[SQUEAKING] [RUSTLING] [CLICKING]

BEN OLKEN: So just to remind you where we were, we were talking about the separation test. And we had gone through all the different-- all these different cases, testing for separation, or whatever. And the idea of testing for separation-- the idea was to say, if we have changes in household labor supply, does that affect the amount of labor demanded on the household farms?

And then, we talked about some of the challenges with some of the empirics in the Benjamin paper that you guys read for last time. So the idea of this LeFave and Thomas paper is basically it just literally takes the model, et cetera, off the shelf from the Benjamin paper, but redoes it using panel data. And so why-- so I don't-- did we discuss this at all at the end of last class? Yes, no? Yes, we did, a little bit but not a lot.

OK, so just to sum up then-- OK, so, Tara, so since you answered the question. So what was the idea of why was the panel data helping us?

AUDIENCE: [INAUDIBLE] effects?

BEN OLKEN: Sorry, Shawna.

AUDIENCE: It's OK. We change [INAUDIBLE] effects for individuals or farms.

BEN OLKEN: Exactly, and why is that helping? What is the empirical concern that that's helping?

AUDIENCE: That there are unobserved individual characteristics of households.

BEN OLKEN: Yeah, exactly, so there's unobservable characteristics. So maybe there were people who had really good land. Those people would have more kids. And they would also have-- we'd be sure they'd have more kids. They'd also have more labor demand. That's the unobserved characteristic that we'd be dealing with, right? So fixed effects basically allow you to deal with that because we can use variation over time.

So the idea is that, basically, we're going to-- so in a very-- if we basically ran this regression but there was unobserved land quality that was correlated with both labor demand and people in the household, with fixed effects, we're going to include that. It's equivalent of essentially demeaning things. And so that concern is no longer there.

So does this strike you as a more plausible empirical setting? Yes, no? What would you think about for this? Yes, no? Well, clearly, it's more plausible. Does it strike you as plausible?

So the ideas we are going to use changes-- so first of all, what's identifying this, to be clear, is changes over time in household composition within a given household. So we have to have a panel structure where household composition is changing. And so the question is, if we get more or fewer people in our household, is that associated with changes in labor demand?

So first of all, does some want to argue that that's not plausible? What would be a problem with that identification idea? Yeah, Whitney?

AUDIENCE: I think we haven't completely fixed the endogenous family size--

BEN OLKEN: So the same arm.

AUDIENCE: So if, for example, your household isn't doing well, you're like, my relative went out to go work. But now, I need more help or something. I'm going to ask them to come home. And so now, my household size has increased, but that was because--

BEN OLKEN: Exactly, so imagine I-- so just so everyone-- if everyone heard, so basically the idea is imagine I have this big labor-demand shock. I need more land. I invite some cousin to come stay with me. Now, my household size has increased. That's a violation of this, so that's exactly right.

So what they're going to do-- so that's an endogenous change. That's what we don't want. What would be an-- so that's a problem. What would be exogenous change in household composition? Yeah?

AUDIENCE: Someone who surprisingly dies.

BEN OLKEN: What? Sorry?

AUDIENCE: Someone surprisingly dies.

BEN OLKEN: Yes, surprise deaths that would absolutely be a change in household composition. What else? What's a more common change in household composition?

AUDIENCE: Unplanned births.

BEN OLKEN: What?

AUDIENCE: Unplanned births.

BEN OLKEN: Yeah, unplanned births, but actually-- OK, so I agree with you. That's more common than surprise deaths, although actually I'm not sure about that. I'm not sure which is more common. But is there something even more common than that they can use? Yeah?

AUDIENCE: Gender of the kid.

BEN OLKEN: What? Sorry?

AUDIENCE: Gender of the kid.

BEN OLKEN: Potentially, if you basically think that-- yeah. But I think that one thing you guys are forgetting is that not everyone in the household actually is of working age, and that's actually what they're going to use. So you guys are thinking of changing, people coming into the household and out of the household, and you're right. Those are the kinds of shocks that we would use for that.

But you get a new-- an infant is not going to affect your household labor supply. Kids have to be a certain age before they can start contributing to labor on the farm. And likewise, at the other end, after people get a certain age, they're typically not supplying as much labor.

So what they're going to do is basically they're going to use the aging of household members as surely predictable changes in your household labor supply, but changes that are uncorrelated with labor-demand shocks. So the idea is even if you thought like, oh, we need more kids in the farm, and we're going to have more kids, there's a time lag there.

So we have a-- they have a kid who starts at age 0. They don't start working until they're, I don't know, say, teenagers or whatever. That's when the-- that shift from being little kids that can't work to teenagers that can work or adults that can work, that's the variation they want to use. Does that make sense? So that's what they're going to do.

They also have just really terrific data. And one thing I will mention, by the way, is this was an example of-- as far as I know, and the authors correct me if I'm wrong on this. But I'm pretty sure this was data that was collected for a totally different purpose.

This was data that was collected because they were doing a long run study of iron supplementation in Central Java. And so they wanted to look at what was the effects of iron supplementation on labor supply and a bunch of other things. And they had all this very rich [INAUDIBLE] data, and then realized they could use it to answer this other question.

And so I actually think one useful thing for you to think about in terms of data sources in general is that there are all of these data sets that have been collected for all kinds of other things. Someone is doing a study of dust and such, and collects a bazillion variables to study it. And those data sets are often incredibly rich and can be used for other purposes.

And so actually, if you look, for example, in the JPL Dataverse, for example, we have lots of different websites from JPL studies or other sites that compile these as well. And it's worth thinking about whether or not you have data sets that can be collected for one purpose and repurposed. I did this a bit in one of my own papers.

I have a paper on social capital, for example, which was repurposing a data set that was used for-- that I collected for a study about corruption. So you can think about how these data sets can be repurposed. And I expect there's lots of good data out there that hasn't been exploited. And this is a really nice example of a paper that was ultimately published in *Econometrica* using a data set that was collected for some other purpose.

So what's the regression? The regression is going to be labor demand on the farm as a function of number of people in the household of different size bins. And they're going to control for household fixed effects, and they'll also control for community times time fixed effects.

So they're controlling for the fact that maybe-- so why would you care about this? If you imagine there were, I don't know, secular growth in household size, not that that's necessarily true, but if there was secular growth in household sizes and there was secular growth in labor demand, you'd want to control for time effects too. So that's going to control for that.

And they're going to parameterize this N into different age bins. And to isolate these exogenous changes, they're going to restrict themselves to households with no change in household size due to membership because of migration, births, or deaths. They're going to just look at the changes due to the natural aging. Is that clear?

All right, OK, so this is a funny regression because there's lots and lots of coefficients. Here are the different age bins like 0 to 14, 15 to 19, 20 to 34, et cetera, et cetera. This is household farm demand on these various things. And in general, you see that farm demand is increasing as you-- log labor demand as a function of the number of people in these different groups is increasing for all the different age ranges, OK, fine. If you use the variation from aging only, it's-- the standard errors are a little bit smaller.

By the way, they do this separately for men and for women. It does look like you get-- for example, in this age range the 15- to 19-year-olds, it looks like, for example, you're getting impacts on labor demand from boys but not from girls. In other age ranges, you get it from both.

But the key thing actually of all of these joint tests. So the thing you want to do is reject-- if we want to reject separation, we just want to do a joint test that-- of all of these demographic variables, not predicting labor demand. And that's what this is right here.

And so you can see they can strongly reject the null of separation that basically-- so on net, all those dummies are-- all those variables of the number of households of different sizes are-- they can reject the null of that. In some sense, this table of a bunch of p-values of joint tests is sometimes the main results. So substantively, it actually looks like we can reject separation with much better data and a better empirical strategy. Yeah, Kyle?

AUDIENCE: Why do you think it published so well? Was it just like a really important--

BEN OLKEN: What? Sorry?

AUDIENCE: Why do you think this paper was published so well? Is that a very important question? Or?

BEN OLKEN: I think so. I think that we'd all been teaching that Benjamin paper for-- since 1982 and having to complain about the empirics not being up to modern standards. And so having someone who can say, look, this is a classic development economics question, and I'm just going to nail the empirical of that, that's a nice paper.

And actually, that could be like-- if you can find great, classic questions with-- that are waiting for a good empirical test, I think that's a great result. I would have totally been happy to publish this paper. It's a great paper. Other questions?

And I was so excited to finally be able to update my slides when it came out and be like, aha, we can now actually say something a little more positive about that. And so I think that's great.

And you'll see lots of papers that have that feel of, here's a classic theoretical test, and then someone's going to really nail the empiric. For example, if you think about the poverty trap section, I think Esther taught the Balboni et al. paper on the estimation of the poverty trap. And that was-- in some sense, what's so exciting about that paper?

Well, we all knew that we could write-- we could all right, do this graph of KT plus 1 versus KT, and drew something kind of funky that looked like that. But actually really nailing kind of what the shape of that curve looks like and establishing whether or not there are multiple equilibria or not is something that I think everyone is really excited about. So I think that's another aspect, that's another reason that paper is so exciting, is because that basic sort of theoretical framework we've had for a long time, but actually being able to show it empirically in a convincing way is really nice.

AUDIENCE: It's kind of a significant amount or something?

BEN OLKEN: I think so.

AUDIENCE: [INAUDIBLE]

BEN OLKEN: On both those examples.

AUDIENCE: Yeah.

BEN OLKEN: Yeah. But again, the key there is finding a classic-- it's not some random theory paper that no one's ever heard of. This is a classic poverty trap model that everybody knows. This one, this Benjamin paper, is something that we've all been teaching as one of these key questions in labor economics and development for a long time.

So coming in and saying, this is going to be the definitive kind of empirical piece that goes along with this theoretical piece? That's awesome. Other questions?

OK, all right, so what's going on? So if we buy that there's already separation failures, what is going on in these labor markets? Why is this happening?

Well, if you think about the theory, it said something is going on that's making these labor markets not clear. That was kind of the key. That's why we got these separation failures. And so why?

And so there was a bunch of papers that have been starting to understand and try to unpack in a little more detail, what is going on in these rural labor markets? Why are they not clearing? So I'm going to talk about a few of them.

So in particular, I'm going to talk about several papers that look at this idea of nominal wage stickiness. So how many of you guys have heard of nominal wage stickiness, like for macro? Some of you? Not everyone.

OK, so the basic-- so somebody's heard of it? You want to summarize it? What's the idea?

AUDIENCE: You can't adjust wages downwards, or the nominal wage downwards, [INAUDIBLE].

BEN OLKEN: Because people don't like it. Right, exactly. So this is the idea that there's an asymmetry in wage adjustments. Everyone is happy to have their nominal wage go up. People do not like their nominal wage to go down. And for some reason, this is true about nominal wages and not so much about real wages.

Because your nominal wage-- and even though it's harder for our models to think about, it is much easier for actual people to think about nominal numbers than real numbers, because your nominal wage is how much you're actually being paid every period or whatever. And that's really important for-- that has really important macro implications, as you may have learned about in macro, because it means that wages don't-- if there are shocks to labor markets, and you have nominal wage rigidity, and you have no inflation, then, basically, you have all these labor markets that are not clearing because you have people whose wages are not ticking downwards.

And so at the macro level, that's often one of the main arguments why we should have a non-zero optimal inflation target, because we want to have a little bit of inflation so that these labor markets can-- because if we have nominal stickiness but not real stickiness, a little bit of inflation means the real wage can move around a little bit and those markets can clear. OK, so you might-- a natural question is, these ideas about nominal wage stickiness are very intuitive in very formal settings with formal wage contracts, and my wage is set-- MIT has set my wage. I have a letter saying what my salary is. It's all very by the book.

It's even more unsurprising in cases where there are labor unions involved. But it's not obvious that this would translate into spot labor markets for agricultural labor in a developing country setting. It's like, we need to go hire some-- we need to go to the village center and pick up a few laborers to help clear our field today. Is that a setting where we expect nominal wage stickiness or not? And it's not obvious.

But you can see why these would lead to separation failures, right? That was actually one of the-- so yes or no? Why would this lead to separation failures? OK, [INAUDIBLE]?

AUDIENCE: If you're speaking of the household, you don't have a formal kind of [INAUDIBLE]. Within the household, you're not setting informal contracts.

BEN OLKEN: Right, but more generally, which of those like various little graphs with the separation failures does this correspond to? Someone? Someone take a stab at this?

AUDIENCE: The market wage is too high.

BEN OLKEN: Yeah, the market wage is too high, exactly, right? So we'd like the wage to come down to clear, but it's stuck too high because of nominal wage stickiness. So it's the case where the nominal wage, the market wage, is too high. So we can't get enough-- we'd like more work. The households would like more work at the market wage, and they can't get it.

All right, so Supreet Kaur has a paper testing this idea of whether there is nominal wage stickiness in rural labor markets. And here's her idea. Her idea is very simple. If you have nominal downward stickiness, then-- sorry, back up. She's going to use rainfall shocks as shocks to labor demand.

So when you have good rainfall, we need lots of labor to clear, do all the agriculture activities. When we have bad rainfall, there's not that much to do because the crops didn't grow that well, so we have less labor demand. So we have rainfall shocks as a measure of labor demand.

But with nominal downward stickiness, the sequence of shocks is going to matter. So in particular, through a positive shock and then a negative shock, the wages are going to go up after the positive shock, and they're going to be stuck too high. And they're not going to fall during the negative shock. Whereas if I have negative chocolate and a positive shock, I have a problem.

And in particular, the key thing that you're going to look at is, in some sense, if I have a negative shock now-- so I want a low wage now-- then the lagged shock is going to matter. So if I want a low wage today because I got a negative shock, and I had a low wage yesterday, then we're good. But if I have a negative shock now, so the optimal wage is low today, but I had a positive shock in the past, then the wage is going to be artificially too high.

Is that clear? Yes, no? [INAUDIBLE], no?

AUDIENCE: But what if-- if it's sticky, then why would the first negative shock push the price down? What do you mean. Let's say you have two negative shots. Won't the first one be sticky as well? [INAUDIBLE] it's not going to fit.

BEN OLKEN: It only sticks-- sticks the thing isn't-- so but it didn't go up. It's almost like the key of a negative, , it wasn't a positive shock. That's the thing to think about.

So imagine we're just at some base case, and then, in period 1, we get just a neutral or a negative shock, so there's no increase in the wage. So the wage is over here. And in the other case, we get a positive shock, so the wage is over here. Now, the next question is, if we get a negative shock in the next period, over here, it'll still be the wage will be down here. Over here, the wage will be too high.

So the key point is kind of that if you get a negative-- whereas, if you've got a positive shock in the next period, then this way, you can just flip right back up. So with positive shocks in the second period, the lag doesn't really matter. With negative shocks in the second period, then you have this persistence of the lag. But basically, the positive shock means the wage is way too high. That's the key asymmetry in our test.

Does that make sense? OK, so if she's going to do this, she has data on a number of years from Indian districts where she has rainfall shocks, and she has the agricultural wage. And so this is kind of-- she goes has two-- this is from 1956 to '87, so she's got a lot of years.

And this other, I guess, this from the World Bank data sets, and also this NSF data from '82 to 2008. And so the key thing is she's going to look at shock T minus 1 and shock at time T. So what are the key things we can learn from this?

So first point is, positive shocks in the second period increase the wage. That's not surprising. It doesn't really seem to matter what the wage was last period.

So if the shock is high-- this is all relative to 0, 0, like neutral-neutral. So if the wage was high this period, it doesn't matter if we were in a drought or had a high wage last period. Those are about the same. If anything, actually, it looks a little weird over here. It's negative. But certainly, it's not the case that drought last period is pulling the wage down today.

On the other hand, if we have a drought this period, or a negative shock this period, the wage is higher if we had a positive shock last period. So that's where you can see, for example, like these coefficients, to the ones in this bin compared to the ones in this bin.

OK, is that clear? I know this a little small, but it seems like guys have all your printouts, or your iPads, or whatever. So hopefully, you can read it, right?

So that's-- this is the key asymmetry. Is like over here, it doesn't matter, but down here, it seems to matter. So the wage is stuck too high in the previous period.

If you take this idea that this is about nominal rigidity, then these effects should be smaller if there's inflation going on. So this is a period over a 50 or, whatever, 30-year period. There's different inflation rates, so you can interact this whole thing with inflation. And if you remember, the key place we thought there was nominal stickiness was like negative shock this period, but positive shock in the last period. So you can interact all these shock things with an inflation measure, and the key thing you're going to find is a negative coefficient over here, which is basically that this effect of the wage being too high goes away when there's high inflation.

OK, and the final question-- and I don't remember why they have acres per adult in the household, actually. I don't remember that. But the final point is, does this actually lead to changes in aggregate employment? So they have a measure of aggregate employment, and if you find that the shock is--

OK, so a couple of things. So something, actually, is going-- something is weird with this coefficient. It's negative over here and positive over here, so maybe there's something-- I suspect there's something off on that one over here. I expect that's probably negative.

Yeah, so in general, positive shocks lead to more employment. Negative shocks lead to less employment. But when you have a negative shock today, a shock equals-- this table is a little funky because sometimes it's 0, and sometimes it's negative. I don't exactly know why she does that, frankly.

But in general, when you have positive shocks in the past period and negative shocks today, you get more negative effects on employment than if you just had negative shocks today. So that lagged positive shock is giving you an extra bump in-- an additional negative employment effect above and beyond just the negative demand shock today, and that's because the price is stuck to high, right?

And so what's going on there, what you'd like to have happen-- so what you'd like to have happen if you have a negative labor demand shock, the wage should fall, and employment should also fall. But if the wage is stuck too high, then the quantity of the reduction in labor will be even larger. So this quantity reduction here to here is less than here to here because you don't have the full wage response.

Wait, is that clear? Yeah, OK, clear? OK.

So at some level, I think it's actually a pretty remarkable set of results if you actually step back and think about what's actually going on. What's going on is that we have these spot-- you would think of these spot labor markets where-- but somehow, these norms about not cutting your nominal wages seems to be taking place there, too, OK?

So here's-- this is one way of showing this. They have another paper-- or actually, Supreet Kaur and a different set of coauthors have another paper-- looking at this experimentally, which I'll show you in a second. Yeah, Leslie, you have a question? No? Comments, questions? OK.

So that says-- it suggested some kind of nominal wage stickiness. Another approach could be, suppose you wanted to test for this lack of market clearing experimentally. How would you do that? Not necessarily this pattern of downward-- so broadly speaking, the previous paper was saying, look, we're testing for failures of the market clearing by testing for one particular one, which is these patterns of shocks. More generally, if we want to look to see whether the wage is not adjusting to the market clearing level, what else could you-- suppose you want to do it experimentally. What could you do? John?

AUDIENCE: You could do a shock to either the labor demand or the supply, and [INAUDIBLE].

BEN OLKEN: Exactly, right? You can just shock either labor demand or labor supply. So in particular, we just literally-- if we can somehow experimentally shock this, we can see what happens. And in particular, if you think that the wage is stuck too high, the key prediction is I can experimentally shock, say, labor demand and see if the wage moves.

OK, so that's they're going to do. So that's literally what they're going to do. They're going to work in village labor markets, and they're going to randomly hire 24% of the labor market in some village to go do some work in some factory-- which I think, if I remember correctly, they were using for some other paper also. I think they needed a bunch of laborers to go do some things in a factory to study some other thing. I'm not 100% sure of this, but I'm pretty sure it's right.

And so they said, well, if we can just randomize where we're getting them from, getting these workers from, we can also study the effects of this labor [INAUDIBLE] shock.

AUDIENCE: Yeah, how is it not prohibitively expensive to do?

BEN OLKEN: I think they're doing it for like a month, and they picked small villages. Yeah, I mean-- I mean, I actually don't know the details, but I think it's something like that. So--

AUDIENCE: [INAUDIBLE] or something like that?

BEN OLKEN: I don't know where they got the money. But actually, hold on a second. So suppose that-- I don't know. So I don't actually-- you can look in the paper and see what the numbers are, but if it was-- I mean, I don't know. Suppose it was 24% of the labor market. Suppose the village had-- it was a small village, a couple hundred people, so 50 people. 50 people, and they make, I don't know what the wage is, a couple of bucks a day.

So $100 a day for 50 people times a month would be $3,000 per village. And you do that for 50 villages, that's a lot of money, but it's maybe not like-- it may not be millions and millions of dollars. And the other point is, and they also need these people anyway for some other thing. So they were getting multiple-- they needed to hire some labor for some other activity, I think. And so maybe that's how they did that, I suspect.

Those numbers, by the way, are not the actual numbers. Those are rough and ready calculations. I'm sure if you look at the paper, you could figure out what the rough number is.

OK. OK, so that's what they're going to do? So what are the predictions? Like what should you look at?

AUDIENCE: Wage would go up?

BEN OLKEN: What's that?

AUDIENCE: Wage would go up?

BEN OLKEN: OK, so right. So this is-- so they're going to call this a labor supply shock. So it's going to be a negative labor supply shock.

AUDIENCE: Why are they calling it a supply shock?

BEN OLKEN: Huh?

AUDIENCE: Why are they calling it a supply shock?

BEN OLKEN: Because we're pulling laborers-- well, so you can do it either way. It actually doesn't matter.

AUDIENCE: So it's a negative supply.

BEN OLKEN: So you can think of there's a bunch of work in the village, of like farming work or whatever, And we're going to reduce the labor supply to that farming work. I think that's how they think about it. You could also think about it as the total labor market in the village and a positive labor demand shock. I don't think it's really going to matter, but his is how I think they think about it. It's like so the labor market for village farm work has just had a negative labor supply shock.

OK, right, so one view is we went from here to there, so the wage should go up and the quantity of labor should go down. that's if we're in a competitive market. What else? And so if it was rationed, what would you expect to have happen? Yeah, Christine?

AUDIENCE: The weight shouldn't change.

BEN OLKEN: Yeah, what about the quantity?

AUDIENCE: It shouldn't change either?

BEN OLKEN: Right. So if we're rationed up here, we're just living on the labor demand curve, and nothing should happen, OK? So any threats to-- anything else that you should be worried about here? Yeah.

AUDIENCE: If your shock is too big, then you're no longer going to be [INAUDIBLE]. If you push your labor supply shock too far up--

BEN OLKEN: Yeah, right, the shock can't be too big. Because if the shock is too big, then actually, then, it'll be off the constraint. So it has to be-- you're looking for the right size shock to be detectable, but not so large that it busts-- it goes over that clearing point. That's one. Number two.

Yeah, what else do you think about? Anything else? So what would happen if-- let's see if I can draw this right. So we have an-- and we had, that's our labor supply shock. And we had totally inelastic labor demand.

So what happened there? This is-- that's right, so what's that?

AUDIENCE: No change in the quantity. But the wage then goes up?

BEN OLKEN: Yeah, no change in quantity. The weight goes up. And likewise, if I drew it totally inelastic, you can have a change in quantity but no change in wage. So you might get one but not the other.

So we need to look at both of them. We need to look at both here, because if we only looked at one of them, we can't necessarily rule that out. And what about-- what if the labor supply curve is totally elastic?

Leslie, what do you want to say? No? You're thinking-- or Patrick, somebody?

AUDIENCE: Nothing happens.

BEN OLKEN: Nothing happens, right? So I think the one alternative-- to my mind, one of the alternative things they have to deal with is, maybe the labor supply is totally elastic. Because if labor supply is totally elastic, then I pull out a bunch of labor, but a bunch of people, they say, it's a reservation wage, a bunch of people just enter, and nothing would change.

OK, so just keep that in mind. And they have-- actually, does anyone remember, did anyone read the paper and remember how they talk about labor [INAUDIBLE]? They have something in the paper to argue this is not the case, but I don't actually remember at the moment exactly why they argue this is not the case. But it's good. They do have a paragraph or so on this one. To my mind, that's the one I would look a little harder at.

OK, the other point is, do you think the response should be the same throughout the year? So here is the key fact. A key factor about agricultural labor markets is there are peak periods and non-peak periods.

Peak periods are when we need to plant and we need to harvest, and people are doing a lot of work then. And then, the slack period is the rest of the year. So do you think there would be likely to be differences between these two different parts of the year? Aaron, what do you think?

AUDIENCE: They would be more likely to pick up the effect that they're they're interested in the peak period, when there's an active market for--

BEN OLKEN: Which one?

AUDIENCE: In the peak period.

BEN OLKEN: Why?

AUDIENCE: Because this is when there's an active market for agricultural labor, and we are less likely to get an excess labor supply situation, where taking--

BEN OLKEN: Sorry, more likely to pick up which effect? What do you mean?

AUDIENCE: Like if we're in a slack period, and there's a bunch of excess labor, and you take 24% of the labor market away, it's probably not going to lead to a detectable effect on wages.

BEN OLKEN: OK, sorry. So going back to our little graphs, so what do you have in mind? So another way to think about it is this is our labor supply, and this is our labor supply shock. And maybe this is labor demand in peak, and maybe I should use different colors. I'll get those different colors going.

OK, this is labor demand peak, and I guess I only have two colors. Oh, no, we have more colors. That's excellent. Let's sort. OK.

And that's labor demand in the slack. And this is labor supply. And we're going to shock. And this is our labor supply shock.

OK, so in the peak period-- this is all market clearing. Here, we're going from here to there, and in the slack period, we're going to go from here to there. So that's all on the market clearing.

OK, so now, everyone-- and by the way, so one assumption I've made in drawing these graphs is I haven't tilted the labor supply. First of all, I've assumed labor-- sorry, I've assume labor supply is the same in both cases. I also haven't changed labor demand to make it more or less elastic in these different periods. You might also imagine that labor demand-- I mean, another thing you could also imagine is maybe labor demand is less elastic in the peak period because I really, really, really need to get those crops cleaned now, otherwise they're going to rot, whereas in the slack period, I have a couple, some stuff to do around the farm, but it could get done now, it could get done later-- like, whatever. And so that would make it more elastic in the peak period.

But just for the moment, with this, with this graph like this, does that give you any thoughts on how things should be changed, or whether you'd expect anything different in the peak or the slack periods? Yeah, Whitney? Oh, I mean Jenny, sorry.

AUDIENCE: So If you are thinking about rationing being like a price floor, you're more likely to have that the case in the slack season.

BEN OLKEN: Yeah, exactly, right? So if we draw the rationing thing to be, say, you suppose, for example-- sorry, let's draw it-- ba ba ba bum. Let me redraw it a little bit. Yeah, let's try it over here, OK?

And let's put the-- Wage floor there, that, right? So if that's the graph-- hopefully, you all can see my graph-- if we have the wage in general is going to be much higher in the peak season, and if I just fix a kind of nominal wage floor, then, in this example, in the peak season, everything behaves nice and competitively. The wage goes up, quantity goes down, because we're above that wage floor. But in the slack season, we're the whole activity is happening below that wage floor, and it looks as if we're rationing. And this is going to be, basically, what they're going to argue is happening.

Is that clear? OK, so this is how they-- so this is basically just showing you this is [INAUDIBLE] true. So, say that? OK, fine.

So what they're going to then do is they're going to then look at the hiring shock-- that's the negative labor supply reduction-- interacting with whether you're in the peak period. And they can measure this just by being the peak agricultural period in these areas, or by just the overall employment level as a measure of how high labor demand is. So OK, so the first thing they're going to do, they're going to look at what they call the spillover sample. So what is this?

So what they do, and this is actually-- I think this is kind of clever-- so what they do is they have to keep track of who's who in these different labor markets. So what they do is, before they start, they recruit a list of workers who are interested in their factory jobs. That's going to be the potential treatment set. And they randomly pick which villagers they're going to hire from. And within villages, they randomly pick people from this list.

The spillover sample is everyone who is not on this list, who did not say they were interested in one of these jobs. And the point that I want to make is, because they recruited this list ex ante, before they randomized, they know who was in the spillover sample in both treatment and control villages kind of identically. Because step one is I split the villages into, say, this half of the villages, if you're a treatment area, we're going to give you factory work. This half of the village, if you were in the treatment area, we're not going to give you factory work.

And we know who you are because you identify yourself beforehand. And then, when I randomize you into the experiment later on, it's then totally OK for me to just look at this set of people and in control treatment areas as the rest of the labor market because I identified them identically in treatment control areas. If I hadn't done this step, it would still be OK to look at the total labor market, but I couldn't just specifically look at what's happening to this other set of workers because I wouldn't know, in the treatment area, in the control areas, who would have been the people who would have taken over the factory job and who wouldn't have been the people.

You see what I'm saying? So the fact that this is all very precise and done ex ante is very helpful. Clear?

So the first thing to look at, the spillover sample. One thing you should know about, by the way, is that rural wages often take the form of a bit of wage and a bit of in-kind. They'll give you some cash, and then, maybe, they'll give you lunch, or something like that. So you actually have to-- and actually, for a rural laborer making subsistence wage, lunch can be a non-trivial fraction of the total compensation. So you have to be careful to measure that, too.

Like lunch and snacks. I feel like in one case, it was like-- I said it was like lunch, a snack, and a couple of cigarettes. It was like that stuff actually can be non-trivial relative to the whole conversation, so you have to include it too. OK, so in the spillover sample, what do you find? So some want to interpret what we're finding here?

AUDIENCE: The shock has no effect in the-- a period of very little effect and--

BEN OLKEN: What are you saying? What, sorry, what was that?

AUDIENCE: The shock, the supply shock, has very little impact on wages [INAUDIBLE].

BEN OLKEN: Right. I don't remember why it's called [INAUDIBLE]. Let's just call it [INAUDIBLE]. So what do you think of that?

AUDIENCE: It seems consistent with--

BEN OLKEN: Our nice little graph over there?

AUDIENCE: Yeah.

BEN OLKEN: Yeah, exactly, right? OK, so that's on wages. And so exactly, that's exactly what You would-- that's exactly this point over here, which is in the non-peak period, it looks like they're constrained, like the wage is unaffected. In the peak period, the wage is affected.

And the second thing is we want to look at employment. So this is still the spillover sample. Ignore this bullet point. I changed the graph and didn't change the bullet, so just ignore that.

So this is the-- OK, sorry, so what do you make of this? This is employment from the spillover sample.

AUDIENCE: Decrease in employment during the peak season, and less of a decrease in employment during the offseason.

BEN OLKEN: Say it again, sorry. I wasn't listening. Say that again?

AUDIENCE: So the hiring shock in the spillover oversample causes less of a decrease in employment during the off-peak season than during the semi-peak season.

BEN OLKEN: Not quite. So I think-- so you're right that there's an interaction effect, but the main effect you shouldn't-- I think you're interpreting the wrong way. So it's positive in the off-peak season.

AUDIENCE: Yes.

BEN OLKEN: Right? So employment goes up in the off-peak season, and in the peak season, it doesn't change.

AUDIENCE: Yeah.

BEN OLKEN: Does everyone see that? So for the off-peak season, it's just positive, whatever, 6%, 10%. For the non-peak season, it's basically 0 because we add these two things. I'm sorry, for the peak season, it's basically 0 because we add these two things, right?

OK, so what is that telling us vis a vis the model over there? Is that consistent, not consistent? Let me show you one more thing. One more piece, and then we'll put it all together.

This one is actually a little harder to think about because this is the spillover. This is what's happening just to the spillover-- I'm sorry, I almost lasered you guys-- to the spillover sample, the people over there who were not, potentially, included in the hiring sample. This is total employment. So now, we're just going to pool the whole village together. And what do we learn from this? Patrick?

AUDIENCE: That the individuals from this spillover sample--

BEN OLKEN: This is everything. This is everybody.

AUDIENCE: Right, right, but does this mean that individuals in this spillover sample increase their level of supply [INAUDIBLE]?

BEN OLKEN: Yeah, OK. First, before we try to interpret what this means, what is this table telling us? It may have been easier, actually, if I started with this one before the spillover sample. So let's take the weighted results-- Fine. No effect in the slack period, big positive effect in the peak period. And now, let's look at total quantities, and what do we find? Hazel, what do you think?

AUDIENCE: No effect in the off-peak, and then [INAUDIBLE] goes down.

BEN OLKEN: And what, sorry?

AUDIENCE: Total quantity goes down in peak season.

BEN OLKEN: Right, total right quantity goes down in the peak season, right? OK, so what is this telling us? Is that consistent with that or not?

AUDIENCE: That's consistent with the left-hand.

BEN OLKEN: Exactly, it's exactly consistent with left-handed graph. So then, so right, so what do we see? And so in the non-peak period-- we're stuck over here. Nothing happens. In the peak period, it's going up.

So it seems that these two results are the key results. So why? What is the spillover result telling us?

Well, you might have just thought, something weird is happening in the labor market. Nothing is going on. Nothing doing in the slack period.

What's nice about the spillover result is that actually, no, no, no, in the spillover sample, in the non-peak period, what do we expect to happen? Well we expect, we just took a bunch of people out of the labor market-- out of this rural labor market. Some of them were working. The total quantity, the total people working in this labor market, should be unchanged. Therefore, we have to pull in some of the spillover people who have to come and take those jobs. So that's the final piece that's showing you that something's actually happening in this period. Is that clear? Yeah, Whitney.

AUDIENCE: So when you're actually designing this experiment, and you want to get this effect again, how do you know the 24% is the number of people you need to be pulling to get this?

BEN OLKEN: Good question! So how do you know that this is a big enough effect? Well, so the answer is you would do a power calculation.

So you would say like- so you don't know. I mean, you don't fully know. And actually, I would love to hear the author's answer to that question of what exactly they did. I suspect the truth was, it the big enough shock they could afford. [LAUGHS] If you actually asked, I suspect that's what's going on.

But I think you can actually do a little more. So you can say, well, look, you can calculate-- so what can you do? So you can calculate, if you have some estimate for the slope of the labor supply and labor demand curves, you can say, if I shock labor supply by 24%, how large an effect would I expect on wages, and how large an effect I would expect on quantities?

That's just, you have these elasticities, so you know the percent change in-- so whatever. If I have a 24% reduction in labor supply and a labor demand elasticity of whatever, I can multiply the labor supply and elasticity times the [INAUDIBLE], divide it, and figure out how big an expected change you're going to get.

So what I would do if I was designing this experiment is I would, A, come up with my favorite labor supply-demand elasticity from some other estimate, or at least try to get it from somewhere. I guess you also need the labor supply. You need both of these estimates. So take your favorite of these estimates.

You would-- oh, no, sorry. Mostly, you just need the labor demand elasticity. I would take my favorite estimate. I would multiply it by 24% to figure out how big a change in wages or quantities I would likely to be expected to get.

And the second thing I would do is I would then do a power calculation and say, well, I'm going to measure wages in this village in quantities in the following way. I can get some estimate from some other data sets that I've done of what the standard errors are likely to be in this regression. And I can put those two things together and say, I expect to find a difference in the wages between these two periods of, say, 7% based on my best estimates of those elasticities and the size of the shock. What's the probability I could detect that given the expected standard errors I'm going to find? And that's what a power calculation exactly does. So that's how you would do it.

Now, that requires two things. It requires having the correct estimate of those elasticities, or at least something that's in the ballpark. But even if you don't get it exactly right, you can get-- the correlations are never perfect, but you can at least get it in the rough ballpark. And it also requires having some estimate of what the standard errors are going to look like. And that's going to depend on the data generating process, and what your sample looks like, and so on, and so forth.

If you've done other samples of wages in these areas, you can do a pilot survey to get some of that data, and you use that. So that's actually how you would do that. Sorry, Paolo, got another question?

AUDIENCE: I just have a very quick, basic question. What is the employment level here in these tables?

BEN OLKEN: So actually I will tell you. I don't remember exactly, but I believe it's a measure of total aggregate employment in the village in that month as a way of-- as an alternative measure of measuring whether we're in peak season or not. Basically, so one is we can define peak season based on season. The other is we can define it based on q, [INAUDIBLE] of q. I don't remember the exact definition I used of q, but that's the idea. Jenny?

AUDIENCE: I think I'm still confused about why it's important to distinguish the spillover group from people who don't get the job but were interested, unless you're worried about--

BEN OLKEN: Composition effects. So for wages-- wages is the easiest case. Wages, it's easier. You see why you want to do that?

Imagine that we have good workers and bad workers, good agricultural workers and bad agricultural workers. So good agricultural workers are the ones who are doing the job in general. Imagine I just take a bunch of those agricultural workers and pluck them out. They're going to be replaced with lower-quality agricultural workers.

If I'm measuring the wage, not particularly-- like the wage I'm going to measure might be your average marginal product. If I replace good workers with bad workers, the wage is going to fall, not because of these effects, but just because I've moved to having worse quality workers who are being paid based on their lower marginal productivity. So the reason for doing it on the spillover sample is we're conditioning on the same set of people.

It's not perfect, actually, now that I think about it because you're getting different-- it's among people who are working, so you are still getting some selection. But I think it's a little bit less. I think that's because it's still conditional.

So if everyone was working, and it was the same people working both periods, it would be totally clean. Now that I think about the explanation I just gave you, you still could have some selection within the spillover sample, but I think it's somewhat less.

AUDIENCE: So the idea is that people who are interested in applying for this external job are more likely to already be in the labor force and have selected [INAUDIBLE]?

BEN OLKEN: No, it's that there's a mechanical selection effect if we pull them out of the labor force. The selection effect is much larger for that sample. Because I'm literally taking these people out, and so there's like a large-- I think the mechanical effect of who's in the labor force is much larger over there.

But I agree, actually. Now that I say it, it's not perfect either. But I think that's, at least, the theory for how we think about it. Patrick?

AUDIENCE: Why not introduce a variation of degree of [INAUDIBLE]?

BEN OLKEN: Oh, OK, so tell me why you would want that, and then I'll tell you why you don't do that.

AUDIENCE: Maybe, then, we would see-- well, in the first case, when you were talking about [INAUDIBLE], I reminded you in the end, if you can make a connection with the [INAUDIBLE], you actually risk overshooting.

BEN OLKEN: Overshooting?

AUDIENCE: Yeah, yeah, yeah.

BEN OLKEN: Yeah.

AUDIENCE: And [INAUDIBLE] to not only avoid that, but also, actually see to what extent it's [INAUDIBLE] over the [INAUDIBLE] talking about.

BEN OLKEN: Yeah, so OK. So Patrick's point, just to be clear, is, in our example, maybe I should do multiple shocks in case one of them might overshoot. So with infinite sample and infinite budget, agree. The reason that they don't do that in practice is a power issue, basically.

So if I split my sample, imagine I-- this is a case, by the way, where your treatment costs a lot more than your control. You have to pay for all these jobs. So the relevant case is not even instead of half-half, like a third, a third, a third, it's probably like a quarter, a quarter, control, is the budget-neutral thing to do. Approximately, not quite, because you have fewer people employed and then a smaller shock group, but let's call it roughly that.

So then-- actually, no, I take that back. That could be, actually, first order, actually. Because you have another-- actually, you have another question here. OK, there's two experimental design issues.

The first experimental design issue is, is it better to have a smaller number of villages with a big shock, or more villages with a small shock? So actually, that is relevant.

So you could-- for the same price, I could have 100 villages with a 24% shock, or 200 treatment villages with a 12% shock. This is the same cost. Which is better?

Forgetting about this issue for a second, just general straight up-- like forget about this more complicated issue. Just straight up experimental design, which of these do you prefer?

AUDIENCE: The [INAUDIBLE] one seems to have more power, I think.

BEN OLKEN: Yeah, so does everyone understand why? Did I go over this already in this class, or was that in the other class I taught this? I can't keep track. Aaron, do you remember which class I went over this issue?

AUDIENCE: [INAUDIBLE].

BEN OLKEN: OK, someone else? Relevant to someone. Someone else, did I explain this already with the beta and the T-stat thing? Yes? OK, great, exactly. Then you were paying attention.

Good, so exactly. This one has a higher power. So the first point is that you-- the lower shocks, you have less power. The second point is related to that, which is if you don't know if you're going to have an effect or be able to detect it, it is almost always better to have just one control group and a really big treatment group, rather than an intermediate effect.

And that's because-- it's like if this is the control, and this is the treatment, in general, we think the response of most things is linear-- like approximately linear-- then we're trying to figure out, we have this kind of small thing here. We have, normally, a lot of power on that small thing. And we'd be trying to test, are we on this linear-- like we're usually interested in, are we on this linear curve or not?

And so testing not just like, are you-- identifying this lower range thing and testing whether you're on this linear curve requires a lot of power. So most of the time, these intermediate outcomes are very, very difficult to distinguish from being on the linear, just on the linear curve.

Sorry, the reason for that, I should say, is like you have less power here, and you're trying to test that this coefficient is equal to 1/2 of that coefficient. It's a very tight test.

Now, so that's the general reason why, if you're not sure if you have an effect, you would need a really big sample for it to make sense to do these 1/2 power intermediate outcomes. In your case, that's not quite right. In this case, it's not quite right because we have this additional hypothesis, this additional issue of maybe we're overshooting. But again, the problem is-- if you ran that test, you would want to say, well, suppose you found in that case that-- the only relevant case is the case where, in the full one, we find labor supply and labor demand effects, and in the smaller, we don't. That's the overshooting case.

But then, you have the other problem of, well, the treatment was also half the size. So you also have to-- you'd have to rule out not only-- you're kind of in the same case. You have to rule out-- you have to rule out, not only are these things different-- like this one's zero and this one's not-- but actually, you have to also be able to reject a harder thing, which is this is not equal to just half of this effect.

You see what I'm saying? The thing you're testing is not just that this would be 0 and this would be the full thing, but you also have to reject that this thing is not just half of that, because that would just be linear. You see what I'm saying? So the amount of power that you would need to do that successfully would be pretty large. Does that answer your question?

AUDIENCE: Yeah.

BEN OLKEN: Yeah, so I think that basically, I mean, another way of saying it is I think on these things, your best bet is-- even though it's very frustrating, very nerve-wracking as a researcher, it is often not the case that the best thing to do is hedge your bets with two different treatment effect sizes.

And by the way, if you see a lot of the-- if you think about how they're running, how they run clinical trials, if you think about the COVID vaccine clinical trials-- which are probably clinical trials everyone is most familiar with, I'm guessing, at the moment-- that's what they did. The first trial of, are we going to approve the vaccine or not, they did a bunch of little small things to check the dosage and make sure it seemed to be having antibody responses and not being toxic, or whatever. And they just took their best shot and picked one protocol and ran with that for the phase III trial.

And then, once it worked, and then in subsequent trials, then they're starting to test different things. And then they have a much better sense of power. And also, the thing works, so they cam-- but the very first trial of, is this thing going to work, they were like, we don't know what the effect size is going to be, so we're just going to max out treatment versus control and just take our best shot based on the preliminary small-scale studies of what we think is going to work. And that's what we're going to do for the big trial and not, in the very first trial, go for multiple different strategies.

As I said, once it works, then you actually have much more information. You know how big the response is. You can plot these things much more carefully. You have much tighter priors. But for that first efficacy trial, it was just like treatment versus control.

AUDIENCE: [INAUDIBLE]?

BEN OLKEN: Yeah.

AUDIENCE: If you're doing-- I guess if you go for one big group, you're not able to ask questions about dosage for the trials. But you can ask questions about other sorts of-- I mean, I guess dosage is different [INAUDIBLE] in clinic. But my point is, you can still do more than you are. I'm trying to think like what subsequent analysis can you do beyond the simple difference of means?

BEN OLKEN: What do you mean?

AUDIENCE: Like-- You lose the ability to do doses, [INAUDIBLE].

BEN OLKEN: Yes, yes. You cannot-- exactly, you can't do dose-response.

AUDIENCE: In a sentence, like what's your thought [INAUDIBLE]?

BEN OLKEN: I mean, if you're trading off-- if you don't do dose-response, you can't do dose-response. That's what you lose. I think that that's the downside. No, so what's the gain? The gain is you're much more likely to pick up an effect by, like, a lot. That's the main.

If you're just trying to say, does this thing have an effect, and the first-order estimate of what is the slope of that thing, the first-order things are going to get a lot more power on identifying what's going on and rejecting the null of 0.

AUDIENCE: Wouldn't you do heterogeneity based on people that you think [INAUDIBLE] might need a larger or smaller dose?

BEN OLKEN: You could if you-- yeah, so if you think of the clinical trial example, if you thought that if you-- maybe you thought dose was proportional to weight. For example, you could do that if you [INAUDIBLE].

AUDIENCE: So that's OK [INAUDIBLE].

BEN OLKEN: But it's not randomized heterogeneity. So it's correlated with other things. Maybe like people who are, whatever, have different immune system, or whoever the equivalent is here. OK, all right.

OK, so what do they find? So they find that basically, in the peak season, wages increase by 5%. Employment declines by 21%. And in the slack season, basically, wages are unchanged. So the peak season looks like a competitive labor market with elasticity of labor demand of around minus 4. And the slack season looks like you have rationing.

OK, so both of those-- in some sense, that paper and the previous paper go together to basically say, look, it seems like we have something not clearing in the labor market. The [INAUDIBLE] was saying that. That experimental paper was saying that. There's something in these global labor markets which is preventing labor market clearing.

So the question is, why? What's going on? And why is this happening?

So there's a couple of recent papers who have started to investigate this. So one is another paper by a similar set of co-authors which runs an experiment to get at this. And here's what they do. They go to Indian villages, and they offer spot jobs, They vary the wage of the jobs, and they vary whether the wage offer is observable or not.

And the way they do this is pretty subtle. They basically do it in the context of survey. They do a baseline survey, and they say, by the, way we're offering this job. Here's the wage. Would you like to do it?

And in some of them, they make that offer on the front porch of the person's house. So people who are-- and if any of you have ever done a survey in a developing country context, it's often pretty common for other people to crowd around and listen in to what's going on. It's pretty exciting. Some person's here doing a survey. Let's listen in.

So they do that on the front porch. Or they say, let's please do the same thing, but we're going to do it inside your house so it's in private. So it's exactly the same, but just varying whether it's on your front porch or in private. So those are the three treatments And so what would you what you expect from this?

AUDIENCE: You're more likely to accept a lower grade in private than--

BEN OLKEN: Why?

AUDIENCE: Pride. Pride, essentially.

BEN OLKEN: OK, so any other theories of why? Yeah?

AUDIENCE: I guess there's also-- if there's a-- I guess the analogy I have in my mind is like a scab, basically accepting a lower wage when people are-- if there's this agreed upon [INAUDIBLE] we're not going to work for [INAUDIBLE].

BEN OLKEN: Right, and in fact, the very first draft of this paper was titled "Scabs." And then they changed title, in part to acknowledge that there were other things, I think, going on.

But yeah, so one theory is like this external social norm theory-- that basically, there's a lot of peer pressure that's going on to prevent, to enforce these higher wages. And so any individual is going to get social sanctions if they accept the lower wage. And the other is your theory, is like this is an internal thing and I'm not going to do it.

But the internal one would have to be only my pride vis a vis the other people. It's not my internal pride. It's that people are watching me.

AUDIENCE: Yeah, I think this is maybe a little bit of articulation of the dynamic that I had.

BEN OLKEN: What, sorry?

AUDIENCE: This is maybe a slightly better articulation of the dynamic that I had in mind.

BEN OLKEN: Got it. So there's something going, right? And so by the way, can you imagine, are there cases where it would be optimal, maybe, for the village to have these norms?

AUDIENCE: [INAUDIBLE] a monopsony, then [INAUDIBLE]?

BEN OLKEN: Yeah, exactly, right? So if there's a monopsony power over the local labor market power, it's like exactly, it's kind of bargaining with the firm. And/or in general, if there's any kind of rents, we might want to extract from those rents, and so we want to keep the wage artificially high. And that might mean, whenever we keep the wage artificially high, there's a little bit of surplus. Individuals might want to deviate. But we're all trying to keep the wage high to extract rents together. So it's kind of a collective action model.

OK, and here's the result. It's super simple. The result is-- this is the main result. So this is probability of taking up the wage at the going market wage. At 10% less the going market wage, when in public-- I'm sorry, in Private-- and 10% less than the market wage in public. So it's pretty stark. So basically, almost nobody is willing to-- and I'm sorry, this is in this the high unemployment periods and low unemployment periods.

In the high unemployment period, basically, lots of people want this job, because there's a lot-- that thing, the labor market, is really-- that gap is kind of large. In the low unemployment periods, the gap is kind of small. And so maybe, actually, we don't care that much.

So in these periods over here, that's the one consistent with what [INAUDIBLE] was suggesting, which is basically we're in a market where the wage is too high. We'd all like to get jobs, so there's some failure of market clearing. And so people are like, if you can do it in private, you're happy to take it, but you don't want to do it in public because you're worried about retaliation or whatever.

And in fact, one thing that's kind of nice is-- this is too small for you to read, but you can see in your handout-- they ask people what's going on, and there are various-- what you would do, sorry, to a worker who took one of these jobs or whatever? You would try to convince them not to. You would impose social sanctions, et cetera, et cetera. You'd try to exclude them from other labor market opportunities, et cetera, or what? So in general, it looks like people-- there are these public punishments for breaking these wage norms.

And so that's the beginning of an explanation for how we might be getting these labor market failures of clearing. We have these prevailing wage norms. Those wages move around, and then they stay high, potentially as a way of getting some rents. Is it efficient? It's actually-- to that question, I'm not sure we have a clear answer from all these papers of whether or not this is allocatively efficient for these villagers or not, whether the additional rents they're getting exceed the unemployment. You may be able to calculate that from these papers. I don't remember them actually doing that, but that's something to think about as well.

So it looks like, actually-- so stepping way back, it looks like these labor markets are actually-- what's going on is kind of nuanced, having these failures of labor market clearing. You have these social sanctions that are potentially impeding that. OK, questions?

So the final set of topics on labor supply that I want to talk about are-- I won't get to finish them, but I'll start them-- is some papers about, I would say, poverty and behavioral issues in labor supply. So which is to say, so one issue that we talked about in general was having these labor market frictions. There may be other reasons, there may be other things that are affecting these rural labor markets due to the fact that these people are really poor. And so how can we think about those issues?

And so I want to explore four different related issues which will take us mostly the next-- I will do this one today, and then the rest of them we'll do next time. So the first is like, as these question are, how does poverty affect our labor supply decisions?

One argument is that basic consumption needs basically change the elasticity of labor supply, and that can have aggregate implications. A second is a behavioral channel where poverty and needing to think about how you're going to meet your daily poverty needs change your productivity directly. We'll talk about identity and labor supply issues. And then, actually, I'll talk about a new paper on the cognitive benefits of work.

OK, so the first paper I want to mention about poverty and labor supply is a paper by Seema Jayachandran, and the idea here is very simple. Her idea is that basically, workers have minimum consumption needs. And if you're really poor, or you have inability to smooth shocks across periods, or to borrow, or to do anything to smooth, that's going to make your labor supply less elastic.

And to be less elastic, if there's less access to credit, or less ability to migrate, or whatever, and less elastic wages are going to mean that there's a pecuniary externality from you to other people. So in particular, if we contrast a labor market where labor supply is elastic and looks like this-- that's elastic-- to one that looks like this, labor supply inelastic, imagine what happens when we have a negative labor supply shock.

When we have a negative labor supply shock, if labor supply is inelastic, the wage is going to fall more than if it's elastic. And what that means is that there's a negative externality-- a pecuniary externality, which is an externality that happens through prices-- of having a lot of poor people around you. So even if we hold my own ability to smooth shocks constant, if I'm surrounded by people who are unable to smooth shocks, they're going to be more inelastic, and that means the aggregate wage that I'm facing is going to fall more in these negative periods. Is that clear?

OK, actually, I'm out of time now. So I'll pick up by just starting with the empirics of this next time, and then we'll go through these other papers on labor supply. And I'm going to try, because I'm running a little bit behind, to finish up the labor supply-- sorry, the labor lecture next class so that we can have spend next week on credit.