[SQUEAKING] [RUSTLING] [CLICKING]

PROFESSOR: All right, so I have a fair amount to cover today, especially if I want to go all the way to what a lot of you were interested in in reading the Indonesia paper. I may not get to it, but we are going to do our best. Today, I want to talk about private and social returns to education.

What are the private returns to education? Well, we discussed them a little bit. They are, on the one hand, the impact on the labor market, and on the other hand, the impact on making it easier to raise a future generation that itself has a lot of education, which we can take as a shortcut for human capital. And the private return also includes the happiness that comes out of educating your children.

There is plenty of evidence that there is a strong correlation between human capital and earnings. One of these come, for example, from one of the descriptive first stats in the Indonesia study, where we see that people who have more education make more money, potentially with some evidence of a jump towards middle school. Might even look like a s-curve if you squint. So that's there.

There is similar evidence, by the way, on health. This is a correlation between people's body mass index or their height and their earnings. And you can also potentially-- you might be tempted to interpret this as the return to physical health. So we're going to be talking today about education, but a lot of what we are doing, both in terms of the intellectual content and the methodological content, can be found again thinking about the return to physical health instead.

And then now, also evidence of all sorts of other health benefits of health and education, which could be represented as forms of consumption, like how healthy your own kids are, political attitude, et cetera. There is also evidence of interaction between different forms of human capital. So educated people are healthier. The children of educated people are healthier. Healthier children means fewer days of school. And then they earn more as adults. For example, when you deworm children, they miss fewer days of schools.

They earn more as adults. I think we saw that in the first lecture. So if we go back to the basic correlation, what's the basic problem? Why was I tentative when I was looking at showing you these graphs, in interpreting these as returns, necessarily? And that is the causal effect of human capital, say, on earnings, education, or height, or BMI. Yeah, Karesh?

AUDIENCE: Selection issues.

ESTHER DUFLO: Yeah, there could be selection issues of various kinds, ability bias, omitted valuable of the kind that people who have enough money to feed their children also have enough money to make sure that they get the right apprenticeship. People who live in towns will have better schools. They will also have better job opportunities, and so on and so forth.

Another issue that we want-- so that's basically what we are going to try and tackle today, what a fair amount of the literature is concerned by. Another issue that comes up very much in the Ghana paper is the distinction between private and social returns. So the private returns are for an individual, how much money they make with or without education. And in fact, from the point of view of the parents, you could add a layer of private estimated return, which is how much they think the child is going to make with or without education. And then there are the social returns, which is the value for society of a particular individual getting more education. And so what are the sources of the discrepancy between private and social returns that we might worry about?

AUDIENCE: In the Ghana paper, there's the displacement effect, if some people get the government jobs, and other people get displaced.

ESTHER DUFLO: Exactly. So there can be displacement effect. In the Ghana paper, it's pretty obvious because education is used as a ticket for rent-seeking jobs. And then there are only so many of those. It's not that you can create more.

You could, to some extent, if the educated people who are so much better at the job of running the government-- you could increase the size of the government a little bit. But the truth is that the government is probably mostly rent-seeking jobs. There are only so many that can happen.

So it's basically a one-for-one displacement. And I don't think that's an exception. I don't think that's a special case in Ghana. I think this idea that we see in the Ghana paper-- I'll go back to that, but I'll show you the graph-- that we see in the Ghana paper, where people think that someone with education, what they might be doing is to get a government job very likely. And that is one key reason why people are looking to educate their kids.

Because there are usually restrictions on whether you can even apply as a function of your education degree. You need maybe a tertiary degree for which you need a secondary degree. Or you need just a secondary degree.

This is something that we are seeing. This queue for the government jobs is something that we are seeing in country after country after country in the developing world. In India, there is a very interesting fact, which is that you can only apply until you're 30, I believe-- maybe 35.

The specific date doesn't matter. Let's say 30. And Abhijit and [INAUDIBLE] have a paper where they show that the unemployment rate falls dramatically exactly at 30. So people basically until 30 still hope that they might get a job. They're holding out for that job.

Live with a parent, figure it out. And then when it's clear that they can't get it anymore, then they go and get a job. So these queues for those government jobs, I've come to believe, has a very destructive impact on the whole labor market. And in the Ghana paper, we are seeing them all over the place.

So that's an extreme example of displacement. But in the absence of government jobs, could you nonetheless have displacement effects in the absence of government, or in fact any rent-seeking opportunity with a perfectly competitive labor market? Would you nonetheless potentially see displacement effects? Yeah.

AUDIENCE: [INAUDIBLE] displacement with tertiary [INAUDIBLE]?

ESTHER DUFLO: Yes, so there is that for sure. There might be a ripple effect on tertiary education. But even if we consider the labor market and just the fact that there are more educated workers in the labor market, might that create displacement of other people? Yeah?

AUDIENCE: In standard search mode in labor, these [INAUDIBLE] people are searching for the same job. It should be pretty easy job [INAUDIBLE] for everyone else.

ESTHER DUFLO: Exactly. And would we find that it's very unusual labor markets, or type of modeling the labor market? Or would we expect this to happen fairly easily? Yeah, Aisha?

AUDIENCE: It's pretty standard, I think.

ESTHER DUFLO: It's pretty standard. It's going to depend, of course, on the production function. So one case where you're obviously going to find it is if people are working with a limited resource-- for example, land. And there's only so much of it.

And if educated workers are more productive, there is going to be a substitution between educated workers and workers in general, educated or not. So we have-- marginally, take any production function. There is going to be, depending on the form, the parameters of the CES function, or the Cobb-Douglas glass function of whatever it is that you choose that is appropriate to model the labor market-- there is always degrees of substitutability between workers.

And then that subsitutability is worse when there is a limiting factor. Land is an extreme setup because there is only so much of it. But in capital also, it depends how well the capital market functions and adjusts to the arrival of these new workers.

If, for example, as [INAUDIBLE] was working on before, eventually came to MIT. The labor market is slow to move from one place to another. Sorry, capital is slow to move from one place to another.

Then there is not an immediate adjustment of the demand of the labor demand to the new effective labor supply for your more productive worker. Then whether or not the externality-- who the incidence of the externality is going to be on, again, depends on the production function. So you could have models where educated and uneducated workers are perfectly substitute, except the educated worker do everything a little faster, in which case, the educated workers are going to be displacing everyone, including the uneducated workers.

If, on the other hand, the educated and uneducated workers are not perfectly substitute, or are very imperfectly substitute, then the educated workers are going to push out the other educated workers. And what you're going to get is not the decline of the wage of the base wage, but the decline of the returns to education, the differential between educated and uneducated worker. So this is what Richard Friedman now many, many decades ago was worried about when he wrote this book, *The Overeducated American.*

It's that a lot of people would go to school. And then they were, at this point, going to college. And they would acquire skills.

And then the return to education would collapse. Didn't happen in the US, but that is a concern. That's in the back of people's minds, something that Heckman also, for example, is talking about. And those are equilibrium effect, or general equilibrium effect of education.

So there is a basic, very standard reason for the social and the private return not to align, even in the absence of very uncompetitive sectors, like the labor market. So that's that. On the other hand, of course, what an institution like the World Bank was mainly motivated by when they were thinking of investing in education-- there are the positive spillovers of educated workers coming from easier adoption of technologies, for example.

That might make everyone more productive and hence increase everyone's productivity wages and so on and so forth. So that's from the A in the production function. And then there could be impacts in other spheres.

An educated population might support institutions that themselves lead to better policies, which again, would increase one layer above-- improve the institution, therefore improve education. So those would be the reasons for the social return to go in the other direction. That might be a little hard to pick up because if they are the very macro level, we don't have so many countries to even think of it.

But people have looked, at least, to see whether it is true, for example, that more educated workers support more democracy, or so on and so forth. I'm not going to do much of that today. But just keep it in the back of your mind that that's an important, that's an interesting topic.

So other point that's interesting in Ghana is, of course, people have a completely wild overestimate of how much of a ticket the secondary education is to tertiary degrees and hence the government. They vastly overestimate, as a result, the probability that sending their kids to secondary school would give them access to this job. And therefore, they vastly estimate the returns to education, which in their opinion, comes mostly from one of these jobs.

So where they are correct is that these jobs are hugely valuable once you get them. They are fantastic, very well paid. We saw during COVID-- very protective against shocks because you're paid no matter what. But where they are wrong is that they are pretty hard to get even after education.

So one way to get across-- we are going to look at two ways today to get around the problem of selection into education. One is random assignment. And the other is the school construction building that provides a natural experiment for increasing education. And both have advantages and disadvantages.

So of course, it's hard to directly randomly assign education. We were able to randomly assign an [INAUDIBLE] program and compare people who get it to people who don't get it. Even then we had some people who managed to not comply to our treatment assignment. But at least, this is something we could manipulate directly.

But education-- you can't really force people not to get an education. You can't force them to get an education, either. So what's randomly assigned in Ghana? Tends to be-- education is hard to randomly assign. Right side [INAUDIBLE].

AUDIENCE: Scholarships?

ESTHER DUFLO: It's a scholarship to attend school, secondary school. In what sample?

AUDIENCE: People who tested high enough to go to secondary school, but for some reason, saying they could not afford it.

ESTHER DUFLO: Yeah. In fact, more than claimed, right? Not only they claimed they could not afford, but they didn't show up.

AUDIENCE: Yeah.

ESTHER DUFLO: The first trimester. So they took some costly steps to demonstrate that they couldn't. And they could discuss this step when there was no scholarship in mind.

So presumably, they were close to could not afford. Exactly. So therefore, if we say in word, what can we confidently estimate by our standard methods of analyzing randomized controlled trial in table 2 of our paper? Shaina?

AUDIENCE: [INAUDIBLE] the scholarship on education?

ESTHER DUFLO: For?

AUDIENCE: For the people who tested but don't go.

ESTHER DUFLO: Yes. So we can estimate the effect of getting secondary school for people who qualify on pretty hard exams, but could not afford to go in the absence of the program. So we had some discussion about external validity later on. That's a pretty specific sample.

It's a sample of interest because, for example, it's a sample that would be relevant if you were considering making secondary education free, but not changing the requirement to get in, which is precisely what the Ghana government did shortly after the study started without asking us. So it's not something we could claim. But that was, at least at the time, very much the debate in India.

We are not getting very necessarily very informative estimate to what will happen. Instead, once they have done it now, well, there are two debates in Ghana currently. There's one party that says we want to revert to what we had before, remove free secondary school. And then there is the party who did free secondary school who wants to go further and eliminate the exam requirement.

So if we consider eliminating the exam requirement, these estimates might not be very informative on them because that's a different group of people. This might be a group of people for whom returns to education are higher or lower. We don't know. So we have to keep that in mind.

So we can confidently estimate that. And so these are your standard, basic way of looking at an experiment table, which you will come to know and love, or at least learn to read and maybe write. Usually, you write the mean in the comparison group, the treatment effect, the P-value of the treatment effect, the AR. And we eventually could have done away with stars to my great delight. And not just delight-- I made that happen because I hate stars.

[LAUGHTER]

This doesn't make any sense that 10% or 5% become magical numbers. It's annoying. So instead, you can report P-value if you think people shouldn't have to make division in their head.

And so that's the comparison mean. That's the treatment effect. If you added them both, if there are no control variables, that's the education in the treatment group. And so the program was effective at increasing completion of secondary school. Some people went anyways.

So we're going to keep that in mind. Almost half of the comparison group goes to secondary school and completes it anyways, not because they were lying about their inability to afford secondary school tuition, but because someone came along. And then basically, it's pretty common for people to try and figure it out during their first trimester or even their first year and enroll one year later.

So eventually, some of them could have found someone to give them a scholarship-- people like us, or the sugar board, or the government. But still, the chance to enroll and complete was larger by about 27%, corresponding to about 1.2 years of education. So that's basically completing secondary school divided by 4. Conditional on starting, most people complete.

So that's why we have 1.2 extra years. Some people get 0. And some people get 4.

If they don't go to secondary school, some kids get to go to a technical institute. There's a small number, but one we need to keep in mind because that means we are also changing the form of education people get, potentially. Once you go to secondary school-- so 27% of them completed secondary school. Of those 27%, over the year, they keep trying.

They keep trying. They keep trying. By 2019, some of those 27% had managed to enroll in a tertiary program and completed.

So this is the people who have enrolled. This is the people who have completed. Some of the people who have enrolled but have not completed-- they are still ongoing because this operation takes so long. That's part of an interesting finding in this paper.

Again, I don't think unique to Ghana-- but it was news to me is how long people wait to get into a degree and to cobble the money together to do it and all that. And in the meantime, they live with their parents or grandparents. And so this is the number of enrolled people, et cetera.

So that's for all. And then we have women and men. Similar effect on education. Basically, larger effect on tertiary school.

In fact, all of the effect of enrolling in tertiary education comes from women, something that's pertinent because although they don't catch up on-- even the treatment girls, if you add to that, the treatment girls are still less likely to have completed, or about as likely to have completed secondary school as the control boys and stay behind the control girls. But it is not true for secondary school, where are the treatment girls catch up with the treatment boys because the treatment boys have much lower treatment effect.

So for secondary school, actually, just giving them the entry, the ticket into secondary school makes them catch up with-- so that says something about the dynamic in the family. And we were discussing that a little bit when we were discussing the model, that parents either dislike their girls, or have less of an expectation that they might do something with their education. They probably underestimate the value of investing in their girls, from their own point of view, which is sending them to secondary school, which is why we have such a larger treatment effect for the girls.

And then they learn stuff. So that's nice. So we did test score. Is Where is my test score?

Yeah. Total cognitive score-- so as we were following them up in 2013, we administered to them a version of the PISA international tests. And reassuringly, people have learned something in the schools.

This, by the way, was a potential criticism of sending people to secondary schools. The schools are so bad people are not going to learn anything. That doesn't seem to be the case.

They've learned something across domains. They also have greater political knowledge. They have greater engagement with the media. They are more likely to be conversant with technology, like having a bank account, know how to use the internet, will have a Facebook account, such marvelous things.

We don't find any effect on fertilizer use. But we find a effect on health that I'm not going to talk much about here. Continuing and moving as this tends to get to secondary school, one of the pretty striking results in the context of this improvement in human capital, which are real, is the fact that there is really not much happening on the labor market.

In fact, there is really nothing happening on whether you're working or not. And few people work more and more over time. But even by 2019, we still have 73% of people who have done any work.

And this kept increasing. This was much lower at the beginning of this period. So people start really, really slowly to start working. They are not making much more.

They are not more likely to be able to cope with a crisis. But they are more likely to have a job with benefits or a wage contract. So one of these jobs was potential rent-seeking job. And if we focus on women, who are the ones who were likely to be enrolled in tertiary, moving down the line, they are also more likely to have obtained one of these precious public sector jobs.

However, note how few of them have one of those private sector jobs. Instead of the 80% or something that parents were expecting for their secondary school graduate, after almost 10 years of this, we are at about 10 percentage points on average, and 10 percentage point boost for girls, and for boys, who are now young men and women.

So in that sense there is an overeducated Ghanaian. The hope that they had of secondary education being a ticket was dashed, squashed by what happened over time. One advantage of having these people go on, and on, and on, and on, and never finishing is eventually, we hit the COVID crisis.

So we could look at the impact on labor market outcomes during the COVID crisis. And for women, who are more likely to have this public sector job, we see the protection at work. Because they are doing much better during the COVID crisis. And that's driven-- they're doing much better on average. But that appears to be driven mostly by these women who have this much more secure job-- in particular, teachers or nurses, employment job.

AUDIENCE: Esther?

ESTHER DUFLO: Yeah.

AUDIENCE: I have a [INAUDIBLE] question about something like that. So whether or not women got a public sector job is itself an outcome.

ESTHER DUFLO: Yeah, yeah. So here all I know is that on average, these women are-- that's why there is nothing here that tells you that it's them who are doing better. And it would be hard to say because that's itself an outcome.

AUDIENCE: Right.

ESTHER DUFLO: So all you can do is eyeball the data and say, well, let me look at-- there are very few of these women who got these jobs. If I look at all of the other women, they are not doing any better than the woman who didn't get the private job. But I can't really do that because that's a selected sample.

So that's why I was tentative in saying it. And it's not in our table. If we wanted to put it in table, we wanted to take this message home, then we would have to do something about selecting into the selection into the tertiary job, and do some bounds about what would the effect be for people who didn't get into the secondary job, which we could do with an exercise similar to the one we did for working at all in the paper.

We have bounds for selection into having any kind of earnings. And we create bounds by saying, well, suppose the people who have any kinds of earnings are the worst or the best. So you could do something like that. That would be the least model driven.

Or you could have a more parametric version of doing that, a Heckman selection equation of some kind for trying to select not getting the government job. And then you could control for that, a bit like I did in the other paper, where I said, well, I'm only going to look at wage workers. Because that's the data I have, but let me at least control for the probability of having one of those jobs, finding another instrument, using a Heckman correction equation.

So you could do something like that, which we haven't done because we left it very, very mellow, so to speak. Yeah? You go ahead, and then her.

AUDIENCE: [INAUDIBLE] so in 2019, I look at the sequence number 3 and 4, I think on jobs with benefits and depending upon this [INAUDIBLE] mode. In males, there is a [INAUDIBLE] anything on getting a better job with benefits for males.

ESTHER DUFLO: Yeah. Yeah.

AUDIENCE: That disappears in 2020?

ESTHER DUFLO: Well, in a job with benefits, it's already a bit-- it wasn't the most robust effect anyways. You have a P-value of 0.13. So--

AUDIENCE: Also, for females, the effect doesn't exist first for during a job with benefits in 2019. But in 2020, there is a stronger effect. Not strong, but close. At least, the coefficient estimate becomes very large, or almost as large as the effect on being a public centered employee.

ESTHER DUFLO: Yeah. So in general, job with benefits should be almost a subset of government employees because all the government employees have some kind of benefits. So the two should really be relatively similar, especially for women. Some men get some of these wage jobs.

So in general, they should look quite similar. But then that also comes from another place in the survey. So it depends a bit on what people tell us.

But if you look at the point estimate, they are almost the same between a job with benefits and government job. And if you look at the means, they also are almost the same. So basically, that's the same variable for women because there are very few women who get a job with benefits, especially probably in 2020, where everybody lost their job if they weren't a government employee.

AUDIENCE: I see.

ESTHER DUFLO: Now if you look at the men, they didn't have-- that's a bit the same thing, which is you have four percentage points more men still have a job with benefits in 2020. Imagine that all of the government employees are in that number as well. And then the treatment effect are gone. And presumably, a lot of people have lost these jobs, both in treatment and control.

So the extra job with benefits, the few, like working with an NGO, or working for the hospital, or being a driver in a big company-- they've gone during the beginning of COVID. So that's why you get this difference.

AUDIENCE: I see.

ESTHER DUFLO: So just so that I don't say good things about my own paper-- on camera, in particular--

[LAUGHTER]

What's pretty cool about this experiment? What's fairly remarkable?

AUDIENCE: The extremely long timeline coupled with very low attrition.

ESTHER DUFLO: Yes. So the extremely long timeline-- I don't know if it's very cool or if it means we cannot get around to write a paper. But we continue finding them. But that's useful because we are certainly learning things over the years.

It's necessary because people-- some amount of time was definitely necessary because you need to let people finish secondary school, start in the labor market. The people who are in secondary school start with a gap four years later. So you need to give them time to catch up in terms of experience and something like that.

So you really need to continue a little bit. But that creates the issue that you might lose them. So we were very concerned about not losing them.

We gave them these cell phones, which we kept changing from time to time, and sim cards, and called them regularly. And if we couldn't find them, we went to look for them in their homes and convinced them to do a survey.

I should say, it becomes harder and harder. They are a bit tired now. So I don't know how long we can keep them.

We are now doing their kids. We are doing a lot of work with their children along with a psychologist at Harvard, Liz Belcher. I'm going to talk more about her work on Wednesday as well.

We are trying to see whether their children have higher human capital. But they are starting to get tired of us. And they don't want to see us anymore.

But otherwise, during all this time, we keep finding them. We pay them small amounts of money for the interview. They're used to us, et cetera.

But it takes a lot of effort. And the low attrition rate is rare for this. So that's what's remarkable about this experiment.

But of course, we don't really have what we want because we are comparing the scholarship to not getting the scholarship. And this is a lecture on return to education, not return on free education. So the question is, can we go one step further?

Now there are good reasons why we are not doing it in the paper, which we will go to towards the end of the lecture. But it's still a good example to go to. It's easy to think about what you might want to do, and to do it for during the class. And then we can discuss the limitation of the exercise in this context and why we decided not to perform it.

So if the question we had was how education improves earnings, or test scores, or health, or whatever and we, again, said that we cannot manipulate the education directly, but we have this thing that has an effect on it, then what would we do? Well, to start with simplest case, assume that we can write this, where the treatment effect is constant. So we don't have this individual treatment effect we started with.

I'm going to change that very soon. But we'll start with that. Then that's easy. We have what's called an instrument in Ghana, which is an emphasis on potential. Because we'll discuss what characteristics the instrument has and whether we think they are satisfied here and so on and so forth.

The scholarships were randomly assigned to students who qualified. So let's call Z a dummy equal to 1 if you get the scholarship in this sample, and 0 otherwise. Getting the scholarship, as we saw, increased the probability to ever enroll in a high school by 25 percentage points. Some kids enrolled anyway.

About half of them complete. Some kids don't enroll even with the scholarship. About a quarter of people say, no, thank you very much. So there is noncompliance both ways.

So still, despite that, we can construct an instrumental variable estimates. In this setup, we can say, well, let's write the treatment effect on school participation-- enrollment, or years of schooling. Whatever we like. And then down the road, we can do the same-- as we have already shown you, the table where we're comparing earnings in treatment and control group, where treatment is defined by getting the scholarship versus not.

And then since I had written Yi as alpha plus beta Si plus epsilon, I can just write the expectation here. And this one as well-- take the difference between the two. And I'm getting this-- and once I have taken the difference between the two, I'm getting that E of this difference is alpha plus beta times this difference.

So if I'm interested in beta, I just have to divide this one by this one. And there's some assumption. And the assumption, of course, is I want to make that E of epsilon given Zi equal 1 minus E of epsilon given Zi equals 0. I would like to be able to assume that it's 0.

So when can I make this assumption? What underlies this assumption in this case, in words? What helps? And is it a good assumption? Yep.

AUDIENCE: In this case, the instrument Z is randomly assigned, so you wouldn't have confidence that whether you're assigned or not is independent of any unobservables that are captured and is epsilon.

ESTHER DUFLO: Yeah. So that's good. That helps.

But remember that the epsilon is in the education equation. So from the discussions we had last time, we know we have a reliable causal estimate of the impact of the scholarship on Y and the impact of the scholarship on Z because the scholarship is-- and so the impact of the scholarship on Y and the impact of the scholarship on S because the scholarship is randomly assigned. So that's a good start. But what is hiding? There's one more thing that's hiding in this assumption.

AUDIENCE: It seems like you might be worried about unobservables that would determine the treatment--

ESTHER DUFLO: Would you mind speaking up?

AUDIENCE: You might be worried about unobservables that are affected by the treatment. So it could be the cash that you get to go to school. It could also be having a--

ESTHER DUFLO: Exactly. So what we might be worried about-- we know that the people who got the scholarship are no different than the people who didn't. So that's your point. I'm continuing that.

I randomly assigned them. I've done this work. But then they got a scholarship. So the further assumption that is necessary to make that claim is that the scholarships themselves don't have a direct impact on earnings, which they might, in Kailesh's example, for example, because I gave cash to the family.

Maybe for the half of the kids who would have gone to school anyways, it's now much easier for them. There is not this financial strain, or it's not there. And so it's easier for the family.

They're better fed. And because of that, they earn more money later. So that would be one. Because I give the scholarship to everyone, even those who would have gone to school anyways, right?

So in randomized experiments, we have to be careful because the first one is for free from the trial. But when we use the experiment as an instrument for something else, the second one doesn't come for free. And it's not always satisfied. And you get a lot of dicey people from going one step too far and saying, well, I'm going to use-- I'm estimating my first stage.

That's the effect on education. My reduced form-- that's the effect on wages, for example. And then I can divide it, too. I have the returns to education.

It may or may not apply, depending on the setup. And here this is something that we might want to discuss. Well, we already started the discussion. There might be other reasons.

Oops. Nonetheless, if we believe them, then we obtain the effect of education-- so on knowledge, earnings, and anything else by dividing the effect of the program on, say, cognitive test scores, if that's knowledge by the effect of the program on education, or earnings, or whatever you want. We call that guide the reduced form.

This one is the first stage. The simple ratio of the two is the Wald estimate. If you estimate it-- and that's a simple form of the instrumental variable estimator.

Instead, you could run two SLS, which basically involves projecting an education on the scholarship, and then regressing the outcome on the projection. And then you'll get standard error. That's your two [INAUDIBLE] square.

But let's calculate the Wald estimate, just to make sure. So for example, if we wanted to do it for completed-- so first of all, one choice we have is what independent variable you want to use. You have one instrument for all of these things.

It's a little bit cosmetic, in a way. It's whether you want to-- because the information that's underlying it is always the same. But you might want to, say, let me do it for years of education, or let me do it for completed secondary school. Maybe it makes more sense to use completed secondary school as a mental frame because again, the years-- it's not that you're giving people an extra year.

It's you're giving a quarter of people four extra years. So I prefer to use that. So let's say it's 27%.

And then I want to take the cognitive score. It's 0.15. So if you recall, the 0.27%, 0.25 that makes it easier to do the calculation. What's the Wald estimate of education on test scores? I made it easier by saying the effect on completed primary school is 0.25. Yeah.

AUDIENCE: 0.16.

ESTHER DUFLO: A point?

AUDIENCE: Or sorry, so that one's 0.1 times [INAUDIBLE].

ESTHER DUFLO: Yeah, 0.16. And then first set test scores.

AUDIENCE: Yep, and then you divide it by the other one that we had. So you get 64.

ESTHER DUFLO: Yeah. So you divide it by 0.25, which is like multiplying it by 4. We get about 0.64. So we can do the same thing for-- so that gives us now the effect of completing secondary school on standard deviation of cognitive test scores.

That's very helpful to think, even for cognitive skills. Now we can compare that with other things we could do that change test scores. And we say, well, that's a pretty large impact.

0.6 standard deviation on test scores-- there are not very many things that move standard deviation on test scores by 0.6. We can also do it for wages. So that's going to be less impressive because nothing is significant. It's not going to give you a bigger ratio, which once you do it as standard error with huge standard error-- but this is the exercise you could do.

Now once we have done that, so suppose our Wald estimate-- so from the Wald estimate-- oh, we don't have cognitive score. So that's OK. We can do it with something else-- for example, with technology adoption.

If we compute the Wald estimate or the IV of education on adopting technology, if you remember, the reduced form was very small and insignificant. So the IV, similarly, is small and insignificant. But then if you note, the OLS is simply regressing whether someone adopt technology, and whether they completed secondary school.

If we had run the naive regression that we started with, what we have concluded education does to the probability of adopting a technology-- technology in general. Fertilizer, internet, health technologies, et cetera. You understand what this OLS is?

It's regressing the outcome on education. So what would we conclude from that with 0.15 with a standard of 0.01? Yeah.

AUDIENCE: That education increases technology adoption.

ESTHER DUFLO: Yeah, that education increases technology adoption. And it's very precise. And it's pretty large. But then when we do the IV, it's a much smaller coefficient.

So that's relatively intuitive. It's saying, well, people who are prone to adopt go to school, even without scholarship, of all of the people who have money. And then with the money, they can do other things. So we have a natural selection bias. So that shows what the IV can do sometimes.

I'm not going to talk now about the GML, but I'll go to that. Now you can see from this equation the danger with IV is that that thing is in the denominator. And if you have anything in IV, for example, that doesn't have a huge impact on your variable-- 0.25 is already pretty good. But sometimes in IV, the effect is much smaller.

You have an instrument that really doesn't move education that much. If it has a direct impact on your outcome of interest, that's going to be divided by a small number. That's going to blow up.

So any bias that exists in the reduced form due, for example, to a direct impact of your instrument on your outcomes, or in nonexperimental cases due to selection biases that are not properly taken into account, any bias like that is going to explode, which is one reason why that, combined with a certain amount of selection bias of what gets published, that an editor will not publish a huge negative result-- so the huge negative results get canceled. They stay in drawers.

In general, there is more bias. So the estimate goes all over the place because the bias adds to the noise to create-- estimates can be very large or very low, depending on which side the bias goes. If you only publish the positive one, then you're finding yourself in a situation where all the IVs that are published tend to be much larger than the OLS. So that is an observation that people have done.

Colm Harmon and David [INAUDIBLE] have a paper on that, showing how the IV is always bigger than the OLS for returns to education in particular. There could be substantive reason. You can definitely think of that. And we'll come in a minute for why IV might be better than or OLS for legitimate reasons.

But it might also just be because the IV are biased. All the negative ones-- the fact that they are biased doesn't mean they are necessarily positively biased, but just have bias, which is enlarged. All of the negative ones are dead. And then you only are left with the positive large ones.

So that's the problem with IV. And that's why you want to be very careful handling IV. And in general, I like it because it gives you a good sense of magnitude. But in general, I'm wary of it.

So that's why, for example, in the Ghana paper, we are accepting that one table, which is meant to compare with the OLS. And therefore, we need comparable magnitude. I often don't tend to print them. A lot of the literature is more going in this direction, at least in development, in saying, even when you have something that affects what you might be interested in and the other, and then that affects the outcome of interest, it's a little dangerous to do an IV because of the bias.

So in this particular case, we were already worried about-- one potential source of direct impact and therefore bias is giving money to the families to attend school. Now this is counterbalanced by the fact that you're taking away money from the families that do go to school that would otherwise not have gone to school, right? Some families now need to pay for the school expenditure, and the scholarship doesn't cover all of it.

So on the one hand, you're giving extra money they wouldn't have had for people who would have gone to school anywhere. On the other hand, you are taking money away, or you're making people spend more on education than they would have. Those are the compliant.

They end up spending more than they would have. In terms of budget, it balances out, given how the point estimate shook out. But that's one effect going one way, one effect going the other way.

You have no guarantee that they even out. So that's one problem. What would be another problem, another potential violation of exclusion restriction coming from the scholarship? Yep.

AUDIENCE: You followed them for 10 years. So maybe they feel very motivated to do the things because a team of international researchers followed them around?

ESTHER DUFLO: Well, I'm following the control group as well.

AUDIENCE: Oh, but they know they were given--

ESTHER DUFLO: Yeah, yeah. They might be grateful. I wish they were even more grateful so we can follow them even easier.

But they might be grateful, yes. And in general, and this gratefulness might be even bigger because we are following them. We are calling them all the time. We are monitoring them.

They feel they have to perform. And you mentioned got given, I think. So in general, even in the absence of the experiment, they might feel psychologically, oh, I got a scholarship. I'd better do something with it, and therefore be motivated to do more, which might bias their earnings either way.

Because they might say, well, I'm not going to take any odd job. And therefore, I'm going to wait for something better. I already had this luck. I'm going to have some more.

Alternatively, I'd better perform. So there might be direct psychological effect of getting the scholarship. There might be direct-- in this case, the scholarship was not tied to performance. But it often is, for example.

There is a famous paper by Josh Angrist and Michael Kremer on estimating the return to private school. But they gave people a scholarship to go to private school that were removed if people repeated. And I always found it very funny to use that as an instrument because you might expect a direct effect of motivation effect just to get the scholarship irrespective of-- so that's one issue. The incentives, the feeling very confident, you want something, et cetera-- any other that you can think of? Yeah.

AUDIENCE: Maybe since they feel grateful that you give them a scholarship, they feel tempted to say they're doing better than they actually are.

ESTHER DUFLO: Yeah. So then there might be a social experiment effect, or social bias in reporting, which I don't think is a huge deal in this case. But it's a very good thing you're pointing it out because in a lot of experiments, it might be. And by the way, this is something that it's not even a problem of instruments.

Or it could be even more because the people who are on scholarship could feel even more that they feel they need to report higher wages because of that. Yeah. So you could have a social desirability in the reporting. For example, it's there in the reporting of their wages. And then for example, on their test scores, they might put a bit more effort. So in fact, for effort, on the test scores, we try to track that-- how much effort they are putting on the test itself.

So all of this would create issues. And I'm sure you could cook up some more, given a bit more time. Forgetting that for a moment, let's assume this is not an issue. Some of this can be looked at in the data.

Many of them cannot. Any money can be looked at. It can be looked at only in circumstantial ways on the back side. For example, with the effort thing, you can try and measure how much effort they put in.

Now forgetting that issue for a moment, at least for a pedagogical purpose, we can ask ourselves for whom the estimates-- we are calculating them from whom. So already, we know it's for a sample of kids who qualified but didn't have money. So that's already a first restriction.

Now with an IV, there is one more. As soon as you stop assuming the constant return to education, some kids would have gone to school anyways. And we call these guys the compliers. They always take your-- they always go to school, so they always take your treatment.

So some kids did not go to school even without the scholarship. We call these guys the never-takers. And some kids were moved by the scholarship to go to school. We know how many of those exist-- 25%. We can also know a lot of things about them by looking at the average characteristics of those who go to school with or without the scholarship.

This is from work by Alberto Abadie. But we never know who they are, exactly. We call these guys the compliers. Those are the people who go to school with the scholarship but do not go without.

So one question of interest is how the returns line up if they are not homogeneous, if these different groups have different groups. Who are we estimating the return to education? So that takes a little bit of algebra.

So we have now our instrument. We go back to the-- and a little bit of notation because we have to keep track of both the instrument and the treatment. I'm going back to the [INAUDIBLE] notation of Wi for the treatment.

So we have now a treatment of 1 if you got the scholarship, treatment 0 if you didn't get the scholarship. So we have potential outcome with respect to scholarship and potential outcome with respect to the treatment. The observed treatment is like we had before-- Zi, Wi of 1 plus 1 minus Ci, Wi of 0, OK? We observe people-- either they won or they lost.

And so that's where the notation becomes a little bit tricky because we have Yi of 1 is the potential outcome if treated by Wi, not treated by Zi, OK? And similarly for Y0. The identification assumption-- so this is from a very classic paper by Imbens and Angrist, which one day will win them the Nobel Prize-- sometime soon, hopefully.

So the assumption is that all the potential outcomes have to be independent of the instrument. That's the parallel of the assumption that we were discussing before about our epsilon, right? And again, this implies not only random assignment of the instrument, which ensures directly that this is independent, but also no-direct effect on the Y's, which gives you independence of the Y's. Otherwise, you don't get that.

In addition, we are going to make one extra assumption, which is monotonicity. What monotonicity says is that you're either as likely to be treated with the instrument down to 1, or more likely to be treated, but not less. So people either never go to school regardless, or always go to school regardless, or are more likely to go to school because of the instrument.

But there are nobody who is a defier, which is they would go to school if they don't have the scholarship, but they refuse to go to school if they have a scholarship. So that's an extra assumption. In further work by Chaisemartin in particular, you can work on some relaxing on this monotonicity. It's also not particularly testable, so it's good to relax it.

But for now, we are going to assume that's the assumption. So you are either compliant and always take care, or never take care, but you are not a defier. In most experiments, honestly, that's a much more plausible assumption than independence anyways. So there is not much to worry about, but it's good to know that it's happening.

So we now have our groups. We don't have our defier. And now we can start doing some algebra.

So like before, we are going to add and subtract some terms. And we are going to express Yi as a full function of the W. So that's what I'm doing here.

This Yi is this one. And then this is this one. And then I can replace some-- I can move things because the expectation can be combined nicely. To go from here to here, you can just write things down and combine them.

You will be convinced that it works. So I get here this product of Yi of 1 minus Yi of 0 multiplied by Yi of 1 minus Yi of 0. And these guys-- so what can I do with these guys?

I can get rid of them by independence. That's nice. So now that we have that, we have to assess how that is. And it's easy to deal with because this is really only 1 or 0 or minus 1, right?

In fact, can it be minus 1? It cannot be minus 1 by monotonicity. So I'm left with-- so this one goes away.

Then there is this one. So basically, I'm taking all of the cases one by one. Either this is minus 1-- the property of that is 0, so that's going to go away-- or this is this one. That's also going to go away simply because it's 0.

And then I'm left with only this one. And so that's the product of-- so then we want to modify that a little bit. And that turns out to be-- this probability is simply the reduced form on education.

So if I divide this by the reduced form on education, which is my Wald estimate, it is equal to the treatment effect for people who have a Wi of 1 minus Wi of 0 equal to 1. So it's a treatment effect for-- so the ratio of this difference, the reduced form by the first stage, which is the Wald estimate, is the treatment effect for this group of people, who are who?

AUDIENCE: Compliers.

ESTHER DUFLO: Who are the compliers. So that's something you probably know. But I think it's good to go to why it's true once.

And I even spared you that. Normally, I do it on the board, but I want to have a little-- just very simple, beautiful pieces of algebra that summarize why this gives you the effect on the complier. So that's very nice.

There are some special cases. If the controls take control, for example, in our case, in the case of the [INAUDIBLE] program, where some of the treatment refused to take the program, then the compliers are the treated. Then it's the effect of treatment on the treated.

In the general case, these are those who are compelled by the instruments to get the treatment. And so there is further issue of external validity-- all above the fact that we were in a small sample. In addition, within that sample, we selected the people who would have gone to school just because of the instrument. And then that tells us that-- it gives us some guidance on how to pick instruments.

Because if you pick instruments that are going to select some exotic group of people, then your estimate is also not going to be of great relevance. Here we are selecting the people who, because of the scholarship, are going to school. So that's a group that is of interest.

So that's nice. We like that. So that means that you want to go beyond finding an instrument that is going to-- not only your instrument has to satisfy independence, satisfy independence of the potential outcomes with respect to the instrument to satisfy monotonicity-- in addition, it has to be interesting in the sense that it moves an interesting group of compliers.

So this one satisfies because in the real world, if a government were thinking of removing fees of education, presumably, they would do that, keeping in mind the fact that it's going to help the people who are credit-constrained today. And those are those people because they would like to go to school. And they go to school as soon as they have the opportunity.

So I think we went already through this. So that's good. So an advantage of doing the IV is that we can compare with other strategies.

So we already compared with the OLS. Already did that before showing you what the IV was. We can do a more fancy version of the OLS, which is the oils after controlling for everything we have in the database that could lead to selection.

We have a very, very, very, very rich database. So you might think that maybe this is enough to capture all the selection. And in that case, we don't need to go over all of this complicated machinery.

But there is one issue in doing that. So what we did in that paper is we applied the double machine learning algorithm that we developed in the paper by Victor, and Whitney, and many, many other coauthors, which is basically a fancy method to do Frisch-Waugh. If you have taken first-year econometrics, that's one of the first things you learn is Frisch-Waugh decomposition, which is basically you can project what you're interested in on what you are not interested in. And then what you're interested in is the Y and the S.

And then when you regress Y on S, you have-- when you regress the residual from the projection on each other, you get the same beta coefficient of interest. So this is basically an application of Frisch-Waugh and saying, well, we are going to do this projection in as flexible a way as possible using machine learning, and then run the regression. So that's double machine learning.

Then you have to do that in two samples, hence the double. But the idea is that. So that was nice. We were very happy with our paper.

And this is, in a sense, the first test of this paper because it works very nicely in Monte Carlo simulation, where all the assumptions are satisfied. But we never know whether all the assumptions are satisfied for DML. So what we have here is comparing the double machine learning estimates to the OLS and then to the IV.

And it's really not that great because most of the DML look much closer to the OLS. And despite throwing all this machinery at the problem, you don't make so much headway. So it seems that whatever is the selection, it's not easy to capture it from even very, very rich observers.

Now you might think, well, all of that is not correct because in fact, the IV is the effect for the compliers. So they can't be compared to the double machine learning estimate or to the OLS, which is the effect for the population at large-- or at least, some other weighting scheme. The weighting scheme in OLS gives more weight to the observation that contribute the more to the variation.

The weighting scheme for DML is who knows? I don't know. It's a mess.

So what we did here is that we said, well, let's do a double machine learning estimate. That is going to mimic as much as we can from the observable variable-- the LATE estimate, which is we are going to estimate the treatment effect for a population that looks on the observable variable, of which we are many. Remember, this is the hope-- the HD row here is that the observer will do a good job in capturing the endogeneities in the sample.

Trying to estimate the machine learning estimate, the machine learning corrected estimate of education on a population that is weighted appropriately to look observationally, like people who are affected by the instrument. So what we need to do now is to apply our machinery in our other people to look at heterogeneous treatment effect by axes. We calculate the heterogenous treatment effect by axes.

We don't actually want to give them a causal estimate. But we just use machine learning to estimate, to get a sense of the heterogeneity in the treatment effect. And then we use those estimates to reweight so that we have the DML LATE is a machine learning estimate for a population that looks like the complier in terms of all of the observables that we have. That is, it puts more weight on people-- on the type, on the kind of people who seems to be more affected by the instruments.

After all of this effort, this is almost the same as the DML and almost the same as the OLS. So this is not-- this didn't make much of a difference. Yeah?

AUDIENCE: I just don't understand-- if you don't believe the exclusion restriction, why did you use the IV as the--

ESTHER DUFLO: Oh, yeah. Until now, I believe it. Yeah, so you could say, well, this is all-- all I can tell you now is that the DML is very similar to OLS despite the addition of the controls. The ML LATE is still very similar to the OLS despite weighting appropriately for the observable.

And they're all quite different from the IV. But you could tell me I don't believe the IV anyways. So there is nothing here that tells you you have to believe in one versus the other.

AUDIENCE: Could you compare it to the [INAUDIBLE]?

ESTHER DUFLO: No. You could, but you would get nothing from it. Well, what you could do-- and would be interesting, but we couldn't do in this data set-- is if you knew who-- you could compare a DML estimate of getting a scholarship to an OLS estimate of getting a scholarship in the control.

Not going to school, but you don't want to compare the going to school to getting a scholarship because that's not the same thing. You definitely have to scale them somehow. And then whether or not it's appropriate depends on whether or not your exclusion restrictions are valid.

But to not have this problem, you could find a setup where you are going to compare-- you would like to compare an OLS to an OLS. So in my case, it would be, can I compare the effect of getting a scholarship not randomly-- in the control group, for example, some kids got from the sugar board. Some kids got from a Christian church.

Some kids got from whatever. Can I compare the effect of getting a scholarship to the effect of getting a randomly assigned scholarship? And that would be a much easier comparison.

But we don't really have that in our data set. In fact, a lot of people scramble the money somehow without, really, a scholarship. So we don't have a way of doing that.

But that's what you would like, exactly. In the original LaLonde exercise-- Bob LaLonde is a researcher at Chicago who compared a lot of non-experimental estimates of training programs with the experimental estimates. And the way he did it is by-- this was comparing like to like because he was comparing assignment to training to going into training of your own will.

He didn't have DML, but he did matching and whatever. Actually, he didn't do matching because he didn't have it at the time. He did OLS, OLS with control, diff and diff, et cetera. And then some people added matching. And one could now add this.

But I haven't seen people doing that. So you're absolutely right that ours is a bit more fragile because you have to believe in the IV to believe in the comparison being valid. OK. So a minute on the substantive findings.

So we find positive effect on all sorts of sides of education. Financial return to education-- they are in the sense, maybe not so much on earnings, but on this-- some people get these perks. But they seem to be mostly private returns.

I'm almost running out of time, so I'm going to skip most of describing how we build an IV from the diff and diff estimator in the INPRES paper because you all read it. And once you've done that, that's just the same except that's a valid diff and diff instead of a valid from simple difference. And then you construct your instrument as the cohort interacted with the treatment effect-- or cohort interacted with the number of schools built in the region of birth. So I'll skip that part.

Then we find that differently from this setup, we find that the IV is very, very similar to the OLS. The IV and the OLS look really almost the same. So in that case, we didn't find much evidence of-- but what I want to talk about for two minutes is once we've done that, so we estimate the treatment effect.

One thing that-- the follow-on on the [INAUDIBLE] paper, follow-on in the compliers-- if your instrument takes a lot of discrete value, not only the effect for the complier, but you can show that it's a weighted average of the effect of each year of education, where the weight is given by the fraction of people who are moved to complete each extra year, which is why, if you'll remember, when I started this project, I wanted to estimate the nonparametric return to education. But that didn't work out because all of my variation comes from primary school.

So there was no way I was going to look at the effect of secondary school. I wouldn't worry about this negative estimate. These confidence bends are hugely too small because they are point-by-point. And they don't take into account clustering.

And when you do all that, the standards are actually large. So you can treat that as a whole big 0 on both this one and this one-- the small spillover on junior high and the negative on senior high. Really, it's driven by the fact that they built more senior high schools where they had more primary schools to start with.

So the school building of senior high school was negatively correlated with school building of private school, which itself is an issue for the exclusion restriction. But anyway, all the variation was coming from my complier, or my poor kids, who are driven from doing one more year of education, two more years of education. So there was no way I was going to estimate a nice, nonparametric curve like the one I showed you at the beginning of the graph, which had been my hope at the beginning. So that never happened.

But what did happen is something many of you asked. So I want to at least advertise thinking about the social returns to education. And something which we talk about in the Ghana paper, but couldn't do because we have an instrument that just affects individuals-- but here are the effects, the education policy-- it [INAUDIBLE].

And unless people migrate like crazy to undo all of the variation in education, which we know they don't, then that really provides, in a sense, a shock in the education level at the market level. So we at least have one hope to estimate not only that the effect of education on myself, but it should not be a beta. It should be a beta 1 or a gamma-- the effect of the average education of other people in my district on my own earnings.

So the estimation problem-- now we have two things. We need an instrument for my own education and an instrument for the average education in the neighborhood. And here we have a nice setup for that because if you consider a cohort that was 12 or older in 1973, and that's not exposed by the program, then what the program does to them is that when they enter the labor market-- let's say they start entering the labor market in 1980-- every year they age, they get a new cohort of educated kids who are coming to compete with them.

Themselves-- never benefited from the school. So for this group of people, there is no worry that the education level, the same thing that moved the average education level also moved their own education. We don't have that bias that might occur in other setups because we know they weren't exposed to the schools.

So we can-- but starting in '79, there is a slow increase in the average education of the other guys. So this is the effect of my own-- of the program on the education of the older cohort. So this is the same pattern of coefficient we had before except that now I had to take one labor survey for every year. Because I'm now looking at the education of-- all cohorts are observing-- '86, '87, '88-- as a function of DD relative to the first 0 here.

So the old cohort doesn't get more educated as they get older. They have missed out. But this is now the average of the young around them. And you can see that the average of the young around them keeps increasing because these new guys finish-- or girls-- guys, in this instance, finish school.

And therefore, they show up. So now we can do the same thing for the wage of the older. And if we find any effect, positive or negative, that's going to tell us that if there is a rising effect of the program with every year that passes on their wages, it's the fact that they compete with those guys.

And what we see is that there is actually a decline-- the coefficient-- the first one is higher. But then they tend to go down, and down, and down, and down, and down. This is super imprecise.

So the actual tables are mush. But if anything, what you're looking at, what you're finding is, in fact, a decrease in wages for the old educated, for the old cohort as these other people enter the labor market. And this decrease in wages is as large for the educated and the uneducated worker.

So it doesn't go to a decline return to education. So they appear to be substitute and substituted away by those guys. And even though the table-- the results are super imprecise, you can hold your nose and do a calibration to say, what does this imply about the movement of capital to this region when more people got educated?

And basically, you can't rule out that the capital is a fixed factor. So no one realized that there were suddenly all these educated people around. And they should have built more factories, and so on and so forth. It's either they didn't realize or they didn't have the credit to do it.

And therefore, the capital end up being a limiting factor in this case. So I write down a model, which has land. And people can work in factories.

And then I calibrate how much capital-- land is fixed, and then how much capital could have come in. And I show that this is compatible with none in some simulation. This was more of an intellectual exercise than a contribution that I want you to take as-- don't take, necessarily, the substantive lessons home.

Because it's very imprecise. But technically, it's a useful exercise to do. OK, let's stop here. Sorry for running over to some extent. And we'll talk about the supply side of education next week.

[SIDE CONVERSATIONS]