefficient that multiplies it. But that's what I was saying about that coefficient of variation. It's a variance divided by a mean.

So where have we ended up? Well, with that quadratic we now have pretty much an exact expression that the expected rate of return on project i run by household j should just be a linear function with a common constant term that has nothing do with i and j and then the covariance of that household's return with the village average return, which is something we can compute. Because if you start thinking about this like a regression, we'll know this from the data. We'll know this from the data. I'll tell you how to get it in a minute. And we'll just estimate Lambda and Gamma. Well, not quite, Gamma we can do something special with. But that's the idea.

So I've said this already. But it bears repeating. Our stuff looks a lot like this cap m. But the mechanism is very different. We do not have households trading assets that are priced in centralized markets. We're instead optimally allocating assets across households, but then not forcing them to eat the return stream. Instead, they enter into this risk syndicate as if a social planner were reallocating consumption.

Whereas, in finance you kind of assume you're an investor. You could do means and variances and take some risks on your portfolio. And then you're going to eat it or at least dynamically optimize, eat some, save some for tomorrow. But this is very different. And it's key that they don't have to eat the returns off their capital stock or make their own independent savings decisions.

And you've seen last time building up to this that households are sharing a lot of risk with each other. And these informal village money markets are quite active. We called them a money market last time. So we've already been using that language. That's what this slide is saying. That's just a reminder of what we talked about last time.

So data, it's monthly. I don't know what they're called, four villages grouped together in a tombone. So we kind of started adopting this South Africa language. I don't know whether I like it. We're calling them townships.

AUDIENCE: It has a slight connotation [INAUDIBLE].

PROFESSOR: Yeah.

AUDIENCE: It's rural and maybe not the best place to be.

PROFESSOR: Yeah. So tombone is Thai for the fact that we picked four at random in this geographic area. But no one knows what a tombone is. But clearly, townships distinguishes it from villages. It's four villages in the data anyway. You know all about where they're located. Reminder that we started in August of '98. We're going to ignore the first four months basically. And use the data from 1990 on, 1999. We're still out there at month 174.

What we're going to use-- and Chris and I just basically redid this in the last few weeks-- we're going to use month five to 160 or basically 156 months, or 13 full years of the data at the monthly level for the 541 households who are basically have not dropped out sort of a balanced panel.

Now, about the townships, so what percentage of households have relatives living in the same village? It's a bit low into only half in [INAUDIBLE]. It's 3/4 in Logburi. And Buriram is already up to the 80s if not to the 90s. But if you go to the township level, then 87, really 88 is the lowest percent. In other words, almost every household has a relative if not in their own village then in a village nearby.

So we stopped kind of looking at networks after this. And networks confuse people, rightly so. Because who's in it? And who isn't? And is there selection bias, and so on? So we just made it easy for ourselves and said, ah, they're all really related to each other anyway. So let's use the whole sample.

But I'll show you what we go back and check. This looks like risk and return in village India. I mean, you can see the cultivation activities, livestock, fish, shrimp, wage earning, and so on. Wage earning is a bit heavier in the Northeast. And in [INAUDIBLE] only [INAUDIBLE] has fish and shrimp. It sounds like a restaurant. I know I've already told you though. That's what you're eating when you buy shrimp at Costco. And all the business, livestock is pretty heavy. And Logburi, also [INAUDIBLE], those are the dairy cattle.

So that's just a quick reminder of the "projects." My project is a cow, et cetera. It's a cash cow. That's terrible. And here is some just reminder of descriptive statistics, males, females, education levels. Some of this you've seen before, income levels, assets and liabilities, and so on. So we're going to use a total return. And we're going to include financial assets. We're not going to separate-- by the way, it's a bit problematic to know what to do with a savings account. It's a financial asset. But on the other hand, maybe like inventory, it's kind of contributing to the business. And if we didn't have it in there, then people would ask. So it didn't make much difference.

So we put in all the financial assets. Actually, subtracted off the debts too. It doesn't really matter. If you want to think about this as real rates of return, you would not be misled in real projects, real people, real assets. OK, you know we constructed these financial accounts. This is problematic. We have to sort of subtract off the cost of labor. And when it's household labor, it's not priced. This is probably our single biggest problem.

And when I show you residual rates of return in a minute, you'll see some of them are low. We probably subtracted off too much. But it's very hard to know. Because we try to grab instruments and adjust. But it's really kind of not terribly compelling. You'll see this again when we get to the labor paper.

So we adjust the income. We divide by fixed assets. That, as you know, is rate of return on assets. We may get real. The data is nominal. It's all measured in Thai [INAUDIBLE]. So we deflate by Bank of Thailand regional price changes.

This is the table of the rates of return mean standard deviation and the so-called-- so I forgot to say it. But that thing I kept saying was sort of like a coefficient of variation, that's like the inverse Sharpe ratio. So you can add that to your vocabulario. Now, each household has a sample average. So you're going to have a mean standard deviation, et cetera, those statistical moments for every single household. So you have a histogram in the population.

So you can look at the median mean if you want, if it's not too confusing, or the 25th and 75th quartiles of rates of return. Some of them are high. Many of them are not that high. But in any event, these are not adjusted at all. They're just the crude rates of return.

OK, so now we got to decide what to do with the risk free rate. Well, these are supposed to be real. So we're basically going to assume they have something like inventory that goes up with the price level. And there's nothing stochastic about it. So we're going to assume the real rate of return is zero. Now again, if we're guessing wrong about this, we're going to be guessing wrong about the intercepts. Because everything we do is subtracting off zero.

But anyway, that's what we do. The market return is the average township rate of return. And again, that's just looking at the income over all the households divided by all the assets. We're about to run a regression though. So we will take out the household's own contribution to the village average. It's like a leave out mean. So we don't you know bias it in the obvious way. Yes?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Depreciation is in here. Because depreciation is a cost that we're subtracting off in the financial accounts.

AUDIENCE: So that's where it comes [INAUDIBLE]?

PROFESSOR: Uh-huh. So these numbers are basically coming off of the income as in the income statement. And the assets as in the balance sheet, which you've seen before. Although, it is true that we cover a lot of ground. So I'm happy to remind you.

So how do we get the relationship of the household's rate of return to the village average? We run a regression. We run a regression of the household's, over all of its projects, household's return at day t. But there are lots of dates t with all those months-- more on that in a second-- regressed against the village return at day t.

OK, so this beta is just literally the regression coefficient that we wanted conveniently. Yes?

AUDIENCE: We are assuming [INAUDIBLE] is constant over time for each household?

PROFESSOR: OK, so yeah, there's several adjustments. I think it's on the next page. But anyway, so actually, what we're going to do is do five year increments. You could just take one household over all the many years and months and ignore possible shifts in their portfolio. But like you're saying, that might be a bad idea. So we redid it as if the household had returns on a portfolio from year one through five then returns on a portfolio from year two through six, seven through eight, and so on all the way through.

And it turns out that idea, as many others, come from the finance literature. So that's what we're doing. The other adjustment I should have mentioned earlier. And I didn't. But you remember when we use a quadratic utility it was a plus b times something, right? And there was no t on those things. So it was like the marginal utility of wealth just depends on wealth. And it doesn't depend on anything else.

But in many intertemporal models, those things would be moving around. So I'll tell you about that adjustment too. We actually did it all the different ways. But these are the key steps. We run the regression, or regressions, for each household. And then we get the beta for each household j. And we get the average return, either over all the years or over these five year increments. And then we run this regression essentially with this already coming from a previous regression using the panel.

This thing is just a cross-sectional regression. There's no time date here. This is the average return. This is the extent to which the village that household's returns are correlated with the village average. And we'll get expressions for lambda, alpha, and so on. And there are restrictions. Although, you no doubt can't remember seven and a half slides back that and this lambda ought to be equal to something like the expected market rate of return. Just take my word for it.

We don't impose that restriction. So here's the first result. Beta is on the right hand side. And it's being regressed against the household's mean rate of return and three of the four provinces. It's quite positive and significant. These are the regression coefficients, the lambdas basically. And this is the estimated constant. And r squares. You know what it is. It's kind of worth noting. Buriram fails and has a very low r square. And Buriram is one of the places I've mentioned, where the farmland got plowed over with roads. And the town was being built up. So there arguably were a lot of changes in occupations. We'll try to fix that. Yes?

AUDIENCE: Is this for one five year period? Or is this like an average of the--

PROFESSOR: This is actually, with treating each household multiple times.

AUDIENCE: OK, so--

PROFESSOR: I think. Let me just see.

AUDIENCE: [INAUDIBLE] it's like putting 20 years.

PROFESSOR: No, it's not yet.

AUDIENCE: So that's like the whole thing.

PROFESSOR: This is the whole thing.

AUDIENCE: So do we have to worry about [INAUDIBLE] issues or [INAUDIBLE] problems with us? Or is that not-- do I misunderstood how the regression [INAUDIBLE]? So is there some possibility that there's like going to be trends in the errors that are not picked up? Because you set the constant. [INAUDIBLE] So we don't have to worry too much about using the whole--

PROFESSOR: We worried quite a bit about potential econometric biases and so on and so forth. Then we went back and read the classics in finance, [INAUDIBLE], French, all these guys. And we're doing what they're doing. So we felt better.

[LAUGHTER]

AUDIENCE: [INAUDIBLE] if we regress the errors on the [INAUDIBLE], does it pause the test? Or is it--

PROFESSOR: Oh, we haven't-- yeah, I don't think we've looked for autocorrelation and stuff like that.

AUDIENCE: [INAUDIBLE] we get hammered three hours and make that [INAUDIBLE]. That's what we're doing [INAUDIBLE].

PROFESSOR: The HH factor. OK, we'll check that. Yes?

AUDIENCE: So do you think it's worse to joint tested r [INAUDIBLE] j for each household, like the finance guys?

PROFESSOR: We're about to show you the alphas.

AUDIENCE: OK.

PROFESSOR: I haven't gotten there yet. We'll see when I get there whether you have something different in mind. Anyway, so if you want to visually see sort of a mean variance frontier, you the beta is down here. That's kind of the risk factor, if you believe this model. And the expected returns are here. And you can see a pretty close fit.

Now, this I've shown to the finance guys at lunches and things. And they're like, wow. Because this is actually pretty good relative to New York financial markets. Oh, it's an easy target, OK. So so far we are showing that just to repeat, when someone has a higher return, it's partly not talent. It's just that it's correlated with the village average. And it's high because you need to adjust for the risk.

Now, we can do this at the village level, in which case nine of the 16 villages we're going to have this positive correlation that the theory says ought to be there. And when we go to the network level, we have five out of nine networks. And I'll show you about the adjustments, remind you here's a kinship network. 206 is married with 207 in some way. And 214 has all these. So can it literally connect the dots? Here's a big dynasty down here. Here's a little dynasty there, another little dynasty there.

So we went through all the villages. And they had to have enough data to do this. So we ended up with networks. Now, some villages have more than one, et cetera. So of the nine with enough data, we're picking up the theory doing well here and not in some of the other ones. So as I said, we go micro, go macro, zooming in and out. There's a lot of explanation of the rate of return. That has to do with the village average. This is villages.

OK, so this has to do with changing projects, basically. So it's what I said. We use these sort of five year intervals and redo it. And kind of reassuringly, Buriram pops back in there. Because Buriram was the place for sure that I know that they switched occupations quite a bit and moved from lower to higher rates of return.

So it's kind of reassuring, given the previous work to see that this adjustment allows us to recover something. Although, the r square is nothing to brag about. Now, another hard thing, easy and hard at the same time, is human capital, not physical capital. So as you know, a lot of these households have laborers. They're earning wages. And you say, well, what's that the fruit of what tree is that?

Well, the tree is human capital. Of course, you never really see that. You can write big H. And you have the notation for it. But what is it really? Unfortunately, we don't need the measure of their stock. All we need to do is measure the flow. So basically, it's not quite right. But largely, the return on human capital are the wages these households are earning.

And we will now put that in the equation as well.

AUDIENCE: Can I ask you a question about [INAUDIBLE]?

PROFESSOR: Yep.

AUDIENCE: So the way I understand this is rather than something like skill and talent, you're saying people [INAUDIBLE] returns, because they're somewhere all along this return frontier.

PROFESSOR: That's right.

AUDIENCE: So once we think of this model, like how do we then respond to the big questions of the puzzles of heterogeneous returns? Do we know what--

PROFESSOR: Well, we don't know yet. I haven't shown you. Maybe there's no puzzle. That's the point. People aren't adjusting.

AUDIENCE: But there's still--

PROFESSOR: But anyway, not to keep you in suspense. It is going to turn out that there is a residual. And then we can see what it's related to. And I actually said it's wealthy females who have the high residual rate of return.

AUDIENCE: So essentially the question is, given that they're on that [INAUDIBLE] regression line, which ones are where and why? And part of the answer might be explained by--

PROFESSOR: Well, there is part of it that we're not explaining, even in the definition of the model, in the construction of the model, and another part that's outside the model altogether. What we're not doing very well is why household j is running project i and some other household is not. We're just acting like the social planner knew and efficiently allocated projects across households.

So that part is sort of related to the talent thing. But we never see it. We can only look at the residual return, which actually ought to be zero if the model is true. Anyway, so another adjustment that I've mentioned is that a plus bw thing. And you might want to have time subscript. Sorry, I'm jumping ahead.

Here's the human capital thing. So now the rate of return has two factors a la, basically, [INAUDIBLE] type thing, OK? One is the physical capital. And the other is human capital. And we can try to get both betas. Another is this problem with the varying shadow prices. Well, another way to say this mu, act like mu is a constant. Mu t is the value of resources at day t. In the derivation, it was always today versus tomorrow. And it kind of got suppressed.

But suppose a village has a run of bad luck, like [INAUDIBLE], with those shrimp ponds going bad. Then they have less and less income. So you would think mu t is going up. Because the marginal value, the shadow price is going up. So we really do need to adjust for that. But there again, there's a way to do it.

It turns out that time varying stochastic discount factor is related to the consumption wealth ratio. And in turn, the law consumption wealth ratio depends on three factors, consumption, physical wealth, and basically, labor earnings. Rather than getting bogged down in the details, I think intuitively you could just think, well, you know, there's physical capital. So basically, the consumption wealth ratio is kind of picking up the degree to which they have wealth today relative to consumption. So that's kind of like an obvious argument to have in today's shadow price.

And these things that determine the consumption wealth, or the law consumption wealth ratio, kind of look like the accounting variables that we've had before. You know, how much you eating? Again, how much is your physical capital today? And then how much is your labor earnings? So that's kind of the spirit of it. Again, do our homework and we can find this exact derivation in the literature.

So this consumption asset ratio just looks like this. So we find that in the data by regressing it on to the arguments that I just mentioned, consumption and so on. And then we substitute in that estimated value into the regression that already was adjusted for human capital. So now we just did a double adjustment. We put human capital in the rate of return on human capital. And now we have extra data terms that have to do with the covariance of those physical and human capital returns with this.

So basically, we're just interacting. We have levels of things. And then we're interacting it with the rates of return. So we do that. We end up with this. It's a longer equation. And we run it. And we're still getting the beta on physical capital. It's quite hardy. It's quite robust. Human capital is in there sometimes, sometimes not. But at least putting it in doesn't harm this.

In other words, as I said, we do all the adjustments that we can. And we're still getting the basic result back out. Now, you want to know about risk? Well, this is the equation that we're running. So you can do a decomposition of variance. The total variance in returns is just the variance having to do with the regression times the variant, or basically, this thing. This is the contribution of market risk to the variance of household j.

And there's a residual here that we were just chatting about. So we estimate its variance. And so that's the contribution of idiosyncratic risk to the overall variance. So now, we have this measure, finally. But note that it didn't just come from using the data. We used the structure of the model to back it out.

And that idiosyncratic part is large, at least half and in many instances. Substantially more than half of the risk is due to the idiosyncratic component, rather than the aggregate. Now, one word of caution. We still haven't rejected the model. No one in finance says there aren't idiosyncratic returns. The question is whether they're priced.

If you believe in the capital asset model or the consumption based risk sharing model, then idiosyncratic returns can exist. But they're completely pooled. So this per se does not reject the model. But when we now stick the variance, sigma j, into the benchmark model. And we're still doing OK with aggregate risk. But here's the rejection, if you like. The idiosyncratic risk is positively associated with expected rate of return and for all the households.

Now, this is the if you like, the disappointing part, or the tension, which is you go to all this work with a benchmark. And it seems to be doing really well. But there are features of the data against which it's not doing well. So it's clear that it's not an entirely perfect markets world from the point of view of the rate of return of these assets.

But is the glass half empty or half full? Because it's also true that a substantial amount of the rates of return had to do with market risk. And most development researchers would have ignored that market part. I think it's fair to say, even though it's kind of a natural corollary to risk sharing.

So let's go back to these alphas. So Jensen talks about having portfolio managers who are somehow smarter or better than others. And they have abnormal returns. I mean, I don't know. This is kind of hard for economists. It's like, well, if you do so well, why don't everyone just go in there? Like, I never make money. I don't trust my instincts.

But I'll tell you story about Walmart. So one of my colleagues who was at Chicago, a sociologist, was in Arkansas. And he saw this big square flat building in northern Arkansas. And he sort of checked it out and said, well, this is a really novel way of doing business. Well, it turned out to be practically the first branch of Walmart. And he bought stock. He did quite well obviously.

Anyway, so maybe some portfolio managers are better than others. Maybe some households are somehow better than others. And we've now got a way of adjusting for all those risk factors, both the idiosyncratic and the aggregate. By the way, you will be relieved a bit maybe to know that when we rank order the households by the unadjusted rates of return and the risk adjusted rates of return, the ordering is largely preserved.

But the orders of magnitude are not. The second point is it's not just completely downshifting the distribution after you subtract stuff off. The tails come in. So it has a different level of skewness. So it's kind of a serious filter on the data. But some of the earlier results that poor households have high rates of return and reinvest in their own projects, we're going to rerun those. But we're pretty confident those kind of results don't go away. And I haven't misled you in the earlier lectures.

So finally, here is running these regressions, where we have aggregate risk, also allow idiosyncratic risk, where we have the betas on aggregate physical capital and the coefficients on sigma. So the first two slides I said households who are older take on less risk, which means basically, both the idiosyncratic and aggregate risk factors are negative. You might be willing to tell a demographic story for that.

We don't really have it embedded in the model. It's more like a residual or behavioral thing at this point. Because these residuals aren't supposed to be there. And here is basically the wealth again, negative. So poor people have higher correlation of their activities with a village average. And the residual have a higher residual variance. I'm going to come back to that. Poor people seem to be exposed to a lot of risk.

However, if you take out both sources of risk and regress on demographics and all these things, wealth finally goes positive. And also, this is a dummy for male, so to speak, is negative. So that the female headed households and the high wealth households are the households with the highest sort of residual rates of return.

OK, so now what I don't have in front of me unfortunately, we didn't quite get close enough. But you may remember that QJE paper from lecture one on volatility, growth, finance, what's related? And we said, well, we've got to look at sectors. We've got to figure out how much each sector is contributing to overall change in value added. And there was some really cool decomposition.

So we're running that on these Thai data. So we can see at the level of income how diversified are these households? It's similar, but not identical to this CAPM type approach. And so far at least, we're getting the same answer, that it's the relatively poor households who are specialized. And the things they specialize in have high volatility, and so on.

So poor people seem to be exposed to risk at least on the income side. So here's moving to Sweden, more sort of standard asset stuff, which John Campbell and co-authors have done called *Down or Out.* And what they're trying to do is very similar. They're trying to look at household management of financial affairs using measurement in this amazing Swedish data. Sweden is a socialist country largely. And they measure everything. And everyone has a common ID, and your whole portfolio, labor supply, everything.

So I've been collaborating a little bit with a Riksbank largely influenced by this paper. Now unfortunately, it takes us away from development. That's not my point. My point is the commonality. They're going to look at mean variance frontiers and see what households are doing. In particular, what are the poor people doing or the rich ones? The focus is a bit different, which is households can make mistakes.

In other words, let's just assume the model is right. And the mean variance frontier is what it is and then see households in the interior who could get less variance or a higher mean if they only move to the frontier. And then they look by characteristics, like education, wealth, and so on to see who is under diversified-- those are the down people-- and who's not participating as in the stock market. Those are the out people.

And hence the title, very clever, *Down and Out.* And they measure the welfare costs of those deviations. As I said, they have amazing data. And here you can see sort of how households manage their growth wealth in Sweden. This is the cash. Oh, look at that. Poor people have a lot of cash. I think cash just doesn't mean Swedish krona. I think it also means bank accounts.

But stocks, poor people don't hold stocks. Mutual fund has this funny hunch-- now, it's interesting that this is real estate. So this is both financial as well as real. And on a household side, there are some businesses here actually. But most of these households are salary employees. But they can hold durable goods in housing.

If you take the physical housing wealth out of it, then the cash thing is still going down. The mutual funds is largely flat, except at the very low and the high end. And then you see this stock. It's the 90 percentile up are the people who have reasonably substantial holdings in the stocks. I ask myself these questions a lot in Thailand. Sort of maybe they shouldn't be investing in their own projects. Maybe they should be investing in Bangkok money market funds, safe funds.

We should be sort of looking at the whole mean variance frontier for these. And we are asking them now. We're actively involved. I was on the phone again last night. We're about to be asking them a bunch of questions about, are they aware of other assets? But in Sweden the point that Campbell is going to make is you've got those categories. And households can make mistakes by picking volatile assets, holding a concentrated portfolio, picking correlated. See the parallel here with what I was saying happens in the Thai villages in terms of relatively poor people not being diversified and for some reason doing riskier things.

Or if you go in the stock market, you could pick your own stocks. It's not a particularly clever thing to do. And this is the picture that makes the paper. So here's the sort of mean variance frontier, standard deviation in means. Here's where they ought to be. Adjusting or not for international asset holdings, that's another. They do both. But they're basically saying, small guys in some sense are unsophisticated and cautious. They're pretty far off the frontier. But they don't have kind of a lot to lose.

And these richer guys up here, they're more sophisticated. Their line is quote unquote above this brown one. But their gap is huge. So I'll spare you the details. They're actually measuring the welfare loss. And some of these richer households are-- because they are not diversifying-- actually have fairly substantial welfare losses.

Now, John goes behavioral in this and maybe rightly so. He's saying should we really be telling households in Brazil about investing in whatever is [INAUDIBLE]. I can't remember the acronym. Or are they going to make mistakes and they're better off not knowing? There was a question.

AUDIENCE: This is already answered by the model, but this is welfare. This takes into account the marginal utility of wealth.

PROFESSOR: Yeah, actually, I'm sparing you a lot. Oh, I should say. This is on the reading list. But I knew I wasn't going to have time to cover it in much detail. So I picked like six slides. But I'll post the whole paper. If you're interested, you don't necessarily want to or have time to read the whole paper. You might as well look at these lecture notes that I created previously. And the same thing for this paper, which is [INAUDIBLE] with [INAUDIBLE], *Dynamic Risk Management.*

I've only got a few slides. I think I'm going to jump and show you something. Fly Southwest.

So this guy's loading that airplane that capital stock up with this variable input, the fuel. They buy forward positions. Because the price of gasoline moves around a lot. So they basically buy forward, OK? They commit now to buy gasoline at a pegged future price. Now, what happens if the spot price is lower than that? It's like, oh, darn. We should've waited. Actually, you could renege on that contract.

And so what do they do with those things? You've got to post collateral against these forward gasoline hedges. Where are we going with this? We're going to get to sort of thinking about something that Chris [INAUDIBLE] was doing in Ghana, which is, do you want to use your scarce collateral to borrow and get more aircraft with potentially high rate of return? Or do you want to use your scarce collateral and tie it up hedging against future states of the world s prime?

And this paper is about how in practice this may not happen. Now, when does it happen? Where is it? It happens when the airline isn't doing very well. So if you think about relatively poor households, it's like, oh, they're behavioral. They're not buying insurance. What's wrong with them? They need it most. They've got the most marginal value. Well, no, not necessarily. They may have better use for their existing wealth. So we're getting again, very explicitly now into obstacles.

So when you start writing down models with obstacles and barriers to trade, you can get at this issue about borrowing to finance projects versus borrowing. So maybe it's still possible that those relatively poor people who seem to be doing really risky things and aren't hedged against it, maybe there's a reason for that.