

[SQUEAKING]

[RUSTLING]

[CLICKING]

**BENJAMIN
OLKEN:**

So stepping back-- what we were talking about on the tax section was we were talking about two different approaches to how you can think about solving the challenge of the infor-- OK, of the information challenges in the developing-- in developing world and the tax structure. One of them was through thinking about how do you design a tax code that allows you to get information revelation. And the other is how do you actually maybe through the tax administration system just get better data on people. And those were the two approaches.

So we talked last time in the examples about the VAT about how do you think about that through changing tax administration. The other thing I want to talk about now is, how do you think about that-- how do you think about whether the tax administration margin is kind of interesting or is relevant? And what are some of the challenges in doing that. So I want to talk about two papers on this, and then I'll more or less wrap this up and switch to the next topic.

OK. So this is a paper actually-- this is also a paper that-- these are papers that I've worked on, so for better or for worse. So this actually is a recent paper-- it's actually now forthcoming-- on tax administration versus tax rates, to understand just at some very broad level is tax administration a really big deal or not?

And so-- and we want to do that in a way in which we can benchmark it or compare it to a traditional tax rate change. So the typical thing you do if you're reading about debates about taxes, the typical thing we think about doing is raise the rate, lower the rate, change the base, those kinds of things. Can we compare just a straightforward change in improvement in tax administration to one of these more traditional things like changing the tax rates?

And so we're going to do this in Indonesia studying two reforms. These are corporate taxpayers. One was a corporate-- a taxpayer administration reform. And here, the idea is that corporate taxes tend to be very skewed.

So people on the-- basically, the firm distribution is skewed. So there's a small number of really large firms that pay most of the tax. And therefore, you may want to invest your tax administration kind of on those firms.

Lots of countries do it for the very biggest taxpayers. Indonesia, when we're studying this, was kind of rolling out this same idea to medium size taxpayers. So, in particular, they rolled out these medium sized tax offices to serve the largest 330 taxpayers in each region. And we're going to study the impact on firms when those were created. And we're going to find really big impacts.

The second thing is the tax rate reform. And basically, here, the identification is going to come from the fact that they-- I don't think I brought the graph actually. But they used to sort of-- actually, I may not focus on this one. I want to focus on the tax administration reform.

But you can read the paper for interested. They had a switch from a tax rate that was based on your taxable income to a tax rate that was based on your gross income. And so that gives you identification for the change in tax rates. OK.

So I want to focus on the tax administration reform. So the typical diff-in-diff assumption that you guys have seen in multiple papers throughout the semester. And I'm sure Ed talked about it in a diff-in-diff lecture. Did you-- how much did you go into matching-- propensity score matching? Is that-- not at all. Great. So I'm going to talk about that right now, is the control groups-- the control group is on the same trends as the treatment group, other than the treatment. That's the assumption.

So that's-- in this context, we're going to study firms that-- so what's the identification? The identification is in each of these-- or the thing we're studying is in each of these regions they create a special office for large firms. And they move the kind of biggest firms into this special tax office. And we want to study the impact of that.

The problem is that the firms that are moving in are large firms. They literally take the largest firms in each area and move them in. So if you saw this, and you said, what would be the-- they're going to treat the largest firms in each area, what would you think of in terms of identification strategy?

If you were rummaging around like the Indonesian tax office and came about this thing where they're creating a tax office for the largest firms in each region, what would you be thinking about for identification? Yeah?

AUDIENCE: Well, one thing you could do is if they had a particular cut-off for the size of the office to be [INAUDIBLE].

BENJAMIN OLKEN: Exactly. Right. You would think you would do an RD, and that's what we thought, too. Unfortunately-- so we thought we would do an RD. Unfortunately, we couldn't do that because, basically, the RD was done on some combination of different other variables-- gross income, which is your revenue, tax payments, maybe some other variables in some formula. They had some Excel sheet in each of the 19 different regions. And the Excel sheets were lost to the recycle bin of time.

So unfortunately-- and I say, despite our many, many efforts, we got graphs where we could try to replicate the formula. We got graphs that look like this. Probability that you were assigned to the new thing looks like this. This is firm size. And it was beautifully continuous and there were no discontinuities.

And I think that that's because-- so clearly, we had the right variables. They were definitely doing it based on this assumption. But because we couldn't reproduce the exact thing that was used, we couldn't do an RD. And I have to say this is not the first time this has happened to me that-- so we tried, but ultimately, it kind of looked like that.

So what we get instead is we used-- yeah?

AUDIENCE: [INAUDIBLE] cannot produce exactly [INAUDIBLE] RD.

BENJAMIN OLKEN: But if the graph looks like this, there's no discontinuity. It's all continuous.

OLKEN:

AUDIENCE: I see.

BENJAMIN I think that basically what happened is because we didn't know the formula well-- I think that the issue is we didn't know the underlying running variable. Nor do we know-- I think what happened is a couple of things.

Number one, we couldn't reproduce the exact underlining running variable. For example, maybe they used an older version of the tax data that wasn't exactly right. Maybe they had other things going on. We couldn't find RD, basically.

AUDIENCE: Propensity matching?

BENJAMIN Huh?

OLKEN:

AUDIENCE: What about propensity matching?

BENJAMIN That's what we're going to do.

OLKEN:

AUDIENCE: OK.

BENJAMIN Yeah. That's the answer. Yes. That's where we're going. So instead, we're going to use propensity matching-- match differences-in-differences. And so this is what I wanted to talk about.

OLKEN:

So instead what we're going to do, is we're going to say, look, it's pretty obvious with this kind of design that if I look at, say, these for-- over here, everyone is treated. Over here, nobody is treated. So it's kind of hopeless there to find kind of reasonable comparison sets.

But in this group kind of in the middle, your probability of being treated is somewhere in between-- is in some intermediate range. And therefore, we may be able to compare similar firms that just-- some of them happened to get treated, and some of them didn't-- happened not to get treated. And that's kind of the idea of propensity score matching.

So what do you do? You basic-- so how do you do propensity score matching? OK? So-- or in general for matching. There's a couple of steps.

The first step is exactly this. You have to restrict-- you take your sample, and you want to restrict to the area of which-- where there's called common support, where basically there's going to be reasonable fractions of people who are both treated and not treated. Or in this case, firms that are treated and not treated, so that you have reasonable matching sets.

So otherwise, you end up-- there's nobody in the control group who's going to be a good comparison here, and no one in the treatment group who's going to be a comparison to this firm here. So you just cut that sample out. And you should recognize that the late you're estimating-- the [INAUDIBLE] you're estimating is for that kind of medium set.

So in our case, for example, we're going to be estimating, an effect, for the smaller firms that go into the medium tax office because the larger firms that go into the medium tax office, they're just getting treated with close to probability 1. So in our case, actually, it's a little more complicated because we're actually doing this on two different variables. We have-- because we know that they matched on your revenue, which is your gross income and your taxes paid. So we're going to do it on both. But it looks kind of like this.

So this is the distribution of-- this is the distribution of gross income in the baseline year. This is the year that you used for assignment. And you can see that the treated groups have much-- are much bigger than the control groups.

And so what we're going to do is we're going to say, I think, do, say, 97th and 2.5th percentile. What does that mean? We'll say, well, when I get down over here where there's only 2.5% of the distribution here, we'll make a cut-off there in the treatment group.

And then we'll go the other way around, and we'll get through the same cut-off here in the control-- in the 97th and a half in the control group. And so that gets us this intermediate set, where we have reasonably-- like a reasonable number of firms from both distributions.

Now, one thing you'll notice if you look at this graph, of course, counterintuitively, is that over down on this part of the distribution you have way more controls than treatments. And over on this part of the distribution, you have way more treatments than controls. So the next step, after you made this common support restriction, is you just deal with that fact and make these two groups look similar by weighting-- by weighting.

So as I mentioned, we actually do this in two dimensions. We do in this dimension. This is on taxes paid. It looks kind of-- similar idea. The firms assigned to the treatment group are paying a lot more taxes. So you do these restrictions.

The next step is you use the preperiod data to reweight the treated-- in control groups-- so that the weighted distributions look the same. So instead of having this-- because if I just took this regression on this sample without doing any weights, it still would be mostly bigger firms in the treatment, and mostly smaller firms in the control. But you can easily see from this picture how I could just reweight them to make them look kind of the same. And that's really all propensity score matching is going to do.

So how do you do that? So how do you actually compute the weight? So there's a couple of different options. So one is what's called a propensity score. So you know if the functional form of the assignment rule, you can just estimate it, and you can use it to calculate the weights.

So for example, say, it was like a probit-- like probability that you're in the treatment group is some $X \prime \beta$, where the X is, and you know this functional form, you can run this auxiliary regression, predict the probability you're in there as a function of these multiple different variables. And then you can use weights $1/P$ for the treatment unit, and $1/(1-P)$ for the control group-- control units using the predicted values. And that will just make the weights balanced. So that's option one.

The other option-- and this is kind of a newer approach-- is a more nonparametric approach, which just solves for the weights directly. And basically-- and this is, look, if I don't know the functional form, what I'm going to do is I'm going to take the-- I'm going to look for a set of weights in the control group, so that I balance-- so that I just match whatever moments I want to match in the treatment group on a variety of different characteristics.

And so this-- there are a couple of different versions of this. This method by Hainmueller is the one that we use in the paper. The idea of it basically is to say, look, let me just solve for a set of weights that matches on whatever things I want it to match on that are going to minimally deviate from uniform weights.

So, of course, many different combinations. If you have more observations than moments you want to match on, there are many different ways of doing it. It chooses the one that is like least-- has the least deviation from uniform weights.

And there's some arguments actually that in cases where you don't actually know the propensity score, or the propensity score formula isn't inexact, these nonparametric methods can actually do a better job kind of matching things. And that was our experience, actually, in doing this, which is why we went for the-- It doesn't really matter. You can-- all the results in the paper we'll go through if we do the regular propensity score, but I just wanted to mention this, as well, because it's a newer-- and there's a nice-- it's a newer method and there's a nice canned data package that will do this for you.

So in any case, first step-- common support; second step-- estimated propensity score and rebalance. Third thing is then you can just estimate your difference-in-difference model on the reweighted data. And that's what we're going to do.

You might have said, well, why do I even care? Why even-- one question you may have is why even bother? Why even do a different-- why even bother with this reweighting? Why not just do a regular difference-in-difference?

And I think the reason is that you might think that these different-- there's two-- these different groups might be on of different trends or whatever. And so you really want to have comparable to comparable to make sure you're balancing the trends properly. Any questions? That's mostly what I wanted to say about propensity scores, but I just wanted to make-- to throw it in there. Yeah?

AUDIENCE: [INAUDIBLE] and weighting methods compared to, for example, a exact [INAUDIBLE] placement that match a treatment unit with one to three for every single control unit.

BENJAMIN Oh. Oh, so like exact matching?

OLKEN:

AUDIENCE: Yeah. So it's [INAUDIBLE].

BENJAMIN So I don't have-- it's a good question. I don't have a-- let me think for a minute and see if I have a good answer for you. So another-- so you're saying another option is for every option in the treatment group I find the closest match in the control group, and I just put them together, and put it in a matched fixed effect.

That's a good question, actually. I don't have a good answer for why-- whether that would be-- well, if you have more-- it's a good question. I'm not sure I have the right answer. I think if you have more treatment-- if you have more of one-- if you have unbalanced numbers of observations, and you end up having to reuse the same ones multiple times, it may be inefficient compared to this.

That's my best guess. But I don't know the answer, actually. Does anyone else have any ideas? I don't know the answer. I do-- any-- it's a good question you-- I don't know the answer.

So this just shows you the raw data. This is like total taxes paid. This is the year when the thing is partially-- the thing takes two years to fully turn on. It's partially turned on in 2006-- fully turned on in 2007. And what you can see is you can see that there-- this is just the raw data weighted.

So you can see that the pretrends looked kind of nice and balanced between the-- once you-- in these two weighted samples. And then clearly, something pretty dramatic is happening when we turn on the medium tax office. They seem to be paying a lot more taxes. And this shows you too a normal diff-in-diff estimation form.

So, in fact, actually, their taxes go up by a lot. They go up by more than 100%. And the government raised a lot of revenue. And comparing this to tax rate changes, it's staggeringly large and equivalent to raising taxes on-- raising the tax rate by, I think, 6% or 8% on all taxpayers.

So the substantive point is that tax administration can make a really big difference. And the thing I want to talk about methodologically is, how do you think about doing these matched difference-in-differences? Any questions?

OK. I'm going to skip this. Oh, actually, no. This I wanted to say, too. OK, sorry. So why might this make such a difference? So one thing that's kind of striking and surprising, I think, is that it continues to actually go up over time.

So you might have thought that once tax-- once you put in a new tax enforcement regime there will be like a direct effect where you'll get more tax-- collect more taxes. And then firms will start to figure out how to adapt. So you might have predicted it would go kind of like up and then down over time as firms start to adapt. That doesn't seem to be the case. In other measures, it seems to always be going up.

So one answer is it just takes some time to learn stuff. And the tax-- and these tax inspectors are figuring stuff out. That could be. But one thing that we think is kind of interesting is that there may be also a change in size-- a reduction in size-dependent taxation. So what do I mean by this?

So imagine you have a simple model. This is actually not the model in the final version of paper. This is a simpler model that was in a previous version of the paper, but I think it just makes the point very clearly. So imagine you're a firm that you-- who can both make decisions about how much to pay, and they can maybe evade some taxes, e.

They're taxed on-- this is a model where there's distortion. So they're supposed to pay taxes here. They pay taxes on this. But they can kind of evade some revenue here. So the firm's production will be distorted by the taxes. Because the tax is not-- if the tax is a profit tax, it doesn't distort your revenue. But if it's-- but since there's this wedge here it may be-- there may be some distortions.

What I wanted to say is suppose that in addition to that, the cost of evasion now-- and you have some cost of tax evasion, where α is kind of the level of tax enforcement. Suppose you have the case where the level of tax enforcement is dependent on firm size. So it's pretty intuitive that kind of in a-- as I mentioned before, they're going after the largest firms because they're biggest.

If that's true, then that's going to create an additional-- there'll be a general distortion of taxation. But there'll be this additional distortion from taxation coming from the fact that as you grow larger you're more likely to get taxed by the tax authority.

So if that's true, once-- so then the question is, once you've been put into this medium-sized tax organization, does that change not just the level of enforcement, but also does it change this slope of enforcement with respect to firm size? And if so, that could potentially explain part of what's going on.

And so we actually also looked at the relationship between the probability of enforcement as a function of firm size-- a function of firm size doesn't have that change when you move to the MTO. So for the regular people in the small tax office, the probability of being audited, for example, is very strongly increasing in firm size.

Once you're moved to the medium tax office, it's flat. Basically, they have enough resources to audit everybody, essentially. Or to audit not everybody, but with high probability-- with much higher probability.

And this is true on a couple of different dimensions that we can measure. And the reason I wanted to mention that is I think there's this issue where-- I think there could potentially be an issue where in a lot of these developing country contexts if the tax department is only starting to pay attention to people who are kind of larger, that that in itself can create distortions. And so I just wanted to mention that.

So in this case, we actually have the opposite happen. Once they got sufficient-- once they were put in the MTO, at that point, that had already-- that was kind of sunk. They were already getting the high tax treatment regardless. So at that point, they may as well sort of-- they lose that-- that incentive not to grow too big kind of disappears. But more generally, that point about size-dependent taxation creating distortions for a firm size is one I just wanted to mention. That's all I want to say about this paper. Yeah, Aaron?

AUDIENCE: I'm forgetting the paper, but it was there-- when these MTOs were created was there displacement of people who were previously working in the smaller taxpayer offices? And was it-- were there effects on revenue from that displacement?

BENJAMIN OLKEN: I think it was small. Because the number of staff they needed for the MTOs was pretty small. And so I think that they didn't really have to-- and they could also promote other people and so on and so forth. Yeah.

So if tax administration kind of matters, then there's now a whole series of questions for how do we actually do it? And what are the particular kind of like-- and that's true kind of everywhere. But I guess a question for a development classes is are there particular development-related issues in doing this-- in changing tax administration? So I wanted to mention one last paper, which is a paper of mine, on this issue.

So one thing people have thought about doing to improve tax administration is to just give incentives to tax officers. The whole idea is we need these people to go out and collect taxes. Like in this low-information environment, we rely on the people. Maybe we need to give them incentives to do this.

So, in fact, giving high-powered incentives to tax collectors is a very, very old idea. In fact, it goes back to the Roman Empire, or in the French-- prerevolution France, they had people who are called tax farmers who basically paid a fixed fee to the King, and they got to keep all the taxes on the margin. So it was like 100% incentive. This was very unpopular.

So, for example, the tax farmers were beheaded during the French Revolution. So that was not-- people didn't like this very much. And why didn't they like it? Well, because these people were super zealous in terms of collecting taxes because they got to keep it all on the margin.

So you don't want to go too far. And, in fact, if you were to-- in fact, if you look at the-- you all live in Massachusetts. If you look at the Massachusetts Department of Revenue's website, there's a taxpayer bill of rights. And one of the things that Massachusetts states is that they will not under any circumstances create any incentives or even informal targets for revenue for any of their tax staff, in part because tax payers don't like this.

So I think-- so we wanted to understand in a more modern context-- on the other hand, though, if you to a lot of these developing country contexts where the information problems are really severe, and you have to motivate these people to go collect taxes, maybe the state should think about revisiting this trade-off. So how do we think about this?

So we did this by doing a randomized experiment in Pakistan where we were working with property taxes. And property taxes are collected by a team of tax collectors. There's a team of three tax staff who-- and together, if they were put in the treatment group, they receive between 20% to 40% of all the revenue collected above some historical benchmark. So you can think of each of them as getting a 10% kind of payment on the margin.

So any thoughts on what would likely to happen if you did this treatment? So you see why they would do it, right? They would do it because like-- so let me back up. Why do they need this incentive in Pakistan and not in Brookline, for example, where I live?

So in Brookline, they have very good information about all the houses. And everyone's required to build a-- to have a building permit. The building permit is automatically linked to the tax people. If you change your house, the building people know. The town knows about it through the building permits. They can send the tax people out. They sort of-- and it's all kind of linked together.

So they do rely on tax inspectors, but they have much more data underlying them. And also, the penal-- if the tax inspector were to come to my house and offer to adjust my taxes in return for a cash payment, the probability that would end up in detection, and they would end up going to jail is probably reasonably high. And so that probably doesn't happen very often. So that's-- but in Pakistan, that might not be true right.

So what do you think is going to happen in this-- so that's why they might want to do incentives. So what do you think's going to happen with incentives? Yeah, Jenny.

AUDIENCE: They exert more effort so there's more total tax revenue.

BENJAMIN OLKEN: OK. Right. So incentives matter. They should exert more effort and collect more tax revenue. Paulo?

AUDIENCE: If there is like a collusive evasion between the tax staff and the person being taxed, you might not observe any change in revenue, and you just have a change in how they bargain for, like--

BENJAMIN OLKEN: Exactly. You've seen this paper before in a previ-- yes. But yes. Exactly. So the other thing is-- so Jenny is totally right. They could collect more tax revenue.

On the other hand, if they're bargaining with-- if they're potentially corrupt and kind of bargaining with the taxpayer, this could just-- all this could do is kind of change the underlying bargaining game. And so that's what we want to illustrate in this paper is how do we think about those two different things.

So we set up in the paper a very simple kind of Nash bargaining model with equal bargaining weights between a taxpayer, P , and a tax collector, C , to collude and reduce their tax liabilities. What does this mean?

This means I'm the taxpayer. I go to, say, Ashana, and I say-- I'm a tax collector. I go to you and you say, look, your house is worth \$1,000-- your tax is \$1,000. But maybe you and I can work something out. And in exchange for a payment, maybe I can reduce your official tax liability. So how does that kind of bargaining work?

So imagine that τ^* is the true amount of tax, which is the same for everyone. We can instead negotiate to pay a bribe, b , and report less tax, τ , to the government. Together we can do that.

The taxpayer's utility is going to be minus τ -- they don't like paying taxes-- minus b -- they don't like paying bribes. Taxes and bribes are going to enter the same here. Those are just money kind of flowing out the door.

And we're going to give them an additional cost like α times $\tau^* - \tau$, which is the-- like a disutility cost from paying bribes, or from underpaying their taxes. Sorry. And maybe there's the different-- that α is heterogeneous among taxpayers.

The tax collector's utility is going to look like this. It's going to have-- they're going to get an incentive payment, r , times the taxes they pay, $r\tau$. There also may be some bribes. They like money. And they let-- money is treated equally here. They like incentives. They like bribes.

But they also get a disutility cost from underpay-- underreporting their true taxes. And you can think of that as like some chance they get caught, or something-- maybe they want to be honest or whatever. And this simple model, everything is linear. There's a version in the paper in the appendix that makes this quadratic, and the linear version is easy.

So what happens in this model? So the Nash bargaining solution says, look, we're going to come to an agreement that's going to max-- that's going to be efficient for the two of us-- for me and the taxpayer. What does that mean? That means we're going to maximize our joint surplus. And then we have to share the surplus somehow.

And different bargaining models get predictions for how we should share the surplus. The Nash bargaining model says, look, we're each going to get our outside option. And then we're going to get some fixed share of the surplus.

So what's-- so the joint surplus from agreement is we add up the taxpayer's utility and the tax collector's utility. We just add them up. And so I can rewrite terms. And I have this $1 - r - \alpha - \beta$ times $\tau^* - \tau$. τ^* is the true tax liability. τ is the amount we report.

So what you can see is that as long as $1 - r - \alpha - \beta$ is greater than zero, then we're going to want to under-- we're going to want to not report any taxes. We'll maximize-- I mean, so you can't have negative taxes here, so we'll set τ equal to zero. And if this thing was less than zero, then we might as well tell the truth and pay $\tau = \tau^*$.

So that tells us something interesting. That says, look, for some combinations of taxpayers and tax-- taxpayers and tax inspectors, they'd be in the collusive equilibrium if r is equal to 0. If $\alpha + \beta$ is less than 1, they'll be in the collusive equilibrium if r is less than 0.

But as I start to increase r , then some of them are going to be like, you know what? It's not worth it anymore. We're going to switch from the collusive equilibrium to the noncollusive equilibrium and just pay taxes.

And intuitively, what's going on here-- what's going on is I'm getting some utility losses from this collusion thing. And the government is sweetening the deal by throwing in this incentive. And it's just-- at some point, it gets to the case where the incentive plus the two utility losses it makes the whole thing not worth it-- the foregone incentive, right?

So that says, look, so one thing that's going to happen in this model is I'm going to stop-- some people will stop paying-- stop colluding and pay more taxes. But the other thing that will happen is that-- that's this case over here. The other thing that will happen, though, is that if I stay colluding, the amount of bribes I have to pay is going to go up. Why is that?

So that's what-- that's-- you can see that in this expression here, that's just-- so how do you solve this expression here? You basically solve for this by saying, what is the bribe?

The bribe is a way of transferring some share of the surplus from the taxpayer to the tax inspector in this model. Because if we didn't have any payment, then the taxpayer would get all the surplus. But they had to share the surplus. So they share a surplus by paying through the bribe.

And well, what happened? Well, remember that each person gets a share of their outside option-- they get their outside options, plus a share of the surplus. So now the tax inspector's outside option has gone up. Because if they don't collude, the tax inspectors getting the incentive from the government.

So now the taxpayer has to compensate the tax inspector for their forgone incentive, plus you have a surplus. So that means that the amount they have to pay, the bribe, is going to be increasing in the incentive.

And that-- if you just-- if you-- the math gets a little ugly with the alphas and betas. But if you just take the outside option and a share, a γ , of the surplus, you get this expression here, which is increasing in r -- the incentive.

So that's kind of the-- and you can do a simple version of this. Like, imagine if there was no incentive-- imagine there was no α , no β . If it was a simple version, say, bargaining over \$1, we split the dollar and pay \$0.50.

Now imagine the tax inspector was getting, say, \$0.30 in incentive-- a 30% incentive. The surplus is now \$0.70. Half the surplus is \$0.35. What's the bribe you have to pay? It's the \$0.30 of foregone incentive plus the \$0.35 of half the surplus. The bribe goes to \$0.50 to \$0.65. And that's kind of the intuition. Is that-- I'm going through this kind of fast, but is that clear? Yes? OK.

So what that suggests is that as we-- I've basically said that. I said all this, right? So we have two different things going on. We have an equilibrium selection effect and an equilibrium bribe amount going up. And that suggests if we look in the data, we're going to want to look for both of those two different things. And that's what we do.

So we chose property tax. So the first thing is we get tax revenue. That goes-- this is actual tax revenue from administrative data. This is in log. This goes up by about 9% in the first year, 9.4% in the second year. So the government really is collecting more tax revenue.

But to look for this kind of other thing going on, we did a survey, and we looked both at the general population and at the small subset of people whose actual official taxes-- tax liabilities change. Let's just call this re-assessed sample.

And for the general population, they report in treatment areas a higher going bribery. What does that mean? We ask the question, what is the typical rate people in your area might have to pay to get your taxes adjusted? That was-- we didn't ask what you did. But I said, what's the going rate in your area? And that's going up by about a third-- by about 30% roughly in the treatment area, as well as the control areas.

How often people in your area tend to pay bribe payments also goes up. We have a subjective perception of corruption. That doesn't actually change, but both these things go up.

We then also asked, how is this different for this group that was actually reassessed or seeing their taxes changed? So these people pay a lot more in tax. So their taxes are going way up. And they don't report any-- this increase in bribes.

So this is like the reassessed people. And the reassessed in the treatment areas is almost like the opposite of this. So they're not actually seeing any changes in bribes, nor are they seeing any changes in bribe frequencies. They're not actually seeing it go down on net, but they're not seeing these big increases.

So that's consistent with the idea that there's two different things going on. One group is seeing the high change in-- the high change in bribes, and the other one is seeing-- like, paying a lot more taxes. I think that's potentially consistent with the model.

We also look at whether they're being overtaxed. They're not. They're basically moving from being undertaxed to being correctly taxed. And we'll just leave that at that-- more of that in the paper.

So what's the point of going through this? This is one example, I think, of how we can think about the-- if you're moving to this question of tax administration as a solution to some of the tax problems in these developing country contexts, then you have to think about the realities of how you do this. And this is some of the challenges that you may actually face. And there's a lot more stuff you could do in this area. I just wanted to leave it leave it at that.

I think I'll stop here and switch to the labor section, there's unless there's any other questions. OK. Actually, let me just give you-- let me just give you one slide on this, which is-- the last topic that I just wanted to mention is this idea of informal taxation.

So in the same way that there are ideas about informal insurance or other informal mechanisms, like supplementing when the formal systems break down, that people are insuring each other or doing other things. I think the same ideas actually also apply to taxation. And we have a paper that discusses a few of these phenomenon and tries to think about how do we think about some of these-- what is this phenomenon, and how do we think about it as a tax perspective?

And basically, there are a lot of cases where in local village context, for example, there are these phenomena where basically people are kind of contributing to local public goods rather-- for things that-- in the US, for example, we would do through the tax system.

So, for example, one of-- I've been in these villages in Indonesia where we need to pave a road. In Brookline when we need to pave a road they set our taxes, and they raise revenue from the taxes, and they hire a road-paving company to pave the road.

In an Indonesian village what they might do is they might say, look, we need to pave the road. So we're all going to get together and have a meeting. And we're going to decide how much we would like you to contribute to the road. And what are you all willing to contribute?

And Sal is willing to contribute two days of labor. And Aaron is willing to contribute a gallon of-- a tank of asphalt, and so on and so forth. And we'll get all these voluntary contributions and use that together to build the road.

At some level, these are like the same thing. We're both collect-- gathering a bunch of contributions from the population to fund some public good. But in one case, we're doing it through the tax system, which is both formal and kind of has these very strong formal punishments for refusal to comply. Or we could do it through this informal system.

And so in this paper we think about, what is this informal system? And what does it look like? And how progressive is it? And what's going on? And this is an area where I think we've done a little work through this paper. But I think that understanding a little more of these voluntary systems providing public goods is an interesting area for future work and development. So I just wanted to leave it out there at that. OK?

So we have some stylized facts in this paper, which I'm not going to go through now. But if you're interested, you could look at that. OK. Other questions? OK. Sorry if I was a little bit rushing, but I wanted to make sure we get on to the labor topics. OK.

All right-- switching gears. So the next set of things I wanted to talk about are labor markets and development. And this is actually another area where there's been a lot of-- I think a lot of movement in this literature. Even though the paper I assigned you to read is an older one, there's been a lot of movement in this literature over the past five years or so. And so you'll see that even though some of this I'm going to start with today is a bit older, some of the stuff I'm going to talk about over the next-- in the next lecture or two is pretty recent.

So here's the overall plan. So a theme that we're going to look at is, how do we think about efficiency in labor markets in developing contexts? I'm going to start with motivation from this idea of the, quote, "surplus labor hypothesis," which is a very old idea, and how do we think about that relating to the separation paper.

I'm going to talk about labor supply. We're going to talk about the separation test as a test for frictionless labor markets. Then next class we're going to talk about nominal rigidities, behavioral issues, and other challenges in labor supply and some frictions in labor demand. And then I'm going to talk a bit about urban labor markets and how do we think about some of those issues as well. So that's kind of the overall plan. All right.

So going even further back than the reading that I gave you for this, in some sense the-- a classic reading in development is this paper by Arthur Lewis from 1954. And Arthur Lewis basically argued that there was, quote, "surplus labor" in the countryside, in developing countries. And, in particular, he said, roughly speaking, "about 25% of the labor has zero marginal value." There's way more people than we need for the tasks that are needed to be done. And they're just kind of like-- just really on the kind of the flat part of the production function.

And his claim, therefore-- this is actually an important development claim-- was he said, look, we can increase aggregate productivity if we can just move some of that surplus labor that's not doing anything useful from the countryside to the cities. We can do that. And he said, we can do it without even decreasing agricultural output because they have a zero marginal product. OK.

So that would mean that either the marginal product of labor is zero maybe because the agriculture production is Leontief. In general, we tend to think-- we don't tend to draw a production functions-- we tend to draw production functions as concave, not actually flat. But maybe it's flat.

And why would it actually be flat? Because at some point it ends up being that you just need more land. So when I say it's Leontief, what do I mean? At some point you run out-- there's only so much we can do with a given amount of land. No matter how intensively we farm it or weed it or whatever, at some point you're running out of-- you run out of land.

Or the other reason that this could be is actually maybe labor supply is totally elastic at some reservation wage rate. And at some point, maybe the marginal product is still kind of like positive over here. But people will just stop working.

Because if people are being paid their marginal product, then at some point if there's some reservation wage rate-- if no one's going to work beyond-- say that the slope of this thing is equal to the marginal product is equal to the wage. If at this point-- if that wage no one's going to work any-- below that wage no one will work anymore, that's how much labor we would get. And all the people kind of beyond that would just be not working. OK? And so the rationale here was that there was an important goal for development was to bring this surplus labor kind of into the cities.

This question of should we have more people in the cities and less in the countryside is a very broad question of development. I'm not going to talk-- I'm going to talk a little bit about that right now, but we're going to talk a lot more about that in 14772 when we talk-- where we're going to talk about migration and the structural transformation from the countryside to the-- and moving people from agriculture into the urban areas. But I just wanted to say it's kind of-- it is related to this idea that people are not doing very much proactive. OK.

So in the spirit of some old papers, this is a paper-- there was a paper by Schultz in 1964 which was one of the earliest natural experiment papers that I can find, which was about the surplus labor hypothesis. So Schultz studies the 1917, 1918 flu pandemic, which killed 6% of the population in India and reduced the workforce by about 8%.

So you can imagine-- given that we were just going through a disease that has a 1% or 2% mortality rate, you can imagine what this must have been like. So Schultz says, well, look, if there was really-- if Lewis is right, and there's really 25% surplus labor, then we shouldn't have any-- this should have no impact on agricultural output.

So what do you think of this as a test-- as an idea? This plausible, not-- is it a good test? Bad test? What you think? Yeah, Ahmed.

AUDIENCE: Call it bias, but I kind of find it hard to believe a pandemic can have no impact on capital markets.

BENJAMIN Oh, on capital markets. What do you mean?

OLKEN:

AUDIENCE: So the test here assumes that, OK, there's a change in labor. And if nothing else changes, this will be a test of that [INAUDIBLE] on the margin, whether labor is needed for these markets.

But if they are getting their inputs from like, say, from the countryside, from other places, chances are this is like epidemic. Things were not going good from all places the other way around, as in the [INAUDIBLE] supply crisis now. So there could still be a reduction in production, even though-- even in the case where [INAUDIBLE].

BENJAMIN Yes, yes, yes. Absolutely. Yes, yes, yes. So let me just also say that I have taught this paper many times, and this is the first time someone has brought up supply chain disruptions in my-- in this. But I think we've all had a lot of experience as to the many different channels that pandemics can screw up the economy. So--

AUDIENCE: I'm just saying. I'm still waiting for my PS5.

BENJAMIN What? Sorry?

OLKEN:

AUDIENCE: I'm just still waiting for my PS5. That's still [INAUDIBLE].

BENJAMIN [LAUGHS]. But yes, yes, yes. Good point. OK. What else? That's a new one for this. I've taught this many times-- never heard that. OK. What's another-- yeah, Christine?

AUDIENCE: The people that were killed by the pandemic are not necessarily equal to the ones that didn't die. So it could have been like people that were more weak or less productive physically, or the ones that were killed. So I don't know. Or the other way around. It's not necessarily random.

BENJAMIN Yeah. It's not necessarily random. Again, we're having so much better insight on how to-- on understanding pandemics as an identification strategy. Isn't it-- yes, exactly. Right? It's not random. Yeah. Totally right.

OLKEN: And, in particular, it would be bad if it was-- he's actually going to reject this. And so if you thought the pandemic was killing people who were more productive than average, right? Yeah. Other thoughts?

AUDIENCE: [INAUDIBLE] assume--

BENJAMIN What?

OLKEN:

AUDIENCE: You would assume-- I would have assumed [INAUDIBLE] to it for the people who died to go the other way. Like, it's the-- I mean, I would assume younger people would have a higher probability of survival. Chances are in a [INAUDIBLE] culture, [INAUDIBLE] it takes a lot of manual labor to do the tasks. I would assume people are [INAUDIBLE] would be more productive than they just-- or [INAUDIBLE].

BENJAMIN Yes. That would go-- that will go-- if that was true, that would go against the direction of finding, in effect. Yeah.
OLKEN: Yeah. What else?

AUDIENCE: Maybe people that don't die but get sick have lower marginal product after.

BENJAMIN OLKEN: Yep. Yep. So maybe there's an additional thing beyond the death of people who are like-- who are sick. OK. Any other ones? Yeah, Hazel.

AUDIENCE: There's also like a demand shock, right? Because people who are dead are no longer demanding--

BENJAMIN OLKEN: What? Sorry?

AUDIENCE: There's also a demand shock because the people who died no longer demand food.

BENJAMIN OLKEN: Yeah. It's also a demand shock. Right. Totally agree. So-- and that would depend on whether or not thought the products-- whether it was world trade or whatever. Although, I guess there's probably a global-- probably demand. This is a global pandemic, so there was a global demand shock for food. So yeah. OK. So-- all right.

So you may not love this. And the other thing I would also-- so I would also-- the other thing I would also think about is the land reallocation process. Would it may have been particularly-- that's the one I've always most worried about, is that basically if someone died-- land markets are not super efficient and fast.

And so if I own my farm, and I'm a victim of the flu pandemic, then it's not clear who's going to take over my farm. And that process of reallocating that land to more productive uses may take a while too. So that's all good reasons why this is not the last word on this subject.

What did he do? He looks at the output and the post-year relevant to the pre-year. We had similar weather. And he looks at provinces that had greater influenza deaths, had greater declines in output. That's his test. And he looks at actually eight total acres sown. Actually, there's enough data and output. So they-- are they like plotting less-- doing less stuff.

And here's his-- he basically has the measure of deaths and has predicted reduction in agricultural output and observed reduction in agricultural output. So he's running a regression of this against this or against that. And I actually had to-- ran the regression on my computer. And you basically get-- with only 10 states, he gets the elasticity of output with respect to population of about 0.4.

So he is actually finding a pretty substantial reduction in output from the flu pandemic. Although, as we've discussed, there are other reasons why this could be an issue as well. OK. That's all by way of background.

But I guess-- I do-- I want to mention this question of are labor markets efficient, and are people operating-- where are people relative to the production function. And are maybe they way past the point which they have low marginal products is an issue that people have been thinking about for a very long time. OK.

I think it's-- sometimes I feel like we only teach really modern papers. And it's sometimes nice to realize people have been thinking about these issues for a while. OK.

So now onto this Benjamin paper, which you guys read for today. And some of you asked why did I choose this particular paper to read? It was not-- as you picked up, it was not for the empirics.

But I think that the-- and actually, I'm going to show you a more-- up until all of this-- I'm going to show you a more modern paper that is a better empirical test, actually, of exactly these issues in a sec. But I didn't want you to have to read two papers.

But what I like about this paper is I feel like it really clearly illustrates this idea of separation failures in going through the various cases and the various models, so I want to talk about that a little bit here. And I think this idea of separation between-- why is it an important development model?

I think that one thing that you see in a lot of developing contexts is you see a lot more self-employment. That basically, as we talked about, actually, in the last lecture the employee share is a lot lower in a lot of developing country contexts-- context. So here, in developed country most people are-- not everyone-- most people are employees, and-- whereas, in developing countries, I think most people are self-employed in some context, or working in their-- or production. And that means we want to think about both for understanding their own household decision-making, for understanding consumption decisions, but also for understanding production.

If we have separation failures, then understanding what's going on on the household side is going to be important for the productive side of the economy, too. And that's going to be really important in a developing country context if most of the people are kind of doing both.

Whereas, if we've separated things out, and we are in much larger context where people are employees, that issue is kind of less relevant for-- in a developed country context. So I think of this as like one of these important concepts that comes up a lot because it describes kind a lot of what people are doing in developing country contexts.

So the question he's trying to answer at some level is how efficient are rural labor markets? And, in particular, he has this test of the separation hypothesis. And broadly speaking, the way to think about the test of separation hypothesis is like is my production decision and is my consumption decision-- are those separate decisions, or are they conjoined decisions?

And the basic theoretical idea is if I have fully functioning efficient markets, then households can freely buy or sell their labor at wage w . And therefore, they're going to make two separate choices. They're first going to choose the labor input for their farms to maximize profits, given the prevailing wage w . And then they're going to separately choose the optimal labor leisure trade-off for the family, given wage w . And with full ability to buy and sell w , there is no reason those decisions should be linked.

So here's the key graph. So this is the key graph. This basically says, look, here's a production. This is the production. And by the way, I found these graphs a little confusing. I don't if you-- did you find them-- were they clear or confusing? [INAUDIBLE]. OK.

I'm going to try to go-- let me say I've taught them a couple-- it's took me a couple of times to teach them before they became less confusing to me. So I'm going to try-- let me see if I can elucidate them.

So what is going on? This is the production function. So how do you choose how much labor to use? We're going to look for the tangency of-- you're going to go to the point where the marginal product of labor is equal to the prevailing wage. So that's this-- that's given by the slope w . That's the prevailing wage.

The amount of labor you're going to want on your farm is given by the tangency of a line with slope w , the wage rate, to the production function of the farm. OK. That's this. So that's the labor used on the farm.

And then how much labor am I going to supply as a household? Well, I'm going to look for the-- I'm going to have some utility function. And I'm going to look for the tangency between my utility function and the wage rate. All right.

So I have some labor leisure trade-off. This is the utility function over here. The bliss point is over here. High consumption, low labor-- that's the bliss point. So the utility function is facing this way. And I'm going to go to the point where it's tangent to the wage rate, and that's going to decide how much labor I'm going to choose.

So what you can see here is these are set-- these are separate decisions. This is choosing-- this is going to be driven-- different people with different preferences. They'll move their labor/leisure trade-off.

Some people really like leisure. They're going to be over here. Some people really like labor-- they're going to, like, over here. Like consumption-- they're going to be over here. Whatever-- people will choose different amounts over here.

And then the farms-- different farms might have-- be steeper or flatter with different marginal products or different amounts of labor to sort of-- different marginal products for a given amount of labor. And therefore, different amounts optimal labor. And that's going to be a separate decision over here. So the baseline model has two separate decisions. Sorry, Chan. Do you have a question? Yeah. AUDIENCE: Yeah.

AUDIENCE: Those wages are [INAUDIBLE].

**BENJAMIN
OLKEN:** Yes.

AUDIENCE: So this is not like-- I was just wondering why the utility curve is not tangent to the production curve the same as [INAUDIBLE].

**BENJAMIN
OLKEN:** Wait, what?

AUDIENCE: Like, I was wondering like--

**BENJAMIN
OLKEN:** Yeah. This is market wage.

AUDIENCE: Which is [INAUDIBLE].

**BENJAMIN
OLKEN:** Yeah. So exactly. So-- and I think the idea-- the way to think about it is there's lots of different people with lots of different u 's around there. And some of them are down here, and some of them are over here or whatever. And there's some equilibrium in the labor market that's determining w . Yeah. OK, Other questions?

AUDIENCE: [INAUDIBLE] the labor-- [INAUDIBLE] labor demand not being called labor supply?

**BENJAMIN
OLKEN:** No. That was related to Vashan's question. No. For an individual household it does not. In the market, sure. But this is an individual household's decision.

AUDIENCE: If households are identical, then--

BENJAMIN Sorry. What? Sorry.

OLKEN:

AUDIENCE: If households are identical, then it should.

BENJAMIN Yes. If households are all identical, then it should, yes. But if there's heterogeneity, then there's no reason it

OLKEN: should in the-- for any given-- like some-- that's what I'm saying.

Some people really like consumption. Some people really like leisure. Some people have big farms, and some have small farms. It's all going to kind of-- for any-- there's no reason those should be the same in any particular household.

AUDIENCE: But a conceptual framework is like I-- so then the production is my own production, right?

BENJAMIN Yeah. This is my-- this is like my farm.

OLKEN:

AUDIENCE: I can work in my farm, and I can-- say, if I have three hours-- if I want to work three more hours, I can go work for three hours and [INAUDIBLE].

BENJAMIN Correct. That is the model.

OLKEN:

AUDIENCE: All right then.

BENJAMIN Yeah. Exactly. OK? All right. So now, what happens if there's rationing in the amount of off-farm work I can get?

OLKEN:

So suppose that, for example, like-- so, for example, suppose there was a minimum wage. So the minimum wage says that basically people don't want to hire that much labor up above some wage, so it's rationed.

So now what's going to happen? Or what might happen? So at this high wage, suppose it's rationed to H. So the farmer-- what is the farmer going to do?

They're going to first work at the outside high wage as much as they can up to this cap H. That's why this graph for the farmer is shifted to the left of the farm. They're going to work at H.

And now after that, then they're going to say, well, I still-- I'm not done consuming. Suppose they're not done consuming at H. They would like to work more hours than that. Then they're going to start working up their own farm. OK?

And how much are they going to do? Well, now they're just going to look to the point where the tangency of their-- they're going to-- how much are they going to work on their own farm? They're going to work up to the point where their utility function is tangent to the production function from the farm.

Because the more they work, as they work-- right over here, they're earning a lot, in some sense, because they're on the steep part. As they work, work, work, work, work, they'll start earning less and less and less. And they'll work to the point where they're indifferent. Or sorry-- where it's tangent.

And the tangency here is going to be at a-- at this high-- at this higher point here. And so-- yeah. So in this case, this is a case where they want more hours than they can get.

And what's critical here is that now the amount of work that's going to happen on the farm-- the amount of labor that's going to be employed on this farm with this production function over here is going to be driven not by the market wage, but by how much labor this particular household wants to supply. Yeah?

AUDIENCE: If the slope and the utility production dependency was higher than the off-farm work, they wouldn't work off the farm at all. Is that right?

BENJAMIN I'm sorry. Say it again.

OLKEN:

AUDIENCE: If the slope of that dotted line was steeper than the slope of the off-farm wage, then would they not work off the farm at all?

BENJAMIN If it was down here? Yeah. That's going to be coming in the next case.

OLKEN:

AUDIENCE: Oh, OK. Sorry.

BENJAMIN Yeah. If they wanted to work less than that, then they would-- then we'd be in a separation case. What they would do is they would work-- I think what would happen is on this farm, actually, what they would do is they would work, say, this amount down here, and then they'd hire some extra labor on the farm.

So it would be a separation case because they would be choosing the-- the amount of production on the farm would be driven by the market wage. So yeah. So it was only kind of an interesting case if the farm wants more labor-- if they want to supply more labor than that. OK.

The next case is kind of the opposite, which is to say, suppose there's some rationing on the amount of labor I can hire because the market wage is too low. Suppose the market wage doesn't go above-- it would clear at some higher level, but it's not going above.

And, by the way, I'm going to talk-- in the next lecture, we're going to talk about some evidence that these market wages do or not-- looking at wages directly, on do they or do they not clear and adjust and why or whatever. So that's another topic.

But what would happen-- what does this look like over here? So now, it's kind of the opposite graph. So in this case, the farm would like to hire this amount of the labor at the market wage, but it can't get it.

So how much-- so in that case, what's going to happen is it's going to get-- it'll hire-- it can only get this much L at the market wage. And then after that, then the household fills in the margin.

And if the household is in this range over here-- like if the household kind of wanted to work more, fine. But if the household is in this range over here, then, again, the total amount that's produced in the farm is driven by the household utility function.

The final graph they have in the paper on this is imagine that higher labor costs more than the farmers return to off-farm employment. What's that? Imagine that basically it is cheaper for-- it's cheaper and more efficient for me to do the work than for me to hire a laborer.

So you might think there's agency problems-- moral hazard. So I'm going to work really hard on my own farm, and kind of slack off and work on someone else's farm. So in that case, if this is the market wage-- the wage at the-- the market wage, if I want to supply less labor in the market wage, then fine. I can supply that, and I can hire some extra at the market. That's separation.

If I want to supply tons at outside wage, I can supply over here. That's fine. But in this intermediate range, again, will be driven by these tangencies.

OK. So what's the point of going through that? If we think that the market is fully-- is going back to this point. If the market is fully efficient, there should be no relationship between how much labor I want to supply and how much labor is used on the farm.

But with these different kinds of restrictions, if I can't hire enough labor, or I can't sell enough labor, I may end up either wanting to-- how much-- it may be that how much labor I can provide on my own farm determines how much labor is available on the farm.

Or conversely, it may be that how much labor I can just supply-- how much work I can get just depends on how much work there is to do on my own farm. And in those cases, the labor supply decisions and the production decisions are interlinked. That's the basic idea. Is that clear? OK.

So then, the idea of the paper is to say, well, if that's true, then do household demographic characteristics, which should affect labor supply, do those affect labor demand for the firm-- or for the farm-- for the farm, in this case. So that's the idea.

You say, look, we have shifters for household labor supply. If we're in the separation world those should not affect labor demand on the farm. And if they do, then we're rejecting separation. Is that clear? OK.

And I think this is a bit-- it's kind of related to a milder view of the surplus labor idea. If there are labor market frictions, you may employ labor on the farm even if the marginal product is below the outside wage. But we'll talk about this in a bit. But then this becomes the empirical test.

So why might there be the misseparation? So here are a few. So one could be literally minimum wages. That would create minimum wages if the market wage is too high. And therefore, I may not be able to-- the farm may not be able to get as many workers as it would like at that wage.

And the minimum wage could be a statutory minimum wage, or it could just be kind of a prevailing wage norm. And we'll talk about that in some of the papers we'll look at next time. It could be imperfect labor markets or agency problems or other market failures. So there could things going on in the labor market, so the thing isn't totally clearing. OK. So any questions on that basic theoretical setup? Yes? No? Clear? OK.

So as you guys noted-- so what did you think of the empirics in this paper? There were various concerns about the empirics in this paper that I read in your comments. So I want to-- what did you think? Somebody. Yeah, Ahmed.

AUDIENCE: Wages can impact population density [INAUDIBLE].

BENJAMIN What? Sorry?

OLKEN:

AUDIENCE: Population density can impact wages, yes. But can it go the other way around? [INAUDIBLE] using population density as an instrument for labor supply, like [INAUDIBLE]? I don't think [INAUDIBLE] any consumption would be satisfied for an industry.

BENJAMIN For using popula-- yes. Right. So-- and why?

OLKEN:

AUDIENCE: So you are interested in the effect on labor supply on wages, right? Say places that have higher wages can attract people. [INAUDIBLE] it can increase population density. Did that cor-- one second. I just don't see any way of--

BENJAMIN So OK. So because I'm on-- so OK. So, as you saw, one of the things he's doing is using population density as an

OLKEN: instrument for wages-- as an instrument for labor supply. His idea is that if there's more people around, that's going to lead to lower wages. So why is that not a good instrument for labor supply? Someone know? What are the concerns there? Yeah?

AUDIENCE: Population density [INAUDIBLE].

BENJAMIN How?

OLKEN:

AUDIENCE: [INAUDIBLE] or [INAUDIBLE] in cities like [INAUDIBLE].

BENJAMIN So one is there's demand for nontradeables. So like if we have-- if labor supply-- if pop-- in people-- where there are more people around, we need more haircuts. So it could affect it through nontradeables. Any other?

OLKEN:

I'll just give you another one, which is that-- so why are all those people there? Maybe there's something productive about the production side, which gave all those people there to begin with.

So, in particular, this one takes place in-- this study takes place in Java in Indonesia, which has a lot of volcanic soil which is super productive. And that's kind of why all the people are there, I think, is because the soil is productive. So it's not necessarily a suitable instrument.

So just as you guys-- I think you guys already saw this, but-- I'm sure. So-- in fact, I'm sure you guys have already seen this, right? So this is from an earlier version of the class where we taught this earlier. But obviously, you need-- the [INAUDIBLE] this is probably not going to hold in this case.

So another issue that comes up is this question of endogenous household size. So what's the issue? Was there [INAUDIBLE] going to say something?

AUDIENCE: If I want more labor, I'll have more kids.

BENJAMIN Huh?

OLKEN:

AUDIENCE: If I want more labor, I'll have more kids.

BENJAMIN Yes. If I want more labor, I'll have more kids. So now I want to ask you guys a harder question. Is that a problem?

OLKEN: And what I mean by is that a problem is can that generate-- suppose that was true.

So let's put us-- so suppose we're under the null hypothesis of δ equals 0. Can endogenous household size generate a false rejection of the null? Do you see what-- do you see the question I'm asking? What do you think? Yeah, Vashan?

AUDIENCE: Does endogenous [INAUDIBLE] to-- that the labor method should have the left-hand side? Then you'll get a correlation. But in terms of [INAUDIBLE], that's fine because you're trying to see if there's a correlation anyway. I don't know. This is only a distant correlation, right?

BENJAMIN Yeah.

OLKEN:

AUDIENCE: So in that sense, it doesn't matter whether you're [INAUDIBLE].

BENJAMIN Exactly. Right? So this is the idea. So in Wesley's story, I have more kids because I want more labor in the household. But the only reason I do that is because of separation failure. So if, in fact, there was no separation failure, the δ was equal to 0, then I wouldn't be having more kids to-- right?

OLKEN:

So I think-- I do think this is a bit of an issue, but I think it's worth thinking through, is this actually a problem in there? What I do think could be an issue is if we had omitted variables. Like, for example, if we have better land quality leading to more income, and more income leading to that kids are a normal good, then that could-- that channel could be a problem.

But that's a more complicated story than I think the one you were thinking of. Because your story is only a problem under-- if the null is false. But, in fact, under the null, it's not a problem. So I just think it's worth being a little more precise about the concerns that we have sometimes. Does everyone understand this alternative story and why this-- this could generate a false reduction in the null? Yes? OK.

What other-- any other comments or thoughts from this paper from your-- what else do you guys think of this from reading it? Any other questions or thoughts? No? OK.

So as many of you noted, the empirics of this paper are a little bit dated. And so one natural solution-- so LaFave and Thomas is kind of literally just updating this paper with better empirics. And, in particular, they have panel data. And so why is panel data going to help? Yeah?

AUDIENCE: Land fixed effects.

BENJAMIN So you're going to have land fixed effects. So what's going to be left? So what else-- right. So now we're going to have land fixed effects so we can control for all the heterogeneity of land quality and [INAUDIBLE]. That's good. What else?

AUDIENCE: Suppose we can find good variation in labor supply.

BENJAMIN Huh?

OLKEN:

AUDIENCE: We can find good variation--

BENJAMIN No. What's the variation?

OLKEN:

AUDIENCE: Over time, even if a place has, like you say, productive land--

BENJAMIN No. It's individuals, households.

OLKEN:

AUDIENCE: There's--

BENJAMIN Huh?

OLKEN:

AUDIENCE: Migration [INAUDIBLE].

BENJAMIN Someone else? So why would panel data help solve some of the problems from before? OK. We're out of time. I
OLKEN: will leave you to think-- ponder this question. We'll come back to it at the beginning of next class.

So you can have a look at this paper, and then I think whatever the next couple of starred papers on the syllabus are. We'll come back to this on-- actually, sorry. There's no class on Wednesday. And we're going to come back to this-- so a week from today.