

[SQUEAKING]

[RUSTLING]

[CLICKING]

BENJAMIN OLKEN: All right, so I just have a couple of things to say to finish off kind of what I want to talk about land, and then we're going to switch to some of the lectures on public finance and redistribution. OK, so the two last points on land. The first thing is I realized as I was looking over my slides this morning that the last thing-- I had four different reasons why we care about land titling.

I never showed you reason number four, which is, is there, in fact, kind of a collateral effect? Right? The original motivation for doing all this land titling was that people could take their land titles and use them to borrow.

Let me just note there's a couple of papers here that show that while there are effects on investment through the time horizon channel, they don't seem to be as much through this kind of collateral for loan channel as you might expect. So I'm not going to go through that in detail, but I wanted to just mention this.

And they put this idea that actually you think that collateral can make things easier because it makes the property rates clearer, but actually if the legal system is not perfect, maybe it can be hard to actually seize that land. The whole point of collateral is that you can take the land in the event of a default. If the legal title is very strong and the legal system for collecting on the collateral is imperfect, it could actually potentially backfire. So I mean, I don't want to go into that in a lot of detail. I just wanted to mention that initial justification doesn't seem to be as strong as we might have thought.

And the final thing I wanted to say about land overall is a bit on the macro facts. And I guess some of you asked some questions about this. What are the implications for distribution of land? How do we think about certain allocation issues? And some questions are like, in particular, one question people look at is are farms too big or too small? What's the optimal size of farms?

And so, given what we talked about in the previous set of lectures on sharecropping and things like that, do you have any predictions as to what would you expect? Is the optimal size of farms big or small? Or what do you think about that? How did you think about it?

From everything we went through in the previous lecture on thinking about sharecropping or whatever, does that give you any predictions? Do you think that small farms or large farms are going to be optimal? Yeah.

AUDIENCE: For clarification, do you mean who owns the land or who works there, thinking about [INAUDIBLE]?

BENJAMIN OLKEN: Let's say who owns the land. Yeah.

AUDIENCE: Maybe if it's too small, because if it is too big, it will be hard to protect.

BENJAMIN OLKEN: What, sorry?

OLKEN:

AUDIENCE: If it's too big, it'll be harder to protect, so you make it efficiently small.

BENJAMIN Hard to protect?

OLKEN:

AUDIENCE: Yeah.

BENJAMIN Oh, yeah, exactly. So in the land titling stuff, I agree with you. It can be hard to protect, so maybe they're

OLKEN: artificially too big. Or maybe they want them small. What else might you be thinking about the Implications for firm size? Yeah.

AUDIENCE: Well, if we're worried about these moral hazard issues, then you might to give the land to the farmer to avert that, but you can only do that if the plot is small.

BENJAMIN What's that?

OLKEN:

AUDIENCE: If you want to avert these moral hazard issues, you might want to give the land to the poor farmers themselves to farm it, but you could not feasibly do that on a very large plot [INAUDIBLE]

BENJAMIN Right. So that suggests the opposite. So if you think the moral hazard issues are really important, you might

OLKEN: actually want small farms, because basically you want the person doing the farming decisions to be the one who's the residual claimant on that farm. So I think that that's one of the reasons farms should be small. Are there other reasons that point in the other direction, that you think farms should be big? Yeah, Whitney.

AUDIENCE: You could have technologies that work better with economies of scale.

BENJAMIN What do you mean?

OLKEN:

AUDIENCE: Like, if I get a tractor, this isn't great if I only have a tiny plot of land, but it's really good if I--

BENJAMIN Yeah, exactly. So the flip side of that is that exactly. A lot of these technologies work better on large land. So

OLKEN: actually, I remember when I took this course many, many, many years ago-- Abhijit, who was teaching it at the time, gave the great analogy of just imagine a giant harvester combine, like the kind they use in Iowa or something. You couldn't drive it down like most plots in Indonesia or Java, for example. You couldn't even turn it around. Like it's just designed for a totally different scale.

And so on the flip side of that, you may think that as agriculture becomes more capital-intensive, there's a lot of technological progress that is consistent with kind of large farm size.

So here's a couple of just data points on some of this stuff. So just to show you, here's an interesting fact. This is from a paper by [INAUDIBLE] et al. This is GDP per capita on the x-axis and log farm size on the y-axis. So what you can see is that richer countries tend to have bigger farms. So what do you make of that?

AUDIENCE: --the costs are very different, so that effects--

BENJAMIN Huh?

OLKEN:

AUDIENCE: Like, labor costs are very different, so they might be more capital-intensive in the high--

**BENJAMIN
OLKEN:** Yeah, exactly. So this doesn't say anything about productivity, but exactly. Like imagine a two sector model where basically you can work in-- you're doing multiple different ways. So one is as we have more capital, we'll switch from having labor on the farm to capital of the farm.

So if you think about the way farming is done in Iowa and the way farming is done in Indonesia, it looks super duper different. Right? You have a couple of farmers in Iowa farming 100 acres with a lot of machines, or you can have one farmer in Java farming a quarter of an acre or of a hectare or whatever basically him or herself using a lot less capital.

And so what's going to determine that? Well, as the relative price of capital of labor as the economy changes, as kind of a place to become richer, you know, you're going to end up kind of with a more capital-intensive and less labor-intensive production technology that's going to look like having larger farms.

Or alternatively, you can think of what is the wage of the farmer? The wage of the farmer is going to be equal to the wage of the farmer in the non-farming sector. So as the economy becomes richer, that one farmer is going to be doing more, and that's going to look like they're doing that and farming more stuff.

So my point is this does not tell us that one of these is more productive, but it is, I think, a pretty striking fact about the world that the way agriculture looks in places like the US looks very, very different than the way it looks in places like India or China.

So this is also the US. This is actually in the US, by the way. Just looking within a given country, it's kind of interesting. If you compare large versus small farms within the US, they look really different. So small farms have higher value added per acre, but lower value added per worker.

So that's consistent with the idea that there's different farming technologies. The small farm is kind of very intensely-farmed by a few workers. The large one looks like they're using a lot more capital. The capital-land ratio is higher on the small farm. They farm more intensively. But on net, they're probably using a lot more capital in some larger places.

So Foster and Rosensweig have this nice paper in 2017 trying to look at some of these issues within India. And then what they basically argue is that they find a U-shape in farm productivity where essentially they say you want to either have small farms or really big farms.

And their argument is that basically-- this is in profits per acre-- the small farms kind of work really well, because they're farmed really intensively by an owner-operator with low moral hazard concerns, and the big farms start to work well, because once you get big enough, you switch your farming technology completely and you start using large-scale industrial farming equipment or whatever. And so that just says that there are these two very different factors.

And they have a piece of evidence of this. And behind this is the idea that a lot of these technologies have returns to scale. So this is the cost per horsepower for electric motors and submersible pumps by the horsepower. So what's the cost per unit of output from the pump as a function of pump size?

And the point is that it's declining, so that looks like there's returns to scale. So you're going to get the same amount of pump-ness or whatever much more cheaply with one big pump rather than a large number of small pumps. Pumping capacity I think is the word I was looking for.

And they show this for a variety of equipment. Like that's for pumps. They also show that adoption looks larger on larger farms. So for example, tractors and sprayers, larger farms are using them. And so on and so forth.

So I just kind wanted to end with this by saying that I actually think that there's a number of interesting questions about how do we think about the optimal farm size literature? I like this paper by Foster and Rosensweig, which helps think about some of these issues, but I wanted to spend a few minutes and put some of this on your radar screen as well.

I think in particular, there's a lot of interest in how do we think about the optimal allocation of land, given the whole mass allocation literature. There are a few papers that I mentioned on the syllabus, which think about these issues, but I just wanted to put a marker on this as well before I move on.

Yeah, Charlotte. I'm sorry, I can't hear you, Charlotte.

AUDIENCE: [INAUDIBLE] for example, if there's misallocation [INAUDIBLE] because the farmers need less [INAUDIBLE] on small farms.

**BENJAMIN
OLKEN:** Sorry, so maybe. But say your story again. Sorry.

AUDIENCE: Say there are wedges of the labor market. And big farmers face lower wedges than small farmers. Then big farmers would be able to access more [INAUDIBLE]

**BENJAMIN
OLKEN:** If bigger farmers face lower wedges. Yeah, sure. Except that wouldn't explain this fact. So that might explain sure like if there's lower wedges then maybe they'd be more capital-intensive, but it wouldn't explain this other fact, which actually looks kind of there's returns to scaling kind of technologies.

AUDIENCE: I thought this captures the cost of capital, and we see that the cost of capital is higher for smaller farms.

**BENJAMIN
OLKEN:** No, no, no. Not this. Not this one. This is like just a literal fact about the capital equipment, itself, which basically says, if I want to buy a piece of equipment, it is cheaper per unit of equipment to buy a really big one than to buy a bunch of really small ones.

So then you actually get into the microstructure of how do you actually share equipment. And that gets to the point I was saying about driving the combine around. But it may be that if we want to do irrigation, for example, maybe we could do it with one big pump and share the water, but that would make it complicated.

And so we might all instead prefer to have our own little pumps. And the point of this is that all of our own little pumps is more expensive. And that's kind of the micro story for this, I think, is that there's something about the equipment itself that has this kind of like returns to scale on it. And that would be the difference. So I agree with you. Just showing more capital on big farms doesn't nail it. You really need to see something about the microstructure there.

The other thing I'll point out is there's also some interesting work, which I'm not going to go through, but which says, well, maybe if this is the fundamental problem, we can break this with rental contracts, and maybe we can buy one big tractor and share it amongst each other. And I'm not going to go into the details, but thinking about that seems like a really interesting thing to think about. And there is some existing work on that as well.

Yeah.

AUDIENCE: [INAUDIBLE]

**BENJAMIN
OLKEN:** What, sorry?

AUDIENCE: Like, a political economy reason for like why landowners might not want to sell small plots of land. They might want to maintain and control the whole land to keep wages down.

**BENJAMIN
OLKEN:** Oh, yeah. Sorry, so there's two different things you basically said. One is kind of a labor market story, which is maybe they want to become monopsonists and drive wages down. That could be. I can't think of a paper off the top of my head that looks at land concentration and monopsony power, although there's been a lot of interest in monopsony in the last couple of years, so there may very well be and I just can't think of one. But that could be.

Second point is it could also be relevant for political economy reasons even separately from this monopsony reasons. Those might be different. But yes, it could be. Yeah, for sure.

AUDIENCE: The two technologies on the other side of the [INAUDIBLE] both of them have--

[INTERPOSING VOICES]

Suppose like the farmers with [INAUDIBLE] for very large [INAUDIBLE]

**BENJAMIN
OLKEN:** Yeah.

AUDIENCE: --technologies that you described, and they look equally [INAUDIBLE] so does that mean they have to same [INAUDIBLE] per acre, but do they-- what are the implications for who they become goes to? Like I'm thinking like one of these is much more labor-intensive than the other, so the implications for wages or rural in conflict or labor--

**BENJAMIN
OLKEN:** Well, it's complicated. So before you start thinking about wages, let me actually, at some level, defer that comment to the labor lecture, because if you think about wages, you got to think of what's happening in labor market.

Right, so to first order, you might want to think about it as workers get a wage, which is the market wage, and the owners of the land get the land rents. And so to the extent you have kind of owner-farmers or whatever, then maybe you're right that that could be affecting their kind of net incomes or whatever, but I do think it's important to separate out kind of the rental returns from the land and the productivity of the land from the wage people are getting, which is going to be determined by the labor market.

The caveats to that are, of course, the labor market-- that's if we were in a world with perfect labor markets, which we are not, which I will talk about in the labor lecture. Yeah.

AUDIENCE: What's the story behind having a middle-sized farm and not [INAUDIBLE]

BENJAMIN OLKEN: Oh, yeah, why does it even exist? I don't have a good answer for that. Right, if you saw this graph, then the distribution should be kind of like bimodal, and I'm not sure it is. But yes, I agree with you if you saw this. Maybe there are inefficiencies in land market, too.

AUDIENCE: Wouldn't this imply like great benefits from increasing credit availability for like middle-sized farmers?

BENJAMIN OLKEN: You know, here's one answer, the thing I said about the Coase theorem last time. So the Coase theorem says that over time, we will end up with the optimal farm size, and maybe it'll be bimodal, but that may take a really long time.

So in the example I gave from Georgia, it took 100 years. So the technological progress is happening at a more rapid rate. So it may be the optimal farm size is happening faster than the Coase theorem mechanisms can reallocate kind of farms to their optimal sizes. That'd be one answer.

Let me also say, actually, I think that another point is that understanding land sales and understanding how that land market works for agricultural land is something that I think there are tons of frictions in that, and I think actually that's not something I know that's particularly well-understood. So I'd put that out there as an interesting question for future research.

OK? Great. So let me switch gears now.

So the next thing I'm going to talk about in the next three lectures or so is how do we think about the public sector in the developing world? And I guess before you even get into the details of this, a bit of an overview, I think that there are going to be several themes.

So why am I teaching this in a development class at all and not in the public finance class? I do think there's lots of complementarities, by the way, and if you're interested in this issue, you should probably take both public finance class and obviously this class.

But while a lot of the issues that we think about in the public sector in the developing world are similar to the things you take in a general public finance class, I think that one common theme you'll see is that the information environment and the state capacity environment are very different in a lot of developing countries. And that means that a way that we tackle a lot of the same fundamental challenges of how do we think about redistribution and how do we think about taxation are going to look really different in a lot of these developing country contexts.

And I think that a lot of those challenges of poor information or whatever, while obviously information asymmetries is the big issue in economics everywhere, I think that a lot of those issues are important themes that come out in kind of a development economics context.

So that's my reasoning for why I want to put in this. And I think you'll see those themes of like low information and state capacity challenges over the next couple of lectures. So that's just a bit of an overview.

OK, so I'm going to divide what I want to talk about into two pieces, kind of the redistribution and safety net piece and the taxation piece. And I'm talking about both of them. So first, we're going to talk about redistribution.

So I think that sitting here like in a wealthy country, people can often think of poverty and development as going together as the idea that most people in developing countries are poor, but this masks substantial inequality within poor countries, for example.

And if you go to a leader of any of these developing countries and talk about poverty, well, they'll be thinking within our country, who are the people within our country that are really poor? And there's a lot of inequality within these countries. And that's true throughout, I think, the developing world.

So for example, when I was writing these slides, I happened to have Indonesian micro-data handy because I do a lot of work in Indonesia and that data was on my computer. So for example, in Indonesia, the 10th percentile household as of a couple of years ago when I wrote these slides, had a GDP per capita of about 1 USD per day, and the 90th percentile household had had income per capita of about 5 USD per day.

Now that's a lot lower than households in the US, for example, but I think the point is that even Indonesia, which, by the way, is a very equal country compared to, say, places in Latin America, which have a lot higher levels of inequality, there's substantial income inequality within these countries as well. And therefore, there's substantial scope for redistribution even within a lot of these developing countries.

And so you could imagine that people who are here or even particularly if there's tales of really extreme more wealthy people, you can imagine that there's a lot of interest in doing various programs to help the poorest members of these communities.

And in particular, as I will talk about in the next lecture, tax capacity in the developing world is starting to increase. I think there's a lot of interest in developing countries in general, particularly countries in kind of the lower end, like low middle income and middle income kind of range, which includes a large number of countries like India, for example, in doing a lot more of these kinds of programs.

And although we hear a lot about NGO-led anti-poverty programs, my own personal take is that the ones run by the government are just orders of magnitude larger in scale in almost all places. Not everywhere. A few countries have very, very large NGO sectors.

But in general, the scale and scope of the government to do this kind of redistribution-- particularly because they can fund it through taxes, so they have this kind of ability to raise a lot of revenue and use it for this stuff-- is really large.

So here are just a few examples to show you. This is from a great report called the State of Social Safety Nets that the World Bank has put out a couple of times. This is from the 2015 one. There's a 2018 one that has some different tables and figures. But just to give you a sense that this is happening in a lot of places.

So cash transfer programs, does the country have cash transfer programs? They, as of like seven or eight years ago, are in 40 different countries in sub-Saharan Africa.

Conditional cash transfers, which I'll talk about in a bit, which are basically programs that condition the assistance they give you on meeting certain criteria-- often it's like human capital investments for your kids like sending your kids to school or getting them health checks-- those were in 64 countries as of 2014 and increasing rapidly.

And the number of countries with like at least one different type of social safety net program is more or less all of them. Not all, but like a lot.

So these programs I think are becoming really common throughout the developing world. And they cover a lot of people. So for example, here are some big CCT programs, control cash programs. Bolsa Familia in Brazil covers 49 million people. Prospera in Mexico is 26 million, and so on and so forth.

Some big unconditional cash transfers, this program in China, the Di-Bau program, 75 million people. Indonesia, this BLSM, this is a cash transfer program, 61 million, and so on and so forth.

So these programs are large. I won't go through this detail. School feeding programs, public works programs.

Actually, this program, which I'll mention a little bit, this is a works guarantee program in India which some of you may have heard of, the Mahatma Gandhi National Rural Employment Guarantee, often either Act or Scheme, so either with an S or an A at the end of it depending on who's talking, this is 58 million people. It's gone up. Now it's over 100 million people are covered by this program every year.

So these are really large programs that are affecting, in many cases, tens or hundreds of millions of people.

So what do we want to know about these programs in a developing world context? So one question that I think is really interesting is, well, who should get these programs, and how do we think about who the beneficiaries of these programs should be?

So should they be universal? For example, there's a lot of talk about universal basic income. Should these programs be universal? Should they be targeted somehow at some set of people? And how do we think about those targeting issues?

By the way, universal programs can also be redistributive because if you tax proportional to income and give everyone a lump sum, then on net, yes, Bill Gates gets his check, but on net, he's paying a lot more into the system than he's getting out of it. So even universal systems can be redistributive.

So how do we think about that? And then conditional on doing a program, what are the different things we should be thinking about in terms of the trade-offs or forms that it should take? Should we be thinking about conditional cash transfers or unconditional cash transfers, worse guarantee programs, et cetera? Straight cash? Like how do we think about some of those challenges?

And I think in general, this is an area that has gotten a lot of attention, I would say, in the past five to 10 years, which I think is exciting, but I think there's still a lot of really important questions here as well. So this is a lecture that looks very different from, say, when I first taught this class some years ago.

OK, so I'm going to start, I guess, today by talking about some of these targeting questions, and that's related to the paper that you guys read today, which was one example of which, which I'll talk about either at the end of this lecture or in depth next time.

So the basic challenge is we don't really know who's poor. The government doesn't know who's poor. And this is actually not just a development challenge. It's a challenge everywhere, by which I mean, if you were to take a public finance class and think about the optimal income tax-- how many of you have seen the optimal income tax literature? Not many of you. Well, you'll see it someday perhaps.

So the optimal income tax literature is basically-- I'm not going to go through that model in detail, but this is the Mirrlees model. This is Mirrlees '71 or there's a nice Saez paper in 2001, which lays this out. So the basic framing of the issue in, say, the US is to say, what I would like to do is target based on your earning ability. OK?

So people have different earning abilities, and I would like to give assistance to people with low earning ability and tax people with higher earning ability. But the problem is I don't observe your earning ability. I observe your income. And therefore, my actual tax instrument is to set a tax schedule as a function of your income.

And kind of the whole name of the game in the optimal income tax literature is to say, what do I think about people's labor supply elasticities, and changing marginal tax rates, how do I think about setting a whole optimal income tax schedule to achieve whatever amount of redistribution I want, taking into account the fact that, as I raise the marginal income tax of various people, they're going to work less. And so the optimal income tax is making those trade offs.

And so the original idea goes back to, I think, Mirrlees '71, and then the Saez 2001 paper is a really nice example of actually if you actually estimate people's labor supply elasticities at different parts of the income distribution, then how do you actually take that and like turn that into an optimal income tax schedule? So that's the idea.

But the basic issue is, we'd like to target earning ability, we observe your income, and we write an income tax schedule as a function of income to figure that out. And that, to first order, roughly speaking, fits the design of a lot of the targeting issues in, say, the US, for example.

So we have some various schemes. Food stamps, the SNAP program, or Medicaid, or whatever. The earning income tax credits a little more complicated, but the child tax credit, for example, that the Biden administration has been talking about now or whatever.

And all those schemes basically are conditioned on income in some way. So the government observes your income through the tax schedule. If your income is below some bar, you get the whole thing. Maybe there's some range in which we phase it out gradually, and then the high income people don't get it. Roughly speaking it looks like that.

And there's a huge amount of literature which worries about the idea that if I make the phase out range too steep of the benefits, that's essentially imposing a large tax on people during the phase out range, because if I earn a little bit more, I lose all these benefits. And so how do we think about getting the design of that right so that we can on net transfer [INAUDIBLE] to the poor and not discourage work too much?

And in some sense, in the US, for example there was a big debate about welfare reform in the 1990s. It was all about changing that schedule of benefits to eliminate kind of very, very, very high phase out implicit tax schedules in the phase out regime.

End of my lecture on US public finance issues. If you want more, go take public finance class.

The reason I spent three minutes outlining that is the challenge in a developing country is very different because we don't observe anybody's income. And we can do additional things. If we have additional information that predicts poverty, we can tag these individuals, and so on and so forth.

In a developing country context, we can't do any of this because we don't observe people's incomes. And I will talk about that a little more in the tax lecture, but there's a nice paper by Jensen, which I'll mention briefly in the tax lecture, which basically says, look, for a variety of perfectly good reasons, most developing countries essentially don't tax people who are not employed in the formal sector, which is like in the top, roughly speaking, because basically you're not going to get any information of them without third party reporting. And I'll about that in the tax lecture. But what that basically means is they don't even try.

Most countries have a threshold, below which you're exempt from income tax. And in many developing countries, that threshold is set at like the 80th percentile of the income distribution or the 85th percentile the income distribution or something really, really high.

So what that means is, if I'm trying to redistribute between, say, like the 80th percentile. If I want to make distinctions between, say, who is the 60th percentile person in my economy, and who is the 10th percentile person in the economy, using tax data or your reported income to do that is just completely hopeless. So I need to do something else. So that's what I get to talk about now.

So what are the something else you can do? There are basically three approaches.

So one approach is you can change the prices of particular goods. So you can subsidize goods that are particularly going to be consumed by the poor. A second thing you can do is you can do a universal program and just distribute the same thing to everybody and just hope it's going to be kind of redistributive through the argument I made before.

Or you can try to solve this targeting problem anyway somehow and figure out ways of differentiating between, say, the 80th percentile person in the 10th percentile person even despite the fact that you don't observe their income.

OK, so I'm going to skip this, but let me just note that some of you may have heard of like poverty rates or poverty headcounts or whatever. I'm just going to note that there's metrics of this. You can look at them if this comes up.

So how do we think about this problem? So let's assume for the moment that we cannot directly identify poor households. There's a nice paper, which I used to teach in more detail but I'm not going to through here. I'm just going to mention briefly by Besley and Kanbur which is, well, how do we think about subsidy policies and price-changing policies in terms of poverty reductions?

And at some level, the reason that prices are really attractive is that it's much easier for a government to think about the price than to do stuff at an individual level. Because if I'm the government of wherever and I want to control the price of some commodities, say rice or fuel or whatever, I can just kind of intervene at the market level.

I can import rice externally and I can sell it into the domestic market up until the point which the price is down. Or I can buy it from farmers and sell it to households, and as long as I can prevent loop where they don't buy it and sell it back to me, which would be like a money machine, if I can stop the money machine, I can just announce I'm going to buy it directly from farmers at this price and I'm going to sell it in the urban or whatever markets to people at this price.

And I don't have to worry about who's who. It's much easier for me to affect prices than to keep track of people and affect quantities. So prices have some traction as a logistically not simple, but simpler way to solving problems.

So the Besley and Kanbur paper basically just go through, well, how do we think about all of these things? How do we think about different kinds of subsidies, and so on and so forth?

And one issue in particular they mention that comes up a lot is if you're changing the market price and you have both producers and consumers of this good, then you have to worry about who the consumers are and who the producers are.

So for example, if you're Indonesia, for example, and you're trying to think about what the optimal price of rice is - like rice is both the most important kind of agricultural production good from rural households, and the most important staple commodity consumed by people--

So you move the rice price up and that's really great for the farmers and bad for the consumers. You move it down, it's bad for the farmers and good for the consumers. What is the optimal rice price? It's like a really heated question of debate exactly because of this issue that it affects both sides. So the Besley and Kanbur paper thinks this through.

And the other general point is, what are the kinds of goods that you want to subsidize? You want to subsidize inferior goods. Right? If your goal is to redistribute, we'd like to subsidize things for which people are going to consume less of them as they become rich.

So that actually could have very nice targeting properties. If I can find some inferior goods, I can subsidize them. The poor are going to consume them, and they're going to get the subsidy. The rich are going to say I don't want that good. I want some different good. And it's going to all work out.

The problem with that is that in practice, it ends up being politically not very popular that both other higher middle class people say, well, we would like the subsidy, too. And number two, I think this idea of we're only going to subsidize low quality food or food that people don't really like or whatever, it just doesn't seem very appealing to people. And for I think very intuitive reasons.

And so in the end, what ends up happening a lot of the time is countries don't subsidize inferior goods. They subsidize normal goods. And the problem with subsidizing normal goods is that they're very inefficient from a redistribution perspective.

Yes, the poor are going to get some of your subsidy from the subsidized gasoline, for example, or subsidized rice or whatever, you subsidized bread or whatever it is, but so are the middle class people. And maybe if it's actually a normal good, the middle class people are going to get more subsidy. So that's kind of a problem.

This is a bit of a long introduction wind up to talk about, but they're giving you a bit of the landscape here. So why are price subsidies worse in general? Well, number one, if they're subsidizing normal goods, they're not very efficient for redistribution perspective. Number two, they also change the market price of that good, so they're going to distort total consumption. So they're inefficient because they're changing people's consumption decisions.

People might really like to eat wheat, but I subsidize the rice, and so they eat more rice than they would want to and that's inefficient because of the subsidy. Or I subsidize gasoline and people drive really inefficient cars. More inefficient than they would be if they were facing market prices.

OK, nevertheless, subsidies are still really, really prevalent. So this is from that same report, which basically argues that half of countries are spending more on subsidies than on targeted programs. So this is still actually kind of like a big thing in a lot of contexts.

OK, so number one thing I could do is I can just subsidize some stuff, but A, I have to choose inferior goods if I want it to be redistributive and I have to worry about the redistribution costs.

A second thing you could do is you could give transfers to everybody. And I think there's been a big argument for this idea that we should give out universal transfers, essentially to answer to basically to everyone.

And how do we think about that? So they're good that they don't have price effects. They're not distorting some price. They don't seem to have actually substantial labor supply effects, because they're not changing your marginal tax rate. So good reason to think so. And the income effect of these programs seem to be small. And we have a paper showing that. And they're comparatively simple.

But they basically need two things to function. The first thing they need is they need a system of unique identifiers to make sure that everybody gets the program exactly once.

So if I'm trying to give out a universal transfer to everybody, now my problem is the government is a lot harder. I not just intervening at a few kind of markets. I have to keep track of, OK, there's like one point whatever billion people in India and who are they and let me make sure that everyone gets their transfer once and exactly once.

So to do that, you need to develop a system of de-duplicated universal identifiers. So you need some ID system to make sure that every single person has an ID so they can access the benefits, but people have one and only one ID, and you need some way to authenticate who's who.

So how do I make sure that I don't get to claim my transfer once. Well, I need to know that I'm Ben Olken, and that I'm not also Esther Duflo. Right? So we can't be both of these people. Otherwise I'd get to claim it probably multiple times. And so we need to make sure that we have some kind of way of identifying this.

So countries have been figuring this out over the last 10 years or so, 10 to 20 years or so. And there's been a big effort to create these national systems of universal identifiers. The Indian one is famous, but many other countries have been trying to do this as well. Pakistan has had one for years, other countries throughout the developing world have been investing a lot of effort into trying to create this infrastructure for doing stuff.

And I think that the challenge is that if you don't get this system right, it creates all kinds of later downstream problems. And so I have some recent work on some of the challenges that happens in Indonesia if you don't have the underlying administrative infrastructure. If it's not perfect, it can create other downstream problems.

There's a nice paper in India that shows this and so on and so forth. But I also want to mention this, because this is actually how do you think about the details of how do you actually develop this kind of underlying infrastructure of state capacity for doing these kind of transfers is actually an important area that people are actively looking into right now.

The second thing you need is you need a mechanism to actually hand out the cash that works everywhere, including in really rural, remote areas. So I want to distribute a universal cash transfer to everybody in some country. I need a mechanism by which the government can say every month or every three months or whatever, distribute a certain number of cash in a cost efficient way to every single person in the country.

And that's also challenging. And some countries have made a lot of progress with e-money, for example, is one way of doing this. Other countries have actually worked to stand up big, complicated systems of distributing money through post offices or whatever.

But I just want to mention that this is actually also not trivial, and I think understanding the role of governments in how do you actually hand out the cash? And it's interesting.

And there's some other recent work that says that actually the government developing mechanisms to handle this problem can actually have spillover effects for the rest of the economy. So actually like if the government actually says, OK, well, we need to set up a whole infrastructure to hand out payments to people, so we have to get them all debit cards or we have to set up a system where they can cash in cash debit cards, that actually may have really important implications for catalyzing the financial system more generally.

And so there's a paper from Mexico that shows us we're actually working and related things in Indonesia on this right now.

So my point there is that the logistics, the details of how do you actually do this are not trivial and actually end up having a lot of interesting issues that are actually a really I think exciting current area of research. And of a lot of papers seem to be coming out on this stuff right now.

The other question, though, is how do we think about this stuff from a redistributive perspective? So are these things redistributive? And how do we think about that?

So I want to make a very simple point, which is how does this interact with the point I made before about who's in the tax net. So suppose people's pre-tax income is y . We can define an after-tax-and-transfer income as α of y . That just says after I tax you and give you all the benefits, I earn y , but I end up with α of y . That's what I can consume or whatever.

So what is that going to look like? First of all, any tax and transfer scheme that has α of zero greater than zero, we can think of that as having a universal basic income. OK, so one kind of confusion about universal basic income is like, yes, the government sends you a check in the universal basic income, but it also taxes you back. So all we really care about is the net.

So any system that has alpha of zero greater than zero can be thought of as featuring a UBI. Alpha of zero is I earn nothing, I get the UBI. And then everything after that, then the government starts taxing and taxing back whatever I'm doing.

And Saez basically thinks about this in the US context, and he makes the point that there are different lessons in labor supply and when the intensive labor supply elasticity of how much I work is larger than the extent of whether I work or not, he argues that the UBI can be optimal. I'll refer you to that paper if you're interested.

But the point I want to make is in a developing country context, we have another issue, and the issue is that, as I mentioned, most people don't pay taxes. So what that means is, if I introduce this thing where alpha of zero is greater than zero, I'm giving the same net transfer much further up the income distribution.

OK, so for example, this is a graph that I drew in PowerPoint. So this is not actually real data. This is like sketching curves. But what does this mean? So imagine that this is my pre-tax income this is my post-tax income, and this is a tax schedule. A tax schedule is just alpha of y . What is my post-tax income as a function of my pre-tax income?

So this is like nice and progressive or whatever. This feature is concave, so it is increasing marginal tax rates in this example. So suppose I introduce a UBI, but I want to keep the total government revenue the same.

So that means I'm going to increase the alpha zero of zero-- so people who earn nothing now get something-- but I have to raise the same amount of government revenue as before. So that means I'm going to have to increase marginal tax rates elsewhere to pay for this.

So maybe it looks like this fine. So that's the trade off. So these people are getting higher net incomes under the UBI scheme, and these people, even though they're getting the check, are getting lower net incomes on net because they're facing higher marginal tax rates.

So what happens now if I were to do this in, say, India, where-- I don't know the exact number, but suppose it was like 80% of the people are outside the tax net? Well, then this looks different, because I have to give the same net transfer much, much higher up in the income distribution, because I don't have this additional dimension at least in the income tax schedule of redistributing people.

Now if you have multiple tax instruments, maybe you can tax some of it back through the value added, through other taxes like value-added taxes or consumption taxes so it's not quite as stark as this.

But the point I wanted to make is that, in a place with a really big tax exempt region, if I introduce these net transfers, the net transfer from a UBI is the same as the gross transfer from UBI. And that's not true, say, in a place like the US where I have much more flexibility over the tax schedule I can fine-tune that one.

So what that means is that in practice, UBIs are more expensive in some sense. For a given dollar of a UBI, it's more expensive in some sense, because you have to give the same net transfer much higher up in the developing country income distribution than you would in a developed country income distribution.

So that's the point I want to make about UBIs. So then actually, we can simulate that. And I'll come back to that later.

OK, questions? I've kind of been talking for like the last 20 minutes. Is there anybody with any questions or any comments? OK.

So the third thing you can do is you can try to differentiate among the poor somehow, or differentiate between people for whom you have imperfect information somehow. And there are a bunch of options that you can do to do this.

So one is you can do a proxy-means test, which I'll talk about in a second. You can do community-based targeting. So proxy-means test basically says, let's try to get imperfect data and use that to predict people's incomes and target base on that.

Community-based targeting says maybe there's local information about people's income levels and maybe we can try to use that to differentiate people. Maybe we can do some version of self-targeting or we dual mechanisms where we can try to create systems by which people will have an incentive to reveal their information.

And the paper you guys read for today is one example of that. Combining it with proximity-means test, actually, but like the idea was thinking about people have this private information. Can the government find a way to get them to reveal it?

You can also do broader things like geographic targeting. I'm not going to go into that right now, but I want to talk about these three options. OK, questions?

All right, so what's option one? So option one is proximity tests. There's another option also, which is regular-means tests, which is ask people their income and you can't really verify it and you hope they're not lying.

And some countries do that too, actually. So I just want to put that out there as well. But that has the obvious problem, which is that if I can't check your income, then you might just lie and that might not be a very efficient program. But some countries actually do have that approach.

OK, so what's the idea? So the government doesn't observe your income because you're in the informal sector and you're not paying taxes and you're below the threshold, so it's not even trying to get your income there.

You might think, I've just seen probably a bajillion economics papers where we have surveys of people's income or consumption. Why can't we use that to target? And the answer is, first of all, they're very expensive, but second of all, they're totally non-verifiable.

So if you've ever done a consumption survey, what is the consumption survey? Consumption survey says, I go to Seana and I say, Seana, please tell me how much meat your family ate in the last week, and how much rice did your family eat in the last week, and how much fish did your family eat, and so and so forth.

You could just make that stuff up. There's no verification of any of that. It's just kind of your recall. And that's actually not a problem in research purposes. Like if I'm doing an anonymous survey that's going to go to the University of Wherever and it's not going to be for anything, you have no particular reason to lie on that survey and most people will be honest and tell the truth and do that thing.

But if I tell you that actually I'm going to use this to distribute transfers and people with low reported consumption are getting a lot of money from the government, lo and behold, we might expect that everyone is going to not report a lot of consumption. They're not verifiable.

So instead, what the government does is it goes around door to door to everybody or some large fraction of everybody, and does a census where it collects information on things that are on observable characteristics that are correlated with your income or consumption levels.

So they collect data on your land, how much land you have? They look at a bunch of stuff about your house. What does your house look like? Do you have a fancy roof? How big is your house? What are your walls made of? What's your floor made of? Do you have a TV? And so on and so forth.

And the idea is these are things which are much easier to observe. Like if the government census taker is standing at your front door looking around, these are reasonably observable characteristics. So that's the theory. And by the way, the things they focus on are things that seem to be correlated largely with your permanent level of income.

And so what the government does, so they have this asset data with, say, 40 different variables of what is your house type, and what is your roof type, and what is your floor type, and so on and so forth. And then they have a census for everybody.

Then they also have survey data that they collect in some low stakes environment-- so like in some research kind of environment-- that has the full consumption data and the same set of assets. And they basically run a regression where they use the survey data to predict people's income based on the assets-- people's income or consumption based on the assets.

They get some data from that regression. They go back to the census data. They get \hat{y} instead of y , so they take all those assets and predict your income, and they target your consumption. Your program is based on \hat{y} . So that is what they do. So it's regressions in real life.

And you can do a straightforward OLS regression if you want to predict consumption. This is a prediction problem. Machine learning might be helpful there. I don't actually think it makes an enormous difference. In most of the contexts I have seen, it helps a little bit. I think, actually, it's been a little disappointing in my sense as people had hopes it would be wildly better. It's like a little bit better, but I think that's basically what people do.

Yeah?

AUDIENCE: But I think that these surveys can also be calcified. At least the problem that we have in Chile is that people already know that the people that are going to survey look for, so they hide their microwaves or their TVs because they know when they're going to get surveyed, so it's not that accurate.

**BENJAMIN
OLKEN:** Correct. So actually, I'm going to talk about that issue in a second. But yes. So it's not perfect. And exactly, people actually might try to cover the TV or whatever. So what do countries do to deal with that?

Number one, they keep the formula a secret. So it's hard to know exactly how much do I have to cheat or what do I have to do exactly to get on the program? You don't know, because the betas are secret.

And even if the betas were not secret-- it's kind of complicated-- they might change the formula every couple of years. So you might have thought we had a secret algorithm of, oh, if I hide the TV and the microwave then I'm good to go, but then they'll change the formula again a little bit so it won't work again in three years.

Yeah, so that's what people do. And there's a really nice paper from Columbia about this exact issue by Camacho and Conover where they get the predicted scores and they plot them. So this is x prime beta. This is like your predicted poverty score-- or y hat, as I was saying.

And the first year they do it, it looks totally normally distributed. And then over time, you start to see it look a little bit more like this. Right? And the paper is like one graph. Like figure 1 is showing how this thing emerges over time. So that's exactly your phenomena.

I think what happened was that that was a case where I think they argued that the formula got out or whatever, and people started to learn oh, what I need to do is I need to tell my enumerators exactly here are the sets of things to write down. And so I think that a lot of that has been combated by changing the formula.

By the way, none of this is perfect. We're in the world of the imperfect. But I think that if you're interested in that issue, that paper is the best one I know of, and it's by Camacho and Conover. But I'm not sure how broad of an issue that is more generally.

OK, other questions?

So what are the problems? Number two is the one that Christine just mentioned. Number one, the biggest problem I actually think is the R squared is much less than 1.

So this is an example of what this thing looks like. So we just replicated this in some data from Indonesia and Peru. So this is y hat, predicted per capita consumption, and this is y , actual measured per capita consumption from the survey from doing this whole exercise.

And what do you take away from this? On the one hand, if this was like most regressions you would run on your computer for a paper, you'd be pretty pleased if you've got an upward slope that looked like this. Like, it certainly has a lot of predictive power. That's good.

The problem is, it also looks like every regression you've ever run, which is like it's not a perfect line. And if you think about it that way, it's a regression. No regression you ever run has an r square of one. There's always kind of noise in the world.

But noise in the world in this case has real consequences. And in particular, what is this program going to do? Well, suppose I'm distributing a transfer. What I would like to do is give transfers to everyone below this line. Because this is actual. What I do in fact is I give transfers to everyone to the left of this line, who has y hat below this level.

So these were the people who are correctly included in the program, and these are the people who are correctly excluded from the program. And these are the errors. So this is inclusion error, which is to say that these people get the program, even though they're actually too rich to get the program based on the truth, and these are the people that are excluded from the program who should get the program based on their actual consumption levels and don't.

And so I think once you frame it this way, it's clear that any kind of regression-based approach is going to have errors like this. We've never seen a graph where this thing is like exactly on the line, and nobody hits and falls in these kind of quadrants.

But what you'd like to do is try to find a system to minimize these two types of errors, because these are really costly errors. And so you would like to work to find ways of minimizing exclusion error and minimizing exclusion error, so that you're really delivering the transfers to the people that you actually want them to be.

I'll make one other note, which is that it's probably not quite as bad as these graphs make them out to be, because this is actual per capita consumption as measured in a survey. And surveys are going to have measurement error.

And so it is worth pointing out that probably the truth doesn't look quite as bad as this, because if I took actual consumption error and just added some noise to it, that would already be just noise in the measurement process. Like if you think about my example of how do I measure consumption, I was going and going through the 300 questions about how much rice did you have last week.

And if you ever doubt how much measurement error there is in a consumption survey, you should try to take one yourself and see how much you remember of how much of what you spent on each of 100 different types of food products over the last week and each of 100 different types of other consumption over the last month and see how good you are at recalling that. And so I think that those things really are measured with error.

Classical measurement error is actually not a problem for this y hat thing, because it's not going to bias our coefficients in that regression. So that will be unbiased, but this thing will look worse than actually it really is, because the actual kind of measurement error.

The other thing I'll point out here is that this creates errors. So it's inefficient. It also creates equity challenges. So I'm sitting in my house with suppose I have a dirt floor but I have a TV and I have maybe like a thatched roof. And maybe Aaron doesn't have a TV, but he's got a nicer floor and maybe has a tin roof, and like I get the program and he doesn't.

And we don't know what's going on. And in fact, the government didn't tell you what the betas were, for the corruption reasons I was talking about. They don't want to publicize the betas, so they just said like, bing, you get it, you don't.

So that, number one, creates real horizontal equity challenges. It also creates I think some transparency equity challenges, because we don't really understand how the government arrived with a list of who got it and who didn't get it.

So I think these are real costs, actually, and I think these are really serious challenges. I think the lesson of this, by the way, is going to be nothing is perfect and there are all trade offs. But those are some of the challenges. Yeah.

AUDIENCE: So in this graph, but also in the paper we had to read, we were thinking about there is like a cut off and there are people that we want them and [INAUDIBLE] But if I look at that graph, like the people that were accidentally included are not that far above the threshold. So like recently they're--

BENJAMIN Correct.

OLKEN:

AUDIENCE: Poor.

BENJAMIN 100%.

OLKEN:

AUDIENCE: So how should we think about optimizing this way when there's like [INAUDIBLE] poverty and I don't mind the people that are slipping and are just above.

BENJAMIN So that's a great question, and it is one that I think is super duper important. And I think that there's-- I drew these graphs in this way because I think the policy community often thinks about the inclusion and exclusion error.

AUDIENCE: Yes.

BENJAMIN But I think you're completely right that, from a welfare perspective, for the people right over here, who cares?
OLKEN: Like they're approximately the same. From what I mean by who cares, I mean if you were to write down a welfare function, this person and this person have approximately the same utility. We shouldn't be discretizing it and saying this is awesome and this is terrible. It should be continuous.

And so actually in this paper, we do it with a continuous function, which I think I have my slides later. So basically, what I think the right thing to do is to write down a utility function, then you've got to pick a curvature parameter, so you pick some curvature parameter that you think is broadly sensible or some other way, which again I'll talk about some more of this, but I think the right answer is you should do it continuously and you should take that into account a little more seriously.

And that's what we do in the paper. Because I agree with you. This is way too discrete. I think that a lot of these measures of like poverty, a lot of them I think people in a lot of contexts find it very easy to think about people who are above or below the poverty line or what is the fraction of people who are poor or whatever, which are easier concepts to think about than like I'm going to run a welfare metric with CRA utility function and rho equals three.

So like, I think that that's why these measures get a lot of attention, but I think that conceptually, you're 100% right. Yeah, sorry.

AUDIENCE: I'm just going to ask, are there any evidence of papers talking about the corrupt change where they did to this parliament aiming to [INAUDIBLE]

BENJAMIN Yeah, there are a few. There's not a huge number. There's the one I mentioned before. There's one called
OLKEN: Targeting with Agents, which is about this issue, which you should look at. There are a few more. In the context I have looked at, it has been less of an issue than I would have expected, I would say.

AUDIENCE: They're so poor [INAUDIBLE]

BENJAMIN

OLKEN:

Well, part of it is if the agents don't know the formula, it's actually a pretty tricky problem. So if I'm the agent, and I think that was the challenge in that Columbia example. It was that the formula got out. Once I know the formula, then I can do this kind of efficiently.

If I don't know the formula, then the corruption problem is kind of tricky because I could say, well, if I'm the person who's doing the thing, I could write you down as like all zeros. I have no assets. That's going to look kind of suspicious. That's going to be very easy for them to get caught.

So they don't want to distort your things too much, because then they're going to get caught. On the other hand, if they don't know what to do, it's very hard to have a contract because you could pay them to distort your stuff, but they'll distort it a little bit, maybe you won't get the program. So I think kind of the opaqueness about exactly how the formula works makes the corruption problem actually a lot harder, in my sense.

And in fact, I will say as a little anecdotal, how does the research process kind of happen? When we started actually working on some of this stuff in Indonesia, that was one of the things we were really worried about.

And actually we had a whole thing of like maybe we were going to define like really simple kind of easily-verifiable things that kind of weren't going to be easy for other people to cross check. And we had a whole plan about that. And then in the end, we ended up concluding that, at least in that context, that was not the first order issue. There were other issues that were more first order. And I think that for kind of the reason I mentioned.

OK, other questions?

All right, so this is positive r square imperfect. So what else can you do? So the first thing that we investigated really seriously was this idea of using communities to make these decisions. And so allowing communities to identify local households.

And I think this is an idea that comes up in a bunch of different contexts in development economics, is that there is a lot of local information about people. People seem to a lot about stuff, so this is going to be a do they know who's poor and who's rich and can we use that for targeting? There's another paper, which I think I'll talk about in the credit section, which is do we know who's like a good entrepreneur? Who has a lot of entrepreneurial ability who would benefit a lot from a loan?

And the whole kind of stuff on informal insurance, which we'll talk about in the next semester, there's a whole literature about do we know who's had shocks and how do we condition transfers based on shocks? There's a whole series of themes where a common theme I think is that if the formal institutions aren't working as well but maybe there's this local information that can actually take their place.

So that's what we wanted to investigate here. And the problem, of course, is that if I actually try to use this information, it may be all well and good like in a survey if I say, well, please here are 10 people in the village, can you rank them from richest the poorest people we actually know. Actually, we do do that exercise in a survey in this context.

And people do seem to have actually pretty good information. The question is, if I tell you I'm going to use that to give out cash transfers, what do I get back from that process? Do I get back the people who are actually poor, or do I get back the village head's uncle or some other thing? What actually comes out of that process? Can I actually get that information out in kind of a high stakes capacity?

So that was our perspective as we went into this project was that there was this trade-off, potentially, between maybe there's better local information, but there's this kind of risk of local capture whereas in the data-driven approach, maybe that's not going to happen as much.

And I think there was some existing evidence when we started this project that-- I was saying we because this is now talking about a couple of papers that I was involved in-- that communities do know more. So these papers from here. So how can we actually test this?

So what we did is we did a randomized trial where we compared three different targeting methods. One was a proxy-means test, which we basically had the government go out and do their kind of data-driven thing. We did a community ranking exercise we'll talk about in a sec.

And we did a hybrid method, and the idea of the hybrid was to say, well, if we're really worried about a leak capture, if we think there's local information but there's a leak capture, what should we do? Well, we should get the community information and then we're on a data screen on this. And we'll basically use the kind of data screen to throw out people who look really rich.

OK so that was the plan. And we worked in Indonesia. We had about 600 villages that were randomized into these three schemes.

And the other thing which we did, which I want to mention a bit as an experimental design issue, we had a whole bunch of sub treatments to tease out kind of why the community and the PMT were different. So for example, we were worried about a leak capture.

So we said in these community meetings I'll describe in a second, should we invite the whole village to these meetings or the whole neighborhood, or should we just have local elites divvy out the stuff alone?

We thought this process is actually kind of hard maybe people get tired of doing this ranking. And so we actually had the kind of community rankings done in random order, which I'll talk about in a sec. We thought preferences might matter.

If we're having a community meeting make these decisions, who comes to that meeting could be really important. So we randomized the time of the day, because we know that if we hold a meeting in the afternoon, more women will show up and if you hold a meeting in the evening, more men will show up in this context. So if I randomized the time of the day of the meeting, we could change the composition of who was going to be there.

And the final thing is, this is a paper where the timing is backwards of what you normally do. We do a survey at baseline. The key data actually is a survey a baseline to measure people's true consumption in a low-stakes setting, then we do the targeting and they see who they picked. And then we do an end line actually to measure satisfaction.

So a couple other things, just important kind of research-related tidbits. One is it is really important to pilot what you are doing. So in the end, what we ended up doing was this is a system where basically how we do the community targeting?

We put up a string across the room. They first had a list of everyone in the neighborhood. They then kind of brainstormed a list of what they meant by local poverty in this context. They could actually do stuff. Then they picked up cards kind of one after another in random order and they said,

OK, first we're going to do Ben versus Abhijit. Who do you think is better off in this village? And maybe we all say one of the two of us, maybe you say Abhijit. So Abhijit gets put to the left or the right of Ben on the string.

Then we pick up the next card and say now we're going to do Glenn Alison. Where do you put Glenn Allison on this list? And so on and so forth until the whole string is kind of full with people with a complete rank ordering of peoples from most well-off to least well-off.

So this worked actually pretty well in terms of actually succeeding and actually getting to the end of this process without a lot of disagreement or discussion. I will say when we piloted this, the one thing we also thought could work was we thought, well, maybe this whole ranking thing is really complicated. If I've got 100 people ranking them exactly, it's really tricky.

Maybe it'd be easier if I first had the group people into bins, four quartiles, say. So like let's do a quick sort first, bin 1, bin 2, bin 3, bin 4, and then we'll do it after that. And so actually that follows some sorting algorithms that have that feature.

That, it turned out, was a terrible idea, because everyone immediately figured out that if I want to get the transfer, I have to get my name and bin 1. And so there was a lot of contention in that process because everyone was like, no, no, no, I want to be in bin 1.

Whereas this one, it was all continuous and we didn't exactly know what the costs were going to be, and you don't know where the cuts are going to be when you're doing all the rankings because we don't know where everyone else is going to come last.

So this one turned out that by not discretizing it that it was a totally fluid process that we could do of hundreds of times. That same process, but just changing the mechanics of how you do the sort. And first, having people put into bins did not work.

And even this one, actually, was only one of I think we had like six or eight things that we piloted in six or eight different villages. The first version of this one, actually, we did it on the floor. And then it was kind of we had to move all the cards around, and so the interesting one came up with the string idea that actually worked a lot better. So my point is, if we had just had this initial idea of oh, we're going to the exact same design, but had gone to the field just with our initial version.

We might have had the four bins, and the thing would have totally not worked. So I think just like before you go to the scale, we were actually kind of testing things and trying to evaluate if they work or not. Just making sure that each of the things you're doing is actually kind of feasible in your setting is really important. And I just want to emphasize that. Yeah.

AUDIENCE: So like, do you usually have one person on the ground who can tell you in advance or help you determine some very local cultural--

**BENJAMIN
OLKEN:** You do it yourself. I mean, so yes, you should always be working with people. So on the one hand, yes, local collaborators are great and can certainly help you with this stuff. And we had local collaborators in this project.

Number two, in this case, we had a research assistants who were local who also could help us. But number three, we went to a lot of these meetings ourselves. And there's no substitute for that process as well. So the answer is yes, but the local collaborators were involved, too, and they thought that the bin thing was a good idea.

So there's no substitute for trying it, yourself. And I mean, obviously knowing the college's local context is important. But I also think like as researchers we also I always when possible. I always try to go to at least some of my pilots for projects because you see stuff that is just very hard to replicate.

The other kind of experimental design thing I want to say-- and maybe I'll end here because we're almost out of time and I'll pick up with the results next time-- is this is what the design matrix look like. So what did we do here?

So I told you we had these three main treatments, community targeting, this hybrid, and the proxy-means test. But then we had a whole bunch of additional things we wanted to learn. And so what we did is we cross-randomized a bunch of different stuff.

So the idea is none of these individual cells was designed to be analyzed on their own. Right? But if I want to say, on average, does it matter if we have the elites running these meetings or the non-elites running the meetings? I can compare this group to this group.

Or does it matter if we have more women in the meeting versus more men in the meeting? I can pool all the days together with all the nights together like this. And so this kind of cross-randomized design can be really helpful insofar as you can learn a lot more stuff from your given 600 observations.

The downside of this is that what you are learning is kind of a weighted average. So when I compare community to hybrid, what I'm learning is, what is the effect of the community treatment with a weighted average of all this other different kind of stuff?

And so that is what it is. But in many cases, I have found that that is a trade worth making. And so if you're ever doing an experiment, you should think through are there additional these cross-cutting of randomizations I want to do, which I don't think are going to affect my overall interpretation of the main thing I'm doing, but could help me learn a lot more from the same set of data points.

They're not without cost. It's complicated, because now if I'm of working with a team of facilitators that have to go run these meetings. There's how many different versions? There's 16 different versions in this table. So they have to keep track of the 16 different versions and which one I'm doing in which place, and make sure they follow the randomization.

So it is not without costs, but relative to the grand scheme of things, you can learn a lot from your same data points with kind of approach. So I wanted to think that through as well.

So I'll stop here. I will go through the answer. What did we actually learn from this paper and how do we evaluate it? And then the [INAUDIBLE] paper, I'll pick up with that at the beginning of next class on Monday. OK, thanks.