

MITOCW | 9. Labor and Development

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

ROBERT So today, we're going to talk about labor and development. There are several pieces to this, just to sort of get into it. One, number one, back to risk sharing, full benchmark, but this time with endogenous labor supply and wage risk. And I'll mention sort of some tests using that framework. I dumped a lot of the excess slides, because I don't want that to be the sole focus of the lecture today, although it's kind of the obvious starting point.

With those lotteries, basically, you end up dealing with a non-convexity. And then that has a lot to do with estimating elasticities of labor supply from micro data versus what you see in aggregate data. And we've learned a lot about that. So there is sort of the complete market's version. And then, a la lecture last time, the incomplete market's version, where the aggregate elasticity is higher, sometimes a lot higher, than it is in the micro data for reasons of the aggregation. And since we've talked about zoom in, zoom out, you know, micro, macro, that seems like a relevant thing. And then we, in fact, end the lecture with Seema's, Jayachandran, paper on labor supply. And I think it was India, with rainfall and agricultural shock.

So sharing wage risk, I just said this. We have non-labor income, which is like the y moving around we've had before. But we also have income from wages, but hours are endogenous. Labor is not supplied inelastically. And we're going to allow for heterogeneity among individuals in a household and heterogeneity across households.

So let me just lead off with an example, because I think it makes the points pretty well. And then we'll generalize. So you have two people, or two households. One is risk neutral. The other risk averse. Why is this an obvious example? Because the risk neutral guy should absorb all the risk and would in every setup that we've seen so far and take it away from the risk averse guy, who's got a strictly concave utility function.

However, this guy cares about leisure and consumption, not just consumption. So typical Cobb-Douglas specification with some curvature, making him strictly concave guy. If you don't consume leisure in your work, that's supplied. You get this wage, W_2 . When we generalize, we're going to allow different wage processes for different people, not modeling necessarily why they have different talent. But for now, there's only one guy supplying labor. Again, y is this non-labor income. T is the time endowment. Consumption good is the numeraire. Wages are real.

And there is a budget constraint for this two-person programming problem, which is like a small open economy in the sense that it's partial equilibrium, and they face these outside wages. In fact, these wages may be moving around. It could be a village with two people, something like that. Wages are determined in the district or national level.

This is the standard budget constraint though. Call this thing full income, if you remember Becker price theory. So basically, it's as if you supplied all the time you have, cap T , hours and days times the wage. You have this non-labor income. The community as a whole has this. And then, you have consumption and Mr. 2 at least buys back leisure.

You know, well, it's obvious, you could just take time minus leisure. That would be labor supply and put it on the right hand side. But it's kind of useful to talk about leisure as a good, because that's what enters the utility function. And we have some standard sort of consumption sharing rules, division rules.

Who's going to bear the risk? Well, it's going to turn out that as we've done so far, this non-labor income risk is going to be shared-- taken away entirely by household 1. But things are much more complicated when it comes to the wage. The wage moves around, and it changes income.

But the wage also represents a rate of substitution between consumption and leisure for the group as a whole, even though person 2 has to supply it. You know, the wage goes way up, there's kind of a change to opportunity set for the group. And you can imagine things should move in order to be efficient to take that into account.

So we solved the problem in two steps. And that's also going to make it familiar in the sense of it integrates with what we've been talking about so far. First, ex post inefficiency, suppose-- ex post efficiency-- this is like the second welfare theorem. But I'll tell you what I mean. Any Pareto optimal allocation can be supported among households with a suitable lump sum taxes and subsidies. Remember that theorem?

OK, so what does it mean for us? You can kind of decentralize the problem of the two people, as if one or later both can supply labor on their own. But this lump sum tax or transfer is like an internal redistribution within the household. So for any target Pareto optimal allocation, you can find a transfer and then have each act in isolation, maximizing utility.

And then step 2 is to take sort of those indirect utility functions, take a step back in time and maximize the ex ante expected utility. And that's going to look like a standard risk sharing problem. And in this case, for this baby example, only member 2 is supplying labor. Remember it was Cobb-Douglas, and the power was basically 2. So basically here it's easier. You take $1/2$ of your full income. And that's your consumption. So those are those easy kind of consumption expenditure rules.

This is actually the same. You put the wage over here. The value of leisure is again $1/2$ of your full income. So those are with that functional form standard leisure consumption for household 2, plug it back into the underlying utility function, raise it to a power, and you may remember working this out at some point in some other course. You get a standard indirect utility function with all this stuff going on in terms of these coefficients.

Well, it's kind of obvious that it had to end up this way. Note that the wage is in here a couple of times, partly because the full income and partly because this guy's optimizing-- as the wage changed, say, the allocation, the maximizing allocation would change. You've got an indifference curve, and you're going to change the slope of the budget line.

So now we take that step back and consider the overall maximum. But instead of having utility of consumption directly of the two members, we can use the indirect functions. I mean, if you believed me, then the indirect function is without any restrictions. It's a natural object that represents efficient ex post optimization. Here's that row. Did I-- I should have emphasized it. Row is the amount that say 2 is getting from 1. Those are the internal transfers. And household 2 optimizes for a given row.

So the ex ante allocation is set the marginal utilities basically equal to each other, or equal to a constant. This constant would be the ratio of Pareto weights. I'm not saying that two members have to be treated equally within the household. It could be very asymmetric. We usually have λ_1 over λ_2 , or μ_1 over μ_2 , or some ratio of those ex ante given Pareto weights. That's all embedded up here in the constant.

Of course, convenient part is the indirect utility of household 1 is just linear, because that person had a utility that depended only on consumption. And it was linear in consumption. So we can now, with this constant, if we carry it around and multiply, you know, keep track that parameters get added or multiplied to it, that we can get a closed form solution from this for the transfer rule. And we already know what consumption and leisure are for person 2 as a function of the transfer.

So we solve that. And I don't know, it's just algebra. What do you want to see down here? Well, person 1's consumption is moving around one to one with non-labor income. So as anticipated, member 1 is absorbing all of that risk and member 2 is absorbing none of it. But the wage is affecting consumption and more obviously labor supply of person 2. And that wage is actually affecting person 2 as well. Here it is down here in-- person 1, sorry, in terms of person 1's consumption.

And it's doing it in a particular way basically. So this paper is called sharing wage risk. And both these members will share the wage risk, even though only person 2 can work more or less hours at that wage.

So in some sense, leisure pooling is a bit different from consumption pooling in the way that they're going to want to adjust their hours is going to depend on the outside opportunities. So that's the summary. And then we'll generalize this case. It's not general, because person 1 was risk neutral.

Here's the general case. There are S states of the world. There's a risk sharing group of households, as in a village or township or something. And every household H has individual members. It could be more than one. It's likely to be two or three. This is meant to be realistic. And we actually go to the data with these specifications. And we have hours and labor force participation for every member of the household.

So we normally only have consumption for the household. So we kind of throw up our hands a bit, but here, we have labor at the level of individuals. And we definitely want to keep track.

The consumption of household H is a sum of consumption of individual members. That's true in theory, even if we don't measure it individually in practice. Aggregate consumption in the village is a sum of consumption across the households. Nothing surprising there. Utilities of everyone stays strictly increasing and concave. And household H faces a vector of wage and non-labor income variation.

You know the S is a subscript on the wages of person 1, the wages of person 2, for how many people there are in this household. And again, here although it really doesn't matter because we're going to added up anyway, you could say different people in the household have different occupations. Some are helping on the farm. Other people run the business exclusively. You could talk about the non-wage income.

If there's perfect markets, which we're going to assume, then the profits of the firm are as if they had supply their own labor to the market and hire other labor to run the business, which is problematic. But, you know, then you get the usual sort of separation. And so these are really like profit numbers.

But this is the aggregate of overall individuals in the household. This is what we've been calling household profits. And finally, the resources available are household profits plus gifts and other things. These τ s are not yet determined. They're given for this household H in state S . In a minute they're going to get determined endogenously from the rest of the village. So, again, we're going to kind of solve individual problems, household problems, village problems, and sort of build upwards.

So what does the household do? It maximizes these new weighted sums of utilities of its individual members, summing over states. And actually, we could sum over dates. But you already know that doesn't matter. There's nothing you know particularly inter-temporal about this-- well, these sub-problems apply even when the household as a whole can borrow and lend with the outside, et cetera. And just simplify by getting rid of that for now.

Common probability over states S . Utilities are allowed to depend on individual i . We're going to try to retain as much heterogeneity as possible, although it's sometimes problematic how much we can allow in the data. And here's the sort of collective resource constraint for the household H . It's the same as before. We're just summing over all the individual members. This is the sum consumption, sum of the valuation of leisure, equal to full income, including any taxes or subsidies the household is getting from the rest of the village.

These Pareto weights μ are kind of taken as given. That's not to say there aren't forces that determine them-- you know, the marriage market, one person has a lot more income than the other, all this bargaining. But any bargaining that takes place here is entirely ex ante. And it is assumed to be fixed and unalterable subsequently as states of the world are realized.

You know, we haven't talked about it in this class. I think you've probably seen it in other classes, this whole literature on the unitary household and, you know, ex post bargaining, when opportunities of the spouse increase and so on. They may take more goods. They may better dictate what household is consuming. That's all suppressed here. It's buried in these fixed Pareto weights. The idea is to see how far we can get even assuming that.

OK, well, then we have the Marshallian demand. I only say that with emphasis because we're going to go Hicks and Frisch in a minute. But this is the regular demand for leisure that solves the individual's optimization problem. This is an individual i in household H . This, apart from the more cumbersome notation, is exactly what we went through before. This is some transfer that individual i is getting in state S from the rest-- giving or getting from the rest of the members of the household. This is the valuation of the time endowment for individual i . And in the left-hand side is the sum of expenditures on consumption and leisure. And that after maximization gives you the indirect utility function, which clearly depends on the wage and on this transfer.

So for the household to have an objective function, we can call it sort of the Pareto μ weighted sum of the value functions of the individuals. Again, we went through this in the baby example at the beginning. And then the problem of the village-- note that now we summed over H -- is to maximize this household by household expected utility, and now making it expected because it's conditioned on state S . And here's the constraint that these transfers stay within the village have to add up to zero if the village is the risk sharing group you want to be considering.

One first order condition, you know, basically says that the marginal utility of the household respective transfers in state S is just pre-multiplied by the probability and then this λ_{HS} -- so this λ_{HS} is like the shadow price on the budget constrain of the household at state S . It reflects how tight or loose the optimization problem is, just like in standard price theory. It's just the Lagrange multiplier on the budget constraint.

It's also going to equal something else. But so far I haven't talked about it in the notation. It's going to equal something else, because these sum of the transfers are zero. So there's like a collective resource constraint that isn't written out here explicitly. And that thing is going to pick up a Lagrange multiplier. And that Lagrange multiplier is going to be common over all the households, just like it was for the consumption problem.

OK, so here's an example, Cobb-Douglas coefficients are depending on i and H raised to a power. This obvious generalization of the example. We have a little corner to worry about here. You know, the way to see this is putting w over here, this is the expenditures on leisure. And expenditures on leisure ought to be equal to-- well, these add up to 1, basically-- I almost think that's a typo. Should be $1 - \alpha$ -- the expenditure share on leisure.

But what happens if this amount, given this full income, this leisure that individual i is supposed to be consuming, exceeds his time endowment? Then you've hit a binding corner. So hence the min, which is this cannot exceed this. When it's less, it's no problem. You just when hours are less, you do those hours. When hours are more, you do the max.

Now, let me just say this is very natural in this problem when the time endowment is fixed. But you could also imagine other indivisibilities. You could imagine that you could only supply discrete amount of hours. Or something like in the US, typically you have 8-hour days and 40-hour weeks, and that's it. And you can't vary your hours. So you can load other restrictions in here.

And in the data, there is enormous variability in non-participation. Month by month, households will sit out for substantial periods of times and then re-enter the labor force. So something is making that happen. And actually, it's the other corner. It's more like zeros, not working up to the max. Yes?

AUDIENCE: So it only gets a corner solution if that household's Pareto bases is too large for that agent?

ROBERT
TOWNSEND: That's one of the factors. No, too small. So-- oh, too large, sorry, leisure. And I just made a mistake, absolutely, yeah. So when the Pareto weight is high and leisure is valued, then they want to do maximum leisure. They set leisure equal to the time endowment and work hours are zero.

So that's one thing, but it depends on the wage. And it's actually going to depend on how binding the household's resource constraint is. So there's two other factors. But you're right about the Pareto weight.

So we can solve out for the value functions in these work and no work branches. I'm not sure that is very enlightening. There is going to be a critical wage, as I was just anticipating, where that constraint is binding or not if the wages are higher than this critical number, the individual would be working. And if it's lower than that critical number, they're not going to be working. So there is a threshold.

And it turns out it's smooth as the wage-- there's no jumps in this problem, at least. You can see it in the value functions. You can see it in the way the wage moves. But again, I don't think it's worth getting bogged down in the algebra to make that point. In other words, hours will kindly smoothly stay at zero until the wage goes above this threshold. And then hours will start to pick up in a regular way.

So here's what I just mentioned to Yeng, which is that households are less likely to participate when the wage is low, when the household is doing well, and when the Pareto weight is large. Or conversely, you'll see greater participation, and condition on participation, greater hours when the wage is high, when the household is doing badly, meaning that the shadow price on the budget constraint is high, and when this person's Pareto weight is small. You know, make the slaves work.

This looks awful. It supposed to be a good thing that we are able to analytically solve. So let me highlight something. There's two equations on the slide. There's a reservation wage and an hour's equation. OK?

Actually look down here first. The reservation wage for individual i in the household h has a constant term and something that depends on the shadow price of the budget. Now, when we did the consumption risk sharing thing, we talked about the shadow price of consumption. I'm kind of like tongue tied. I don't quite know what to call it. It's more than consumption, because we have leisure in this problem. But we did have a budget constraint. So it's the shadow price on the budget constraint. And there's no H on this guy, because this thing represents that all the individuals and all the households in the village are pooling all this risk. So there's like a common community level resource constraint. That's the content of the theory.

And here's hours actually worked when working. And again, it has a constant term. It has this same shadow price object, DS , in it, and also on the wage. Now, the wage just reflects the fact that, as we said before, the higher the wage, the more you're likely to make someone work.

And so now we can maybe get the courage to go back up and look at what all these objects are. Look at all the beautiful restrictions that theory is placing on these regression coefficients. First of all, why all the i 's and h 's? Well, we allowed a lot of heterogeneity in those Cobb-Douglas coefficients and the curvatures for each of the households.

Let's spot some familiar-- this is all preference stuff. This is cool for me. This is the log of the probability weighted shadow price in the whole village. And what's this? This is the risk aversion. So you remember that stuff about risk tolerance? So, you know, the greater γ , the less risk tolerant they are-- sorry, going the wrong way. This γ is the curvature in the utility function. So-- oh, there's a negative sign there, right. So basically what it's saying is the more risk averse that an individual is, the less you should see that its hours are moving around with this aggregate state. So that's the leisure version of the consumption equations we had before.

I mean here's Pareto weights in here, for example, in the critical wage threshold. This is the Pareto weight of the household. This is the Pareto weight of the individual. And the higher this thing is, the higher that wage is going to be, because that's basically the way you take your leisure. Yes?

AUDIENCE: So before we were talking about it's great to see all these restrictions because of being familiar with the data and doing regression. But we will never have data on a Pareto weight, right? That's not a real thing?

ROBERT That's true. But--

TOWNSEND:

AUDIENCE: If that's correlated with other stuff that's going to start moving around the other betas that we're estimating on the stuff we can see, right?

ROBERT Yeah, but basically-- that's right. And that's what this is basically. So you start looking at what the regression coefficients are. For example, here's this one thing, the risk aversion of individual i and household h , and there's an F on there. So where was the F thing? The F thing was what I was highlighting right here.

TOWNSEND:

So now this is a fixed effect. We don't see it. But you can put a time bearing fixed effect for the whole village. And then estimate this regression coefficient. And the theory is telling you that's 1 over the coefficient of relative risk aversion. And so now you can work backwards, you know--

AUDIENCE: Wouldn't that be correlated with the Pareto weight on me?

ROBERT Mm, hmm. No. No, we didn't see that in the--

TOWNSEND:

AUDIENCE: Oh, no, sorry, what I meant is if I don't have the Pareto weight-- so that'll be like fixed effect, right?

ROBERT The Pareto weights is like a fix-- this is a time bearing fixed effect. And the Pareto weights are going to show up--

TOWNSEND: where are they? Over here. Here's one.

AUDIENCE: Can you have some--

ROBERT This is the households' Pareto weight. This is the individual's Pareto weight. Granted there's a little bit of these

TOWNSEND: alphas to worry about. So we'll have to grab them. But part of this fixed effect, this is not varying with time or states or anything else.

AUDIENCE: So you can have like an individual effect is what you're saying.

ROBERT Yeah. So that's sort of-- the algebra is if you had all these regression coefficients and so on, then you basically

TOWNSEND: can identify almost all these parameters. The thing that you can actually get is the absolute value of the risk tolerance. And we actually saw that earlier. It's the ratio of the risk tolerance of household i relative to other members that you're actually able to identify. You can see who's more risk averse than others, but not-- OK, and actually the paper itself goes through-- it's long. I'm afraid it's quite tedious-- goes through, you know, this parametric example. It actually goes through non-parametric stuff. Pierre-Andre has figured out with his co-authors how to identify tons of stuff, just based on labor supply, even though there's all this underlying consumption stuff going on. You use data from Britain. OK. So yep?

AUDIENCE: So you starting out with the multiplication that consumption will have to equal to waste.

ROBERT Yeah.

TOWNSEND:

AUDIENCE: But your paper in 1994--

ROBERT Yeah.

TOWNSEND:

AUDIENCE: So how do you reconcile with this?

ROBERT Well, I mean, the first answer was, you know, I assumed the problem away. Actually, not quite. If you look, you'll

TOWNSEND: see that I had put measures of hours in the consumption equations. I was aware of this even then.

But it certainly would be easy to make a mistake. And people are often critical, like puzzled, how can you say income shouldn't matter? And here, we are distinguishing the income that's coming from labor supply versus the income that's non-labor. And it matters. The theory places restrictions on how. So in some sense, consumption does move with income. If the wage is moving income, then it's moving consumption. And it's moving consumption of everyone.

AUDIENCE: Not quite generally.

ROBERT
TOWNSEND: No, it's just from the first order conditions of the optimum problem, the baby example that I gave at the very beginning.

AUDIENCE: So why are we testing a social planning problem instead of competitive equilibrium. Is that because Thailand economy is centrally--

ROBERT
TOWNSEND: There is no social planning.

AUDIENCE: Isn't just to see how close we get to the benchmark?

ROBERT
TOWNSEND: Yeah. First of all, it's not a particular benchmark, because those fixed effects-- we want to see where we are in the whole frontier, right?

AUDIENCE: I guess there's a menu of benchmarks.

ROBERT
TOWNSEND: Yes. So the idea is somehow like sort of a closed thing, that you wouldn't leave, quote, "surplus on the table." That there must be markets or institutions or other arrangements that lead them toward efficient outcomes. But, yeah, it's not like-- it's interesting that you're against the social planner. It's-- and I often slip, and you caught me the other day, saying, you know, social planner this and that when I did the capital asset model. That's just a metaphor for-- it's as if the community as a whole is allocating resources so as to not lose any surplus.

Now, you know, exactly what the institution-- so I should tell you about what we find. First of all, in these villages, there's something called labor exchange arrangements. And it's partly with the rice thing. You know, the low land gets flooded early. They all go down and help that guy farm his land. And then as the water rises, they flood the other plots. But they don't charge each other for the labor. It looks like free labor exchange. And it doesn't necessarily balance out either.

Now, I'm not saying in any given economy you'd be lucky enough to observe something like that. But in these villages you actually do have a real institutional counterpart that might explain how they managed to do so well. Yes?

AUDIENCE: So I'm not sure it would matter for this model. But the story you told is that there's several risk sharing arrangements for a different domain. That one was labor opportunities. Maybe there's a separate one for consumption. Would it matter-- from what perspective would it matter if there were one unified arrangement?

ROBERT

No, it wouldn't matter. But, you know, it's nice to be able to say something about what the mechanism is. You

TOWNSEND:

know, for some people just looking at outcomes isn't totally convincing. I must say when I started this, I thought it was a virtue to say we don't care how it's done, I'm not going to look separately at savings accounts, selling bullocks, borrowing, running down your grain stores, let's just look at the total. But still people want to know how it could possibly happen.

So what's going on in these data? Well, I spare you all the details. But we put non-labor income on the right-hand side of the participation equations and the labor supply equations. And the coefficient on it is small and quite negligible. I mean, it's not zero. So we reject.

But, you know, if you come at this-- and we'll see this in Seema's paper-- if you come at this from the traditional point of view in development, it would be like labor supply is your safety net. If you're getting your income is really low from crops, then migrate out, work for wages. And this says, well, not so fast. It might be done at a community level.

And the test of it would be whether how well you're doing on your own land has any influence on your participation and your labor supply. If they're pooling risk like this, it shouldn't once you control for, say, village level fixed effects. And, in fact, you reject. But the coefficients are tiny. They really are sharing a lot of this risk, not all of it, but a lot of it.

OK, one last thing just to tie some of this literature to-- oh, god. So this thing, this is like a famous parameter, because it's called the Frisch elasticity. And that brings to the next part of this lecture, which is this tension between the small estimates, micro level estimates, of labor elasticities compared to what we seem to need in a model to get aggregate hours and participation and move around so much.

Now, this literature spans micro and macro and development. From the macro side, Chetty reviewed all of it. But, you know, Sargent, for example, but way back to Rogerson and Hansen talk about a way I'm going to show you an indivisibility that makes the overall aggregate elasticities a lot higher than you would get if you looked at those individuals in the data.

So again, just to anticipate, this is a bit like zooming out and zooming in. You zoom in to the individuals. If you have labor supply and so on, you can estimate their micro level elasticities. And there's just tons of literature, which I suppressed from the slides, because it will overwhelm you of the interest in estimating those elasticities. But you could say all these individuals reside in a village or district. Let's look at what happens when there's a rainfall shock. Let's look at aggregate-- how much the wage is moving as a function of the rainfall shock and try to estimate from there how inelastic labor supply is. And it's natural that aggregate up over individuals to get an aggregate labor supply if you're talking about how district wages move around with rainfall.

Well, that's Seema's paper at the end. She's seemingly not aware, and she didn't make a mistake. But it would be quite easy to make a mistake if you were quantitatively trying to reconcile the data that you use in the aggregate with the micro level elasticities. And so that's, as it turns out, been the focus of this of this macro literature, so many elasticities.

And partly because the different literatures developed different jargon. There is the Hicksian elasticities and the Frisch elasticities. A Hicksian elasticity is basically compensated. So when you have a wage change, you don't just look at total hours. You push people back to the original indifference curve and slide along it. That's Hicks.

Frisch has to do with variation that you would see if you held the marginal utility of wealth constant. Why does that come up for us? Oh, risk sharing. The marginal utility of wealth is basically at each date and state constant over all the individuals. It's a very natural object and related to the risk sharing literature.

And then macro, they call-- just to add Frisch elasticity. But it's of aggregate hours. And every one of these things has extensive and intensive. And that's because there's a big variation in participation, not just hours in any data that people have looked at. I'm not sure Chetty did this. I'm not so sure how helpful it is.

But the idea is the Hicks effect is like shifting a wage. And you're going to shift it-- if you're going to distinguish age, you shift it over the entire age profile. That's the way he thinks about it. Whereas he thinks about the Frisch elasticity as just one little instant. I have trouble relating, but anyway this is-- and then there's this table.

These are micro estimates of the intensive means conditioned on working, how much your hours are varying with your wage. These are, quote unquote, "relatively low" numbers, 0.3, 0.5. The extensive margin is in the micro data also not big, 0.1, 0.2.

What this literature is about is aggregate hours. Aggregate hours have higher elasticities than any of those numbers. And in particular this thing, this macro level, Frisch elasticity is huge. So in order to generate this number that we see in the data, and keep the micro parameters that you're going to fit in, to be consistent with that you have to generate a huge elasticity on the extensive margin somehow. So this is about trying to reconcile micro with macro.

AUDIENCE: Is that bottom line, they're not estimates from the data?

ROBERT Which one?

TOWNSEND:

AUDIENCE: The bottom line.

ROBERT Yeah, they're not. Although this is like that. So let's respect the data on that dimension. And somehow create a very big number here somehow-- and I'll tell you how in a minute-- in order to get what we see in the data over here.

TOWNSEND:

Now, in the original slides-- I'll give them to you-- there's like six slides that now come you know with all these other who did this and which author and so on. And I just felt that was going to be bewildering and certainly time consuming. So I'm skipping it.

But here's what I already went through. And I guess I said this too. So the punch line is going to be if we have a non-convexity in participation, like you either have to work full-time or not at all, is an extreme. And then we go through the mechanics of how people would behave in that economy. We can generate a sort of pseudo representative consumer. So that it looks as if the aggregate are generated as a solution to that representative household. And the trick is that representative household is going to end up with a Frisch elasticity in the pseudo function, which will allow huge variability in aggregate hours. It responds to wages. Yep.

AUDIENCE: This doesn't match the extensive margin micro data.

ROBERT Because it doesn't have to.

TOWNSEND:

AUDIENCE: Why?

ROBERT TOWNSEND: Because you can populate the underlying individuals with actual observed micro estimates and then say there is this inelasticity. I'll show you. I'll show you.

OK, so we have labor, capital, and output. Let's just do a single period. Imagine for a minute it's not stochastic either. We have an aggregate production function, as a function of aggregate capital and aggregate hours. Why are we doing it? This is basically just going to allow us to close the model instead of having it open partial equilibrium. So we're going to get the wage and interest rate off of this thing. It's convenient.

There are a continuum of individuals indexed by i . Each individual has one unit of time, one unit of capital. Time is indivisible. And you got to choose, you're either going to do leisure or you're going to do work. So that's an extreme version of the indivisibility.

They care about consumption and leisure, or the disutility of work. But this work level is either 0 or 1 by assumption. So let's just normalize and say v of 0 is 0. But v of 1 is whatever it is. Let's call it m . It's just a real number, because you're never going to see at the individual level any other levels of hours. So remember m is the disutility of working 1-- putting all your time into work.

What's the consumption set in this economy? Consumption is non-negative. You can't supply more capital than what you own. But you have this non-convexity in the labor decision. It's 0, 1. This is not an interval 0 to 1. It's either 0 or 1. It's the set consisting of just the two objects. So it's a non-convex consumption set. Linear combinations of 0 and 1 are not allowed.

OK, so what would be a competitive equilibrium? Households max, firms max, markets clear. Here's the household max problem. For any individual i , solve for the consumption decision about participation and the capital respecting the budget constraint. Namely, you can work or not-- so this is either 0 or something positive-- at a given wage. And then you get the rental rate on your capital stock. And that's that. Firms maximize profits, paying factors of production. Clearly for a given f , this is going to generate marginal products of capital and labor in the whole economy.

And if we just solve the problem this way, we could get an equilibrium. But it might not be optimal. We can make people better off by sort of randomizing who gets to work. How are we going to do that? We're going to basically create this probability object ϕ .

So here would be the expected utility given a ϕ . The realization with probability ϕ , you're going to be the worker in this economy. With probability $1 - \phi$, you're just not going to work. And at the moment, you can get different consumption for the two different branches. Yep?

AUDIENCE: I sort of understand this is getting at which the market might be flexible. There's not much part-time. There's not much intensive margin. So people turn to things like household and they specialize. But how do you interpret the lottery as-- like we imposed lumpiness of working. And then we use the lottery to break it. Why not have some other way of having other types of labor worked?

ROBERT TOWNSEND: Um-- well, I'm about to show you on the next slide the decision problem for an insurance company. And the insurance companies like pooling over all the individuals. And it's going to maximize, and there's going to be free entry. And that's going to drive the prices of certain things to be equal to their actuarial values.

So it's as if people were buying and selling insurance. And they get to choose-- they're not forced as if by a planner to work or not work depending on how the roulette wheel turns. They actually get to choose voluntarily whether to decide to work full time, decide to work zero or be on call. They'd be on call, they're going to get a wage that's proportional to the probability, because that's the pooling aspect of it.

AUDIENCE: Like [INAUDIBLE]

ROBERT Yeah, so far. Yeah.

TOWNSEND:

AUDIENCE: I don't know that the lottery is so great, because it could be we have like some of us turn up for the job interviews, and the people doing the job interviews like just-- you know, they're picking based on some stuff. Maybe it's just as good as random, right? Like so--

AUDIENCE: That's crazy.

AUDIENCE: No, it's not.

ROBERT Disconcerting maybe.

TOWNSEND:

AUDIENCE: Suppose there are like 100 unskilled manual laborers turned up for a job that requires 30 laborers. I cannot observe past project of these laborers. Like I can maybe observe cleanliness and they're ability to speak. Like that's about it. So what am I supposed to do.

ROBERT So that would be an institutional arrangement, where the markets seem not to clear, because not everyone gets the allocation, not everyone supplies labor. You know, they're lined up on the street corner, and the truck comes by, and there are only so many jobs.

AUDIENCE: But what if we were indifferent, then it would still be--

ROBERT But there are other ways to do it. I mean you could have sunspot. You can even index off of observable things like rainfall. And then it would look like a state contingent contract, where the labor supply is a function of the rainfall. But that doesn't mean everyone's supplying labor. They have some arrangement. If it's sunny, I'm working. If it's raining, you're working.

And it would be very hard to say, unless you were very clear headed, whether it's this sort of extraneous risk or some intrinsic risk, some underlying state that they were trying to insure against. This last interpretation is Arrow's interpretation of the lottery. And he's entitled.

OK, so we're going to have these prices, rental rates and wage rates, normalizing the consumption good. OK, here's the insurance company thing. Basically, if you're working, you pay a premium ex post. And if you're not working, you get this indemnity. Well, I mean intuitively if you're working, you have more income, you kind contribute some back to the pool.

But if an insurance company we're offering this, they would potentially make profits depending on the fraction of the people from whom they get the premium and the residual fraction to whom they pay the indemnity. If you have free entry into insurance, then this is going to drive this equation to zero. OK? So that's going to tie the premia and the indemnity to the probability or fraction of people working.

So now, we're going to decentralize it. We're going to say look it's as if the household were maximizing expected utility, but as I said, voluntarily deciding whether to work for sure or not work at all or something in between. And then they're going to have this resource constraint. But I've now substituted in-- the x's entered into consumption. But the x's we're constrained by that zero profit condition. So we've substituted in the zero profit condition back into the household budget constraint and gotten this equation.

Capital is equal to 1. So this thing's just going to go to 1 no matter what. You're going to get a wage. But you're basically going to be paid the expected value of your pre-committed labor participation.

I agree to be on call. A hospital may call me. I'm going to go to work if I get the call. But my salary is based on whether I'm on call all the time or not at all. And I get to choose that. Somehow that's the flexibility in the contract.

And this is the decentralized version with the household optimization. Firms maximize profits and markets clear. That's where we're going to get the wage, the overall wage and the interest rate.

So if you looked at the first order condition to the consumer problem, you would see some obvious things happening. You're going to have a theta price, shadow price, on the budget constraint that I just showed you. And these probabilities are entering symmetrically in utility and in consumption. So when you take a derivative, they are kind of like cancel out.

And so you get this. But, look, here, they're going to cancel on both sides. So theta is common. So consumption has to be common.

So you're going to get full insurance on the consumption side. There's no reason to let consumption vary. Remember, this was a separable utility function. It was the utility of consumption plus-- so why screw up the marginal utility of consumption? You can get full smoothing out of on the consumption side. So consumption is equal to each other, whether you work or not. But that theme may not be at the boundary. It could be interior. So substituting a common consumption, remembering disutility of work is just m, you have this new optimization problem. And this is equivalent with the other one.

Now, where's the aggregate elasticity? What's the disutility of work?

AUDIENCE: Phi, right?

ROBERT
TOWNSEND: It's linear. It doesn't get any more easy to substitute than that. See here it looks like you're choosing phi. But the coefficient in front of it is just a number, m. So phi is-- you know, why is phi look like hours? Phi is the fraction of time you're working. And if you work, you're working 1. So total hours is like the fraction of the population working.

That varies continuously-- this sort of pseudo representative consumer thinks that it can choose any fraction it wants, even though the underlying population cannot. This is a stand in household. It's not a real household. But the aggregates that are predicted from this will correspond to the underlying problem. So that was the insight that Rogerson had.

And actually, I'm pretty sure it's also 1 to 1 with the right welfare. You can use that aggregated utility as a representation of the underlying welfare. Why is that? Well, even though you have all that diversity, households are really alike in terms of their utility functions. So expected utility is the right metric. And that's what that thing is. It just looks funny when you write it down. OK.

And then quickly, here's Hansen's version of it with business cycles, going to the actual aggregate data. As I said, total hours worked is hours per person of those participating times the number participating. And then you take logs of it, basically.

And then you want to look at a variance. You say the variance of a log of aggregate hours has three components. It's the variance of individual hours, the variance of participation on the extensive margin and this covariance term. And in the US data, those are the numbers consistent with that table that I showed you before. So a lot of the variability in hours, say over a business cycle or something, has to do with varying participation among individual members.

OK, so this is sort of familiar ingredients. You'd have an aggregate production function over hours in capital. You have some Markov process on TFP. You've seen this in the first part of this class as well. Capital depreciates. You can add to it with investment. And you've got this utility function. Again, this is like a Cobb-Douglas version, well, logs off of a Cobb-Douglas version.

And the overall problem is to maximize discounted expected utility. Except we introduce this indivisibility. So you can either work at h_0 hours with some probability or 0. Just like Rogerson with slightly different notation, you'd get α , say the probability of working or not working. You know, the way they've normalized it, the log of 1 is 0. This thing is going to drop out. And you'll get a utility function over consumption and hours.

But remember, again, just to reiterate, h_0 is the fixed number of hours you have to work if you work at all. And α is the fraction of the population that's being assigned to work. α , this coefficient in front of hours, is the thing that's getting determined in equilibrium.

So after you make all those substitutions, just like we did before with Rogerson's paper, we end up with this pseudo utility for the representative consumer. And, yeah, again, it's linear in hours. So depending on the wage relative to b , it would look as if there's just enormous willingness to supply labor, or elasticity, as you go from low wages to high wages, you know, an enormous response. It's just a review of what Rogerson did, but it's in the context of this business cycle model.

So solve the first order conditions. Compute the steady state. Compute approximation to the steady state. Solve for the law of motion of the endogenous variables. And compute the moments.

Now, I don't want to belabor this, because it would just take us too much time. But we did models at the beginning of class that featured occupation choice and development and so on. This is like that in the sense that we have an attempt to match data. In this case, in this literature, it happens to be business cycle data. But you're going to see Seema's version of it in a moment.

So the cool thing is you've got this machine. You've got this model. And if you're willing to take a stand on parameters that you observe in data or calibrate other parameters, you can generate time series from the model. And you can compare the time series to what you see in the data. And this is the algorithm for solving the model.

Again, I'm not going to belabor this. You know, they pick capital share and depreciation rates and discount rates and so on and so forth and generate the data. Now, they're two economies here. This one with divisible labor, this one that we featured with indivisible labor. And you might want to look at the variability of output consumption and hours.

This is the variability of ours in the data. This is the variability of hours without the indivisibility. This is the variability of hours with it. Now, they didn't get it all the way up to 1.66. But they're certainly dominating 0.70. So that, you know, as I keep saying, these economies with certain kinds of indivisibilities are able to increase the responsiveness of aggregate hours to wages, more consistent with what we see. Yeah, Matt.

AUDIENCE: Is this economy with individual labor, is it actually-- like one of your student simulation-- is it identical to the RBC model except the utility function?

ROBERT Yeah. Well, except for the indivisibility, yeah.

TOWNSEND:

AUDIENCE: When you have the utility function with indivisibility, when you have like the reduced form, because don't-- I mean, in fact, again you've got ϕ or the proportion of time work--

ROBERT That's what's going on. See, there are two economies here and kind of two ways of doing business. You could go to the micro data, look at individual labor supply, try to populate the economy with some fresh elasticity that you observe in the data and then go through all the steps. Or you could say, no, no, no, no, you either work or you don't work, get that pseudo household. And it's going to generate its own elasticity and redo the simulations. And that's what's getting compared here. Yep.

AUDIENCE: Compared to the response of working hours is totally driven by the extensive margin.

ROBERT Yes.

TOWNSEND:

AUDIENCE: If you also want to match the elasticity of this extensive margin, maybe--

ROBERT Yes, I'll show you another paper where-- so this is too extreme. But it shows how to get-- OK, so this is Kim. I'll take you through this really quickly. But it does link up to the incomplete market literature that we talked about last time. And the point is you can generate a higher aggregate elasticity than what you see in the data without assuming all this stuff with complete markets and the lotteries.

So in other words, there's something crucial about the aggregation, which is more limited here. We're not going to get as high a number. But we are going to get a higher aggregate elasticity than what you would think if you just looked, as Yeng was saying really, if you just looked at the micro data.

This is a nice paper-- I wasn't aware of it till a few months ago actually-- because they're really into the details of the cross-sectional earnings and wealth distribution. They're really looking at the micro data. Whereas those macro guys, at most, are borrowing a few parameters.

You've got males and females, like a two-member household, maximizing expected discounted utility. This is again this Cobb-Douglas type specification. This is at the aggregate household level for some reason already aggregated with these weights. So it's male hours and female hours.

Gamma is the sort of degree of inter-temporal sub-elasticity that we care so much about. And now, they're going to be serious micro guys, and they're going to actually look in the data, at male and female wages and impute productivity. So these α things-- there's going to be one for males and one for females-- take hours and multiply by productivities. And these productivities can move around over time.

So it's not just a firm having TFP that's moving around. It's households having labor productivity that's moving around. And the household is going to have this sort of standard incomplete markets budget constraints. So this looks a lot like what we had last time sort of. They're going to have assets in the previous period and interest on that. They can save to take into the following period. And the household, as a whole, is going to have wage earnings as a function of hours. Now, we have labor supply. We didn't have that last Tuesday.

And there's, say, a bound on assets. This could be zero, no borrowing. It could be a negative number, a limit on borrowing. Just like the standard incomplete markets lecture from last time, we give them an aggregate Cobb-Douglas technology. And we solve this forward-looking value function.

So it's a bit bewildering. But they're both working. This e is not w_e . It's like w_{ee} . No, ee means employed employed, employed not employed, et cetera, depending on the female and male members. And this is the case where currently they're both employed. Solving for assets, you know, anticipating that they'll decide on their work status next time.

Now, again, you can see this extensive margin here is playing a big role, deciding member by member whether it will be working or not. And at least in the '50s, this was a big issue, increasing female labor force participation. Things have settled down now. It's much more symmetric. But the wife did not always stay home. Sometimes they participated, depending on the household situation.

So family labor supply, and what do I want to say here? So not much. It looks more complicated than it is.

You have this all for a steady state, which is complicated. Someone asked me-- who was it? Can't remember-- about adding state variables. So here, you know, we have a lot of state variables. We've got male productivity. We've got female productivity. We have the distribution of wealth. All these things have distributions in the population. And we have the sort of shock in the aggregate production function.

So they're going to do what those macro guys do, which is just look in the steady state. They have to compute it. And then they're going to look at very-- there's no time subscripts on the wages and interest rates, because you're in a steady state, where all this distribution stuff has settled down. You have a lot of churning from one status to another within the steady state. So individuals and households are experiencing different things over time. But the economy wide wages and interest rates have settled down to constants.

And I'll skip-- they actually start estimating parameters with the micro data. And, you know, their goal here is to get this aggregate elasticity high, higher than you would see as they estimate from the micro data. So, Yeng, that's a little better example of actually looking at the underlying micro data and still using this sort of macro aggregation. Yep?

AUDIENCE: Still saying that from micro data the elasticity, so it's made to be smaller.

ROBERT Yeah. Yeah.

TOWNSEND:

AUDIENCE: Like even though you have model, it just--

ROBERT That's right. So why is it still happening? Because, again, there's this indivisibility going on about whether or not
TOWNSEND: to work.

AUDIENCE: So if you incorporate indivisibility, then-- because I thought that point of Cheng is that if you look at the extensive margin, that's still like very small compared to what the macro model would be in individual labor.

ROBERT Yeah, again, at the individual level, you'll back out certain low elasticities. But the point of this whole literature is
TOWNSEND: when you aggregate up and look at how aggregated hours are moving around, you can get bigger numbers. It's not inconsistent.

Which brings me very briefly, although fortunately, I've said a few words about this, Seema's paper looks at productivity shocks in agriculture And, looks at what happens to wages. Now, you know, think about a supply curve and a varying demand curve. So it rains a lot or a little, that kind of affects productivity. That's moving the demand for labor. If the supply curve is inelastic, you're going to get a high variability in the wage.

Her point is she cares. The workers get hurt a lot when there are droughts and it doesn't rain because the wage drops. Landowners love it. Their output is lower, but they have less of a wage bill. And then she couples that with assumptions about how well or how poorly the banking system is.

So I'll just tell you in words, she has a 2-period problem, which she thinks is enough, rightly, to get what she wants. She has people forward looking. So they're solving 2-period problems. And then you decide whether or not to save or to borrow based on your current situation. She's saying that if they can borrow and save in the banking system is good, they're going to be less needing to work in a bad agricultural year, because they can borrow instead of working harder. So this is an incomplete markets model with limitations and regional variation in the banking system that is focused on how much wages move around with productivity agricultural productivity shocks.

Now, her paper is qualitative in the sense that she shows that there is with less what she called wage elasticity-- there's less responsiveness of wage to rainfall shocks or instrumented productivity shocks-- the better the banking system is. That's what she cares about. Now, interestingly, she doesn't get into the details of the numbers. How quantitatively big or small is this response? So she doesn't do anything wrong. It's actually a very nice paper.

But if she were seriously trying to estimate the elasticities, she would need to take into account whether there is this extensive margin participation. And the aggregate level elasticities could be lower than they might seem to be as in the two papers that we just covered. So again, nothing wrong. But, you know, when you're in practice and development work and you're thinking not just about individual responses, but sort of collective responses and whole regions, and so on then these sort of tricks that we get, lessons learned from the macro literature are very helpful in combining the macro and micro data.

So you can read the details of the slides. But I basically covered it. Thank you.