[SQUEAKING]

[RUSTLING]

[CLICKING]

BEN OLKEN: OK. So where are we? So we're just finishing up the last couple of things I wanted to talk about redistribution programs and transfer programs. And then I'm going to talk about tax, OK? All right.

So did I talk about cash versus-- did I talk about the conceptual versus supply shocks, and man shocks, and all that stuff? OK, great. But I didn't show you the results, is that right?

That's right? OK, right. So just to remind you, so yeah. I think this is what we went through last time, right? So the point about these cash versus in-kind programs is one thing you might think about is that a cash program is a demand, going to be a demand shock, whereas an in-kind program may also be a supply shock, right?

So this is very basic, right? What happens if we do can-- you can actually do it in different ways. It's as simple as-- I should draw it, probably, right?

So this is the supply of whatever the good is. This is the demand of the good. If I give you cash, I'm holding the supply of the good-- if I give you if a cash transfer, that's going to increase the demand.

If I give an in-kind transfer, then the government may not only increase the demand, but also may change the total aggregate supply. And so there may be differences in terms of the price effects from these different programs. Yeah, Ahmed?

AUDIENCE: The system [INAUDIBLE] transferred. Some say [INAUDIBLE] transfer is like rice.

BEN OLKEN: For example.

BEN OLKEN: For example.

AUDIENCE: It has this little holes full of rice and it just send into different provinces or regions. But it can also very well be that our government just buys from a market in that region.

BEN OLKEN: Totally, yes. Exactly. And so you're totally right. If you think about-- it's actually not obvious. Where is it? If you think about housing, so this example works really well if you think about, say, public housing, right?

In public housing, the government goes and builds houses as opposed to just providing vouchers. That's clearly a supply shock, right? I agree with you for something like rice. It's a little more complicated and depends on where is the government sourcing its rice.

It's the cleanest example is they source it outside the country and bring it in. Obviously a supply shock, right? I think in some cases, if they are just going into the super local market and buying it right there, not a supply shock. And so actually, that's related to what they're going to do in this paper.

So they're going to talk about-- they cut this, basically, by the remoteness of the regions. And the idea is that basically, it's related. So the basic idea is in really remote locations, what the government is doing is it probably is a net supply shock. It's sourcing the stuff externally, bringing it into the remote regions.

In less remote regions, some combination of markets are more elastic because they are just thicker. They're thicker markets. So maybe the supply curve is more elastic, right?

If a supply curve, obviously, is totally elastic, then this doesn't matter, right? So in those less remote regions, it could be that. Or it could be that you're buying in those regions and sourcing anyway. So for those two reasons, I think they focus on these more remote locations.

That's a really good point. And just what's the point of this? The point is that if they look at the price of the goods, they find a substantial increase, a substantial change in the price, where the price is-- and by the way, this is relative to a control. Neither of these is significant. But the difference between them is significant.

So they find that prices are lower in the in-kind areas relative to the cash areas in the remote locations. And then that can affect consumption and people's overall well-being from the program because you get less-- in real quantity terms, you get more in the in-kind areas because of these price effects. And that's basically what they argue.

So the in-kind area is people's total consumption is actually higher for that reason. So let me actually not go through the details of this because it's a pretty simple point. But I just want to see if there are any just mention of those two things, which seem like they should be equivalent in principle, but actually have these different reasons for basic price theory reasons. Yeah, Tishara?

AUDIENCE: How do we think about programs where government buys, say, rice at a high price from the producers but essentially gives it out for free [INAUDIBLE]?

BEN OLKEN: What, what's that?

AUDIENCE: So government is buying, say, rice from the producers at a minimum support price, which could be higher than the market price, and then sells it, sends it to the consumers essentially for a very low price or for free.

BEN OLKEN: Yeah, that's still going to have-- it's still going to look like this in the-- oh, sorry. No. That's a different program, sorry.

So you're saying there's another program. I think you're saying-- so this would be the market level. And you're saying the government basically comes in and goes like this, right? I think you think of it like-- if you think of it like that, what's going to happen?

So it's going to raise the price that producers are selling it at. It's going to lower the price people are consuming it at. And it's going to inefficiently increase the amount of the good, or it's going to increase the amount of the good consumed above and beyond the market level.

And so the way you should think about this is like a redistribution program. First of all, we want to think about who the welfare of these people are, right? Are they rich or poor? Is that a good thing to be doing or not? Relative to who are these different people?

Second thing is we should think about, do we actually want more consumption of this good or not, right? So for example, governments often-- and the third thing we want to think about is, who are the people who are induced to buy more of it and get the subsidy? Are they the people you want to be targeting?

So if this is an inferior good, this actually may not look so bad. If this is a normal good, then you're actually giving a lot of the subsidy to richer people. And that might not be so efficient.

And the other thing is, do you actually care that you're distorting the overall consumption of this particular good? This is a good think is being-- suppose you think this is a good that's being inefficiently underconsumed in the market equilibrium. You think maybe people are undernourished, for example, and you want them to consume more of this at the market price-- this looks great.

If this is, I don't know, gasoline or something, where maybe you think that actually there's no particular reason we think we're under consuming gasoline, it's terrible. So I think that's the way I would think about it. But the key difference in that policy is it's changing the marginal price, right?

And that's why it has-- whereas those things are not changing the marginal press. And the [INAUDIBLE] paper works through all this stuff with a poverty-- thinking about heterogeneity in the population, and poverty, and who do we think-- how do we think of all the different agents. And that's a very nice framework if you want to think about that and compare the inframarginal versus marginal transfers.

Other questions? OK. Oh, so the other paper I wanted to mention-- and again, I'm just giving you just a brief flavor. There's an enormous literature, an expanding literature on the design of these programs. But I just wanted to mention a couple of things just to give you a sense of some of the things people are looking at.

So how many of you have heard of conditional cash transfer programs? Maybe most of you, OK, right? So these are programs that are very common now in the developing world.

They condition aid on fulfilling a set of criteria the government thinks are good, like for example, sending your kids to school or getting them health shocks, or whatever. You get the cash grants to do this stuff, OK? So the government is incentivizing these additional behaviors.

And there were several randomized trials that showed that those did increase the desired behavior or the incentivized behavior. Not surprising because it's being incentivized. But it turns out to be there.

And the justification for that is intergenerational, right? The idea is that many parents don't fully internalize the value of their kids' human capital acquisition. And governments can correct for this by incentivizing these behaviors.

What this paper was-- first paper that I know of to, in a really rigorous way, compare what is the impact of the conditionality per se? So those programs before were bundled, where there's overall randomized evaluation of the conditional cash transfer developed from nothing, but didn't say, is this conditioning thing actually a good idea or not? And this paper is nice. It examines conditional and unconditional transfers in the same context.

And so how do we think about that? Well, the first thing you would think about if you condition on something is, of course it's going to change the price, right? That's good if you want incentivize that behavior, right?

But the other point, which I don't think people necessarily-- it's maybe obvious ex post, but wasn't, I think, obvious until people had written this paper-- is if there are people who are really far from the margin who are not going to take up the incentivized behavior, then you cut them off from the cash transfer, right? So a rigidly enforced conditional cash transfer is not going to be providing the transfer to people who can't meet the conditions.

That's the whole point. And therefore, those households don't get the income effects from the transfer. And that may also be important. OK? So how do they show this?

I should say, by the way, there's another paper, which I'm not going to talk about, which is about the labeling effect of the conditional cash transfer, which just says if you tell people the conditional cash transfer it's a conditional cash transfer, but don't actually enforce the conditions, people may think they're enforced and do them anyway. And that paper says a labeled cash transfer versus an actual conditional cash transfer-- you may get a lot of the benefits just from the labeling without this negative cost.

But what do they do in any case? So what do they find in this case? So let's look at this.

So what are some of the things I wanted to note? So the first is let's look at actual enroll-- so they condition on enrolling your kids at school. They have these self-reported enrollment. But then they have teacher-reported enrollment.

And they have some arguments in the paper that the teacher-reported enrollment is the better one. So lo and behold, not super surprising, the conditional cash treatment, which is incentivizing people to go to school, results in more kids going to school. OK? So that's that.

And you can see it for three years or whatever. And those things are statistically significantly different in most of the specifications. And they do better on their test scores, OK?

Fine. So it seems like it works in terms-- the conditionality is important for the incentivized behaviors. On the other hand, they also then looked at early marriage, basically, as another outcome. So early marriage particularly for girls in the household.

And early marriage is not incentivized outcome. But it's something that you might think that preventing early marriage might be something the government might care about in this context even though they're not incentivizing it. And what they find is that actually, the fraction who were ever married in this period is actually lower in the unconditional treatment than in the conditional treatment.

So both of these treatments lower early marriage compared to the control, but it's larger in the unconditional treatment. And what they argue that's going on-- you see it here-- which is if you look among those who are not enrolled in school, so these are the people who would be not getting the-- are endogenous groups, right? So there's fewer people not enrolled in school in the conditional treatment than in the unconditional treatment, right?

This is the share of the group. So that's not surprising. That's the condition that we just saw, that the conditional treatment reduces the number of people enrolled in school-- who are not enrolled, sorry. But there's a big effect of the program on early marriage for the not-enrolled group in the unconditional group, but not in the conditional group.

And why is that? Because these people over here don't get the transfer. This is not-- this cross tab here is not super clean because this is an endogenous variable. There could be samples not shown here.

But I think that this is a pretty natural explanation for-- but these effects are well-identified. And this is a pretty natural explanation for what's going on over here, right? So basically, what happens is there's-- I think the way to think about it is there's a bunch of kids, right?

Some of them are on the margin for going to school or not. And they don't-- in the conditional treatment, some of them switch over here. They don't go to school-- I mean, they go to school, right? But there's another set of kids who are just not on the margin for going to school or not.

And for those kids, the income effect could be really important. And you're taking away the income effect for the kids who are not enrolled in the conditional treatment. And so this basically analyzes this, I think, a little more, basically showing you the same thing it entails.

Actually, maybe not. I don't need to go through that. But I think this is the key insight. Yeah, Ahmed.

AUDIENCE: So what treatments are unconditional--

BEN OLKEN: What?

AUDIENCE: When the treatment is unconditional, it affects a larger number of people. But then does this cost more? So are we comparing-- and is the treatment the same--

BEN OLKEN: It's more expensive. You're right.

AUDIENCE: We are still--

BEN OLKEN: Because you're right. In the conditional treatment, you don't pay the-- well, actually, there's two different effects. Number one is, in practice, actually monitoring all the conditions is actually not free. So you do get a cost savings over here from not having to actually go and do all that monitoring or whatever. So that's one point.

But on the other hand, you're right. You are saving-- quote "saving" unquote. You're not paying the cost of those transfers for the people that are getting cut off, which is 30% of the sample over here.

So you're right. The transfers you're giving out are whatever. Here you're only giving out 70% of the cash compared what you're giving out here.

So you're right. This program is probably a bit cheaper in that sense. I suspect these differences are sufficiently stark that maybe that's not going to do it. But it's a fair point.

Other questions? OK. So I think that was-- I just want to mention that everyone always thinks that the conditions are good. But it's worth keeping those in mind as well. I'm not going to-- I wanted to just put that out there as a pointer-- those two papers, almost, as a pointer to a whole literature, which I don't have the bandwidth here to go through in detail.

But understanding the design-- as countries are doing more of these transfer schemes and redistribution schemes, understanding the design questions involved in doing that properly is a really expanding area that people are thinking through. So for example, some recent work-- we have a paper on, how do you actually deliver the transfers effectively and does transparency that gives people information about what their rights are matter? One of our recent papers says yes on this.

How about technology and reducing corruption? There's a lot of interest in what is the underlying technological infrastructure of these programs? smart cards, biometric authentication, digital payments-- all that stuff is entering. And understanding that has been a really active area of research in the last five years or so.

And actually, even just changing the back-end system-- the nitty gritty of how the government reimburses itself for these programs can be really important. We have some papers on government versus private distribution of this. There's some work on how do you get people enrolled.

There's questions about more stuff on the cash versus in-kind. Should you have transfers be lumpy? Should they be a series of big lump sums or a small series of payments, et cetera, et cetera. There's a number of questions that people are actually working, studying here.

I'm actually in the middle of writing a lit review on this with several co-authors. But I want to mention that if you're interested in this, I think is a really active area, but there's also a lot going on. And so it's worth making sure you get up to speed if you're interested in these topics. But there's a lot of active research on this. OK, any last questions before I switch gears?

And this is general-- oh, sorry. And one last thing. I think a lot of lower and middle income countries-- not the poorest countries in the world, but low middle income countries, are starting to think about social insurance programs now. How do we provide health insurance, disability insurance, weather insurance, crop insurance, those kinds of schemes.

And I suspect-- and there's just beginning to be some research on that. And I suspect that's an important next frontier of how do we understand those challenges in these low information and low governance environments? I think that's-- I just wanted to mention that as an important next frontier of future research.

OK. Any last questions before I switch gears? No? OK. Yeah?

AUDIENCE: How should we compare this spending with infrastructure spending?

BEN OLKEN: I'm sorry, [INAUDIBLE]. You've got to speak up.

AUDIENCE: Sorry. How should we compare this spending with, say, infrastructure spending? Is that possible? Or do we know anything about that?

AUDIENCE: [INAUDIBLE]

AUDIENCE: Yeah.

BEN OLKEN: What? Sorry?

AUDIENCE: Maybe in terms of margin of value of--

BEN OLKEN: Well, that's exactly what I was going to say. Yes, exactly. I was trying to say, well, yes. You can compute the MBPF of all of these things and try to think them through. That's not a-- I think that the-- so that's what I was going to say, yes.

You just figure out a welfare frame. So two points. The first is you need an augmented-- OK. So for those of you who have taken public finance, you've probably talked about the marginal value of public funds as a framework for just thinking about, what is the marginal benefit of dollar of government spending relative to the net cost to the government, where the net cost takes into account the effect of the spending on future government revenue?

So that's the rough-- another way of saying it is, what is the recipient's willingness to pay for whatever you're providing relative to the net cost of the government, where the net cost takes into account effects on future government revenue? So some of this stuff actually looks free, has infinite MBPFs if basically the net cost of the government is zero because it increases people's future earnings so much they pay more taxes and so on and so forth.

In principle, those approaches allow you to compare everything to everything else. The problem is you have to actually estimate-- you have to do two things. Number one, you have to estimate what are the gains, or what is the willingness to pay in some real sense? And what are the real costs?

That's hard to do for some of these programs, particularly some of the infrastructure programs. The other point is you also need to augment that approach to make sure it takes into account redistribution effects. So the baseline MBPF considers a dollar is a dollar for everybody. But if programs are fundamentally about redistribution, you've got to take into account the fact that giving a dollar for a poor person is more valuable than a dollar for a rich person because of concave utility.

And you've got to include that in the welfare calculation somehow. So the long and short of it is yes. You can do it. What you need to do is a careful cost benefit analysis or a careful market value of public funds. There's a bunch of inputs you need to do that.

But you want to think of what is the net social return relative to the net social cost of the government taking-- weighted by some redistribution metric, roughly speaking. OK, other questions? And there's a [INAUDIBLE] and [INAUDIBLE] paper on some of the marginal value of public funds in the *[? KJE ?]* I think last year, which goes through some of this in a fair amount of detail of how do you think about that.

I don't think it talks about the redistribution component a bit as much. But I think it's probably in there somewhere. And also, when our review paper is out, it'll all talk about that, too. Other questions? Yeah?

AUDIENCE: What do MBPFs look like for these programs in developing countries?

BEN OLKEN: I don't think it's been very well calculated for most of the programs. I think it's a great question. I don't think people have calculated for a lot of these programs. The other thing I'll also say is that the-- yeah, let me just leave it at that.

But there's a lot of interest in some-- so there's starting to be some long-run evaluations of some of these programs, which are starting to say what happens to these programs to people 10, 20 years later to see, is it actually-- are they entering the workforce? Can we actually-- because one of the questions is, what is the long-run impact on people's ability to work, and ultimately pay taxes, and all that stuff? And are they in the tax net?

I don't think that stuff has actually been fully thought through, frankly. And I think there's a lot of interesting issues in doing so. Often, for example, when I've seen some analyses like this, they'll say, well, we're going to look at some interim impact. They'll maybe do it in two stages.

They'll say, well, there's an impact of this in going to school. And then we'll take the schooling regression and apply it to that. And we'll back something out of that.

So that's some of the approaches people take. But I think that thinking about some of the stuff seriously would be interesting. And there'd be more work to be done there. Other questions?

OK. So the flip side-- the next part of it is, well, how do we raise money to pay for all of this? And how do we think about the challenges of the public sector in a developing country context, OK? And there are, broadly speaking-- so what am I going to talk about?

I'm going to talk about a bunch of things. First, I'm going to talk about what's different about the public, the tax side of things in a developing country context? And I'll talk about a couple of different theories for what's going on there.

I'll then talk about, I think, two different ways of approaching the problem of tax enforcement. So one of them is related to the paper you guys read for today about the VAT. Essentially, you can think of that as through the tax structure. Are certain kinds of taxes or tax structures is going to work better in these low information environments?

And the second is tax administration. Are there certain things we need to do about the way we collect taxes that are going to be different in these developing country contexts? And ultimately, does this matter?

And then I also want to talk briefly, if I have time, about something called informal taxation, which is, I think, the analog to informal insurance on the tax side. So that's the outline. All right.

So just stepping way back, if you take a public finance class, there's a vast literature on taxation issues, incidence, optimal income tax theory, capital taxation, consumption taxes, dynamic considerations, and so on and so forth. You take an entire semester on this stuff. I think by comparison, understanding how do we think about this stuff in a developing country context is we know relatively little.

We a lot more than we did maybe even 10 years ago when I first started teaching this class, which is exciting. There's actually been a really active area of research, which is part of the reason I want to mention it for people who are interested in development economics questions. And the other thing is I think what we know suggests that there's something that is very much a development economics issue, thinking about something in the tax structure.

And I think there's two. The first is-- and these are the same themes that I think came up in the stuff I talked about redistribution. So one issue is an information problem, right? In the same way we had information challenges of identifying who's poor, we also have information problems about figuring out how to tax and how much to tax people, right?

So how do you levy an income tax on subsistence farmers, for example? Or even do you want to do that? How do you tax labor in a place where you have 50% of the economy is in the informal sector?

The government just has much less information about a lot of people. And so how do you deal with that? And the second is, I think, on enforcement or governance, which is I think also related to the information problems, there's a lot of opportunity for corruption and other problems. And how do we think about the enforcement side?

And I think that suggests that just the way we may want to think about taxation in a developing country context may not just be, we should think about what is the optimal Ramsey tax? And how do we set our taxes-- I think some of those issues are perhaps important.

But some of these other issues may be really first order. And that's what I want to talk about today. OK. Yeah?

AUDIENCE: I just have a question. How much of these issues are limited to the public sector and the [INAUDIBLE] most of them [INAUDIBLE] taxes? The same issues aren't true when these private companies [INAUDIBLE] figure out which customer is good, or something like that. Does that make sense?

BEN OLKEN: Yeah. So I think it's a-- yeah, it's a good question. So let me hold that thought for a minute. I actually think that there are particular issues vis-a-vis tax.

I think they are a bit different. So let me give you some examples. And then you can see what you think.

AUDIENCE: OK, I just have [INAUDIBLE].

BEN OLKEN: So particular one of the big themes, actually, is about third party reporting, which is about formality. And that has a very specific tax structure.

AUDIENCE: OK.

BEN OLKEN: Is that what you're saying?

AUDIENCE: Yeah. I guess what I'm [INAUDIBLE]. Are there any particular concerns about space or moving some things that [INAUDIBLE]?

BEN OLKEN: In general, it's the opposite. In general, governments have a lot better options than private firms, right? If you think about what can the IRS observe about you, they know a lot more than amazon.com.

Maybe they don't, actually, in this day and age. But in principle, they can observe all kinds of things that the private sector can't because they can mandate that people report stuff to them. Right? So that mandatory reporting of information means that they can-- it's interesting.

I just said, in the new tech world, they can figure out ways of backing a lot of the stuff out by what you click on or whatever. But the government has a fundamental advantage, which is it can make you tell them stuff. And they can throw you in jail if you don't or if you lie about it.

So that is a huge advantage the government has over the private sector in getting information. Right. OK. It can make you tell them stuff.

It can make banks tell them stuff. Again, how that's actually a policy question and how much of that additional reporting we want-- and different countries make different choices about how much information they want the tax authority to have. And actually, if you've been reading the news in the US, for example, this question about should the IRS be able to access information about your bank balances and under what circumstances is actually an active question of policy debate, where people seem to actually disagree about this question, apparently.

So OK. So anyway, so I think it is different. OK. So just a few overview facts.

So this paper by Gordon and Li basically made the point that just overall-- the paper is a bit out of date now. But overall, the tax structure looks very different in developing countries than it looks in developed countries. So there's just a few facts here.

Actually, let me pass out the handout. Oh, I want one. Thanks. OK.

So this is too small to read on the screen. Maybe I can zoom in a little bit. OK. So what's the point?

Several just facts. First is rich countries have a lot more tax revenue as a share of GDP than poor countries. OK? So in general, the size of the public sector looks really different.

Number two, where they're getting their tax looks really different. So income taxes in general are a smaller share of revenue in rich countries than in poor countries. Corporate income tax is a much larger share of income taxes in poor countries relative to rich countries.

So in the very poorest countries, half their income tax comes from corporate income taxes. In rich countries, it's a sixth or something. So rich countries are better at taxing individuals relative to corporations. Poor countries tax corporations.

At least in this period, tariffs are a non-trivial share of government revenue. This probably was true in the historical US, for example, also, that tariffs used to be a much bigger deal than they are. So tariffs are non-trivial. They also talk about seigniorage as a function of revenue.

Again, inflation has gone down up until the last six months. Inflation had gone down a lot, I think, around the world. So this may be a little bit out of date. But inflation was more of a big deal.

And seigniorage income, which is the revenue the government gets for printing money, was a non-trivial share of the government budget in some of these countries. And the other fact is the estimates the informal sector as a share of GDP are much larger. So it looks like these countries are doing different stuff.

And a question is why. And so here so there are two theories that can-- here are two theories that can explain why things look really different. So one is this theory by Gordon and Li.

And the other is a theory by Kleven, Kreiner, and Saez. So this one is about The financial sector generating information, which I'll talk about in a second. And the other one is about third party reporting generating information.

And underlying both of them is this question of the government needs the information to tax you. Otherwise, you evade. And where do you get the information? And why might that look different in these two different contexts?

So I'm going to talk about this one very quickly. And this one I'll go through a super simple model to make the point. OK. So the Gordon and Li paper argues that the financial sector is really important.

And the basic idea is that if you use the banks, the banks observe how much money is flowing through your accounts. Because that's what banks do, right? They have to keep track of how much money you have.

So that creates a record of what you're doing. And the government can tap into that record and use that as information to tax you. OK? And they just argue, basically, that essentially, larger organizations need to use the financial sector.

So sorry, you can do it in a couple of ways. Number one, there's more financial development for other reasons in rich countries, just like people are using banks for other reasons. And number two, the corporate sector really-- you can run a local pizza shop, maybe, in cash. But you can't run General Electric in cash because it's just too complicated.

And you need to actually use banks. And there's another fact, which we'll talk about in 14742 that-- firms are just a lot bigger in developed countries for other sets of reasons. So little tiny mom and pop shops that maybe potentially could run in cash are a much larger share of what's happening in the corporate world in India than they are in the US, for example.

And so that also gives you-- they argue that basic fact of the financial sector generating information generates a lot more of the pattern, some of these patterns. For example, they're going to say that's one of the reasons why developing countries are going to tax-- spend a lot of effort taxing corporations. Because at least they can see the corporations and see what they're doing if large corporations are inelastic in their use of the formal sector.

And then they argue everything else has to do that, which is now-- so for example, we have this sector that's now being taxed a lot, which is the informal sector. So they argue maybe you have to have tariffs to protect them and make sure everyone's on a level playing field. And then you use inflation, they argue, maybe as an optional tax on the cash economy.

So basically, if putting your money in the bank means it gets taxed, we also want to have an equivalent tax on your cash holdings since they do that through inflation. This maybe is a little too neat maybe. It's not totally obvious to me that in fact, actually, what's really happening in these countries is the central banker is saying actually, the optimal Taylor rule level of inflation in our economy is 10% because we want to correct this tax wedge between-- so this story is, I think, consistent.

I'm not 100% sure it's how people are actually thinking. But that's the story. So their argument is it all flows from this financial sector. And everything else in the tax structure flows from that. I'm not going to go through that in a lot of detail. But I wanted to mention that as one argument.

The second argument is this argument that Kleven, Kreiner and Saez have put forward. And this is related to a series of papers that they and others have done on this idea of third party reporting. And this is consistent with the VAT paper that you guys read for today.

So the basic idea of a lot of taxation is how does the government know stuff? Well, it gets information from two different parties and compares. And if you get different answers, there's something funky going on.

And so what's called third party reporting of the same information-- so the government is the first party. I'm second party. I report to the government. Having someone else report that same information makes it harder for me to lie.

OK? So here's a simple model that can explain why is this relevant and why is this potentially a development phenomenon? So most taxes, they argue, are collected through third party firms-- third parties, such as employers, which double report.

So for example, I work for MIT, right? And every year at the end of the year, I tell the IRS how much I think I earned. And MIT tells the IRS how much they paid me. And the IRS can just check if I'm reporting something different than what MIT reported.

That's a red flag. And I get audited. Or actually, they send me a correction letter. They would send me a correction letter saying you've underreported your taxes. Please pay, OK?

And there's some work even in modern economies. There's these same sets of co-authors, or some combination of them have a nice paper in Denmark, for example, showing that underreporting of income is very, very low when it's third party reported, and very high, even in Denmark, when it's not third party reported. OK?

And so I think that's probably true in lots of places, right? So this may be for-- yeah. And the idea of this paper is to argue that this third party reporting-- it's easier to make this thing work in really large firms.

And then they're going to rely on the fact I just told you, which is that firms are much larger in the developed world than the developing world. So what's the argument? So imagine there's a firm with n employees. It generates some surplus, big W, and some wages, w1, wN.

And the total wage bill is just the sum of the wages. Third party reporting is that the firm and the employees report w bar to the government. And the government taxes you based on w bar.

And imagine if w is not-- the firm has to keep track of what it's doing. So it has a set of books, right? And people inside the firm know what w is. And they know what w bar is.

And if those are different, people understand that, right? So in particular, if w is less than w bar, any employee can go to the IRS and report that you're cheating. OK? And if someone announces you, the government applies a fine theta for all the evaded taxes plus collects the back taxes.

So basically, if you get denounced, you pay 1 plus theta times the amount times w minus w bar, OK? So if everybody can agree not to denounce you, then we all just report w bar equals 0 to the government. And there's some surplus. And we split the surplus, OK?

The problem is that a single employee can denounce the cheating, OK? And this can happen for two reasons. There can be random shocks, which is if someone makes a mistake or someone gets annoyed or aggravated.

Or you can actually have rational whistleblowing, OK? And both of these are going to make it harder to have misreporting in larger firms. So why is that?

So here's the random shock version. It's very simple. So if w is not equal to w bar, then imagine that each employee denounces with probability epsilon. What is epsilon?

It's I get really mad at you. And I go to the IRS and turn you in. And actually, if you've talked to the IRS, how do they learn stuff, it's really a lot of cases like that. Some employer has some disagreement with the employee.

The employee gets really aggravated and is like, well, I'm going to show you, and marches into the IRS office and denounces the firm, OK? So I think this idea that people get upset with some probability and go to denounce the firm is actually a pretty reasonable summary of what happens. So actually, in fact, we had a student a couple of years ago who was looking at not tax, but but a related thing of Medicare and Medicaid fraud in the US and talked to a bunch of prosecutors.

And they were like, yes. This is exactly how we learn about stuff. Some disgruntled employee walks in and is like, I'm going to tell you what's actually going on. And they're like, please do.

And OK. So OK. So what's going to happen? Well, this is almost obvious when you say it. What's the probability the firm's going to get away with it is 1 minus epsilon to the nth power, where n is the number of employees you have, right?

So everything basically follows from that, right? So what is the probability that you get away with it, essentially, is going to be declining with that. So at some point, it's going to basically be not worth it for you to do that. And in particular, they give this condition that they're going to evade if 1 minus epsilon to the nth is greater than theta over 1 plus theta.

You can work this out yourself. I don't think it's crucial. This is the key point, right?

The more employees you have, the more likely that someone's going to get aggravated and denounce you to the government. And therefore, really small firms can probably sustain this with high probability. And therefore, it's worth it.

For really large firms, the probably you get denounced is really high. And because the penalty is applied to all-- and the key thing driving this model is it only takes one person to denounce. But once I denounce, I catch everybody. And you pay the big penalty.

So the penalty per person is constant. But the probability is increasing with them. OK, so that's that model.

The other model-- oh, I didn't actually put it up. The other model is actually, I think, even a little more interesting, actually, which is this whistleblower idea. And the whistleblower idea is-- and this actually-- the US actually does exactly-- it's actually one of the few countries I know that does exactly this.

So the US's whistleblower policy is if you walk into the IRS and you say MIT has been cheating on its taxes and here's how, they pay you, the whistleblower, between 15% to 30% of what they recover. OK? So why is that a great policy? Because now as the firm gets large, any individual employee can walk into the IRS and take the whole 30%, right?

So if for a firm of size 2, for example, this is not worth it, right? Because you're only getting, say, 30%. But you were getting half the surplus. So 30% is less than half the surplus. So you're worse off.

For a firm of size 10,000, this is a really awesome deal, right? And therefore, equilibrium people-- they don't actually have to pay very much in equilibrium because nobody cheats, right? So in equilibrium, this is free, but prevents cheating in large firms.

So I think this idea of these whistleblower rewards, actually-- this one, I think, is more the developing country version. Because I haven't seen a lot of developing countries do this whistleblower thing. But the same argument, I think, also could be there.

But in any case, both of these have the principle that as the economy grows, the optimal firm size is going to increase. And therefore, you're going to end up in these really large firms. And in really large firms, it's basically, it's just too risky to cheat because someone will probably denounce you, OK? And I think there's probably something to that. Yeah.

AUDIENCE: It's actually the case that it's a firm [INAUDIBLE] individuals who actually have information about the taxes or a tax rep within the firm also [INAUDIBLE] proportionately to the firm?

BEN OLKEN: Yeah, it's a good question. So you're right. So in this model, everybody can denounce. You might think as the firm grows, we have a specialized accounting department who becomes more and more specialized.

And maybe it's only three accountants actually have all the information. The problem is every individual knows they were cheating on your own taxes.

AUDIENCE: Every individual knows?

BEN OLKEN: Yeah. Because even in the model with an accounting department, you know what you actually got paid. And you know what the firm told you to write down on your taxes. So everybody knows that the firm is cheating.

And that could be enough to start the process. You could at least walk into the IRS and say, I don't know how prevalent this is, but here's what they told me to do. And here's what they told my buddy to do.

And maybe that's enough. So you're right that that creation of specialized knowledge may undo this. But I think the key point is everybody has to be complicit in the cheating. OK.

So this paper by Jensen-- I think this is actually now forthcoming in the AER-- takes seriously this idea that basically, individual income taxation only works in formal employment relationships. So another way of saying it is that-- so what's the idea? He says, look.

If I'm an employee in a formal firm, then the firm is double reporting, right? And maybe the government can try to collect it. But if I'm in an informal firm, or if I'm self-employed or whatever, forget it. The government has no hope, OK?

And Jensen asks, well, what are the implications of this for the tax structure more broadly? And in particular, he basically postulates that countries realize that when the employment share-- so what are the implications for how we design the tax code, OK? And he argues that basically, when the employment share in a given income decile is too low, governments basically just give up on taxing the people altogether and just make them exempt from income tax.

They don't even try, OK? And then he basically argues that as countries develop, as countries become larger, the share of people who are working as employees in those larger firms-- that increases and moves down the income percentiles. And therefore, countries are going to change their tax structure endogenously. OK?

So I'm going to focus just on some cross-country facts for now, which I think are cool. Sorry. These are first the US over time. So what is this?

So this is the employee share, OK? The employee share by income decile, OK? So this says what fraction of people are employees in part of a firm or self-employed? The informal sector, we're not really going to focus on here.

So what do you see? So in the 1870s, it looks like basically, most of the people in the United States in most income deciles-- it was only the employee share only happens at the very top of the income distribution. And we have no income tax in 1870. So the US income tax was introduced in 1918, 1919. I don't remember exactly-- approximately in 1918 plus or minus a couple of years.

1912? Does anyone remember? It's in the 1910s. There's a Constitutional Amendment to it. Anyway, approximately then.

Certainly not here and yes here. And so this bar is what is where is the exemption of the income tax? So almost all countries say people below some income don't pay income tax. And he says, well, where is that threshold in the income distribution?

And you can see that basically what's happening in the US over time is two things. As the US grows, the employee share is really changing. OK? So it's always the case that richer people are more likely to be employees.

But where that thing is happening is moving down the distribution. And as that's happening, the income tax exemption threshold also is moving to the left. OK? And that's the US over time.

So the next point is that if you look across countries you see very similar patterns, OK? So India looks a lot more like the historical US, right? Only employee share is really high. Very high income tax threshold.

Indonesia-- a bit richer in his data. It's a bit to the left of the threshold. And again, this pattern is similar.

If I were to show you Mexico, it's what you expect. It's going to look like a little more to the left on both the black line and those crosses and so on and so forth. And then you can plot all these different countries.

And you can see a couple of things. The first is what is the overall employee share? What fraction of people are employees versus self-employed? You can see this general pattern we were seeing that this is moving to the left and getting higher as we become richer-- that seems to be true looking at a cross-section of countries around the world, OK?

So richer countries have high employee shares. Poorer countries have much lower employee shares, OK? You alsc see that the personal income tax base is what share of the population is above that black line required to pay personal income taxes? That's also declining with income.

So richer countries are trying to tax further down the distribution. And this says what fraction-- this is the most amazing graph of this paper. This is what fraction of the total employees are in the personal income tax base?

So it's the ratio of those two things, right? What fraction of the people who are employees are above the black line is almost constant as a fraction of development. OK?

So this is basically arguing that as countries-- what's happening is countries are realizing that they can only tax people who are employees. And they wait until-- and they're setting-- they're designing their tax schedule to take into account the fact that basically, it's only worth it to go after the employees. And they're doing that almost a constant-- not exactly constant. It's still the case that it's a little bit higher. But I think this is remarkably flat by comparison.

AUDIENCE: I think that [INAUDIBLE]--

BEN OLKEN: What?

AUDIENCE: I have a question. What happens with informal employees?

BEN OLKEN: This is only formal employees.

AUDIENCE: OK, [INAUDIBLE].

BEN OLKEN: Should only be formal employees. I believe it's only formal employees. I'm not 100% sure that's how they calculate the data.

But it should only be formal employees. The model is only about formal employees. I don't remember the details of how he constructs this from all different data sets to know if it's formal or informal employee examples, but it certainly should be based on the model. Yeah?

AUDIENCE: What if informality is [INAUDIBLE] imagine maybe everybody shows the same [INAUDIBLE] in the US, but like India, there are some countries that hire the same employee and just choose to have them be formally on the books and social security or just have them pay the cash [INAUDIBLE]. So it's [INAUDIBLE] salary [INAUDIBLE].

BEN OLKEN: Yeah. So let me apologize that I don't remember exactly how he's constructing this data, to answer Sal's question. So I don't remember that question well enough, and I should.

But I think the point is this is an equilibrium relationship. I think that's the way I would think about this graph. I guess the broader point I'm trying to-- so I think this is a, I think, very-- and there's another part of this looking at the US. But I think to me, what I think is really pretty remarkable about this is showing that at least it looks like-- and the data is really very consistent with countries choosing their threshold of who we're even trying to tax, taking into account what's going on in the employment structure.

I agree with you that there's a whole range of additional questions to think about-- formal versus informal, and so on and so forth. So I don't think this is by any means the last word on this. But I just think the point I want to illustrate is that the idea is that thinking about the previous model, the data is consistent with the idea that countries realize you can only tax people who are in these employment relationships in equilibrium and are adjusting their tax code implicitly.

And therefore, as countries develop, one thing that naturally happens is because they work in larger companies and that happens as part of the development process, that almost always has implications for what's going on on the tax side. Questions? OK.

I think one question that this may raise is can you improve on the ability of firms to report on their workers? So can we-- if this whole thing was taken as exogenous epsilon thing, but can we change this, right? Can firms change epsilons, or the probability of getting a report?

Can governments change epsilon, the probability of getting a report? And I just want to mention one paper here, which I think is a really nice paper that tries to look at this. I think this paper doesn't-- it has some data limitations. So I don't think totally is the last word on the subject.

So I just want to mention this as, I think, an exciting, interesting question to think about. So what does this paper do? So this paper says, look. Under Mexico's old pension system, workers' pensions were essentially a flat function of how much your wages were reported. Sorry, they were flat function-- they were not a function of how much your wages reported.

They were just a function of years of work. So you didn't care what the government was reporting on your behalf. In the new system, you basically, essentially get some fraction put into your pension account as a fraction of how much the government is reporting of your income. OK?

So now all of a sudden, you go from a world where you have an incentive to care about what the government is telling about your income from one where you don't. And they try to establish a different diff by saying after the reform, maybe young workers have a greater incentive to make sure firms report wages honestly than before the reform. And they find better match in wages before and after.

The paper, as I said-- there's empirical challenges. They don't quite-- so it's a little bit tough. But I think this question of-- I just wanted to mention this as one example of how do you think about improving-- how do you think about changing epsilon? Actually, this was something that I actually myself tried to do.

I tried to develop an RCT on this topic in Indonesia. It didn't end up working out. But we were thinking of well, maybe the government could just start sending people text messages and say hey, here's what the government-- here's what was reported for your wages this month FYI as a way of increasing that epsilon.

But I think more generally, I think this is an interesting question. The epsilon isn't fixed. And maybe you can do something about that. So I just wanted to put that up there. OK, questions?

OK. So now I wanted to talk about the paper that you guys read, actually, sorry, for last time, which is about value added tax. And the reason I want to mention this is I feel like, as I mentioned, there are two different approaches to what we can do about this problem.

Or sorry, the first set of papers was just like, why does this look different in the developing world? The second set of papers are like, what could we possibly do about this? Number one is on the tax design side.

And the other is on the tax enforcement side. So this paper on the VAT is an example of the tax structure itself, designing the tax structure itself to make sure that it's third party reported, OK? So I think for all the non-Americans in the room know what a value-added tax is, or knew prior to reading this because every country in the world basically has one except for the United States, approximately.

So basically, the idea is that firms are taxed on output but receive a tax credit for what they already pay in their inputs. OK? And the other thing you have to note is that imports are taxed and exports are not taxed. So at the very end of this at the VAT chain, you have to rebate.

And the reason is you want everything to be on an equal playing field domestically. So if you were taxing all domestic production at, say, 10%, we want to tax imports at 10% so the imports aren't cheaper than the domestic stuff, so they're equally taxed. And the way that the world has worked this out is that when things come into your border, you tax them at your local VAT rate.

So you have a uniform tariff at your VAT rate. And when you export things, the exporting firm gets a rebate equal to the sum of the VAT on the input chain. So essentially, everything is taxed at its own domestic VAT rate regardless of where it is produced.

That's the principle. I will note that in a lot of developing countries that I've been engaged with, this step-- it ends up being challenging. So governments love taking money in.

They do not like sending money out as a general rule. And so they put a lot-- and also, you can be worried about a lot of fraud. So they put a lot of checks and cross checks on those VAT rebates at the export stage. And that can be an important impediment to trade, actually.

And so actually, thinking about that question of how do you make that VAT exporting step work properly is at least, in some cases that I've encountered, a challenge. OK. And so what why do a value added tax? If everyone participates, it's equivalent to a consumption tax, right?

So this is the idea. So the idea here of our VAT chain-- so I can't remember the way Dina draws it in the paper. Which way did she draw it? OK, she goes to the supplier on top, right?

So we have the supplier. The chain goes down, right? Middle and retail, OK?

So we have-- so this is goods are moving down the chart, right? So the idea is that this person-- this is the raw supplier pays VAT on this piece here. This person pays VAT on this minus this. So they pay VAT on this and this.

And the retail pays the bottom. And this is equivalent to just charging a sales tax at the very bottom, right? If it's uniform, basically, you're paying a tax on the whole production chain.

It's just that here it's collected in these three different pieces whereas here it's collected all at the end. So in the US, for example, you pay a sales tax when you buy things in Massachusetts, for example, of whatever-- 6 and 1/4%. They're taxing it down here. In a VAT, they're trying to do the same thing-- tax consumption-- but they're taxing it all the way through the chain.

And the idea of this is that basically, you get better reporting because in a sales tax, the retailer has all kinds of incentives not to underreport the total sale price. And the government loses those taxes. That's still true here, by the way.

There's nothing going on over here. But in all the previous steps, you have incentives that go in opposite directions. So this person would like to underreport how much they are selling their good for. But this person over here would like to make that number as big as possible because they get a credit for that amount, right?

And so the whole idea from governments is we get accurate reporting of information if we have information reported by both parties. The whole idea of the VAT is that for everything except for the very last step, we can generate double reporting of information. And the government can actually use that to figure things out. OK.

So that's the basic idea. And this paper is about testing whether that idea actually works. And there are these VAT chains. I think in your comments, I just want to mention a couple of things.

You guys mentioned that one thing that actually has to happen, which you'll see as we talk with this paper, is doing this third party reporting is actually tricky. So it requires the government collecting tax ID numbers for everyone to match every transection, or at least every firm at the firm-to-firm level. It turns out that actually is harder than you may think.

And at least in some contexts that I've worked at, they collect all that data. But actually, they try to actually analyze it and figure out if people are actually cheating. It turns out to be a very labor-intensive and surprisingly hard to automate process because the world is messy, right?

It's a lot simpler if I have one employer. And the IRS can pretty easily check my income tax. For a complicated world where a firm has got thousands of clients and hundreds of suppliers, doing all this stuff can be surprisingly tricky.

So someone mentioned automating. What is the impact of digital payments or automating stuff? I think actually understanding all that and how that implies what governments can tax as firms are developing and becoming-- countries are developing and becoming more electronic or whatever is an interesting area to think about.

I had another thing. What was I going to say? Anyway, OK. All right.

So what she talks about in the paper is that there's apparently two different kinds of tax evasion. She calls them collusive evasion and unilateral evasion. So collusive evasion is we omit an entire transection completely, whereas unilateral invasion is we just disagree on these amounts. OK?

And she also is going to try to understand what the impacts of this are. I think after I went through all this. So she argues that basically, what happens if I shock this firm over here?

Well, if I was just-- and I make sure this firm becomes honest-- well, in collusive evasion, all of a sudden, that whole transaction appears on the books. Yeah, so OK. So what happens, basically-- inclusive evasion-- when I shock this firm, that new thing goes up.

And this firm can actually start reporting it as an input and maybe reduce their VAT. Over here, what happens is they all have to move towards agreement. And so those could be two different effects.

What I want to talk about a little more, though, is what she actually does in terms of the experiments. So she does two different things. The first thing she does is she does a spillover experiment, which is to say she tells these firms that they're going to be suspected of tax evasion.

They're told they're going to be audited. They then go and audit them. And they see what happened, what actually happened.

And the thing I want to note here is this audit is the thing required to actually go and match everything up. So she only knows what happens to the upstream and the downstream people in the audit group because doing that whole matching actually is manual and requires the audit to actually take place, OK? In the deterrence experiment, she basically sends some random fraction of people in different parts of the chain a letter saying, you're likely to be audited, and looks at their responsiveness. OK?

So these experiments are doing two different things. And so this one, I think, is pretty clear, right? This is basically testing for the VAT chain. This one, I think, has the cleanest test.

This one basically says if this VAT chain thing is actually happening and I tell you all of a sudden you're going to be audited, the question is not just what your response is. But does that feed through to your upstream and downstream people through this VAT chain? The second piece, the deterrence piece, is a little less clear, I think.

In the deterrence piece, she basically says, look. In the VAT world in general, where should evasion be highest, right? We think evasion should be highest down here at the retail level.

So if you believe that what's happening in the deterrence experiment is that it moves you from whatever you were doing now to no cheating, then you expect a larger responsiveness down here. Because there was more cheating to begin with. But I think, as one of you pointed out in your comments, this one is a little more complicated because you could also have the view of maybe the auditor can't figure it out down here because there's way of figuring out what the truth is over here.

So I think this one, the deterrence one is only interpretable if you think of the deterrence as moving from whatever you were doing before to the truth. Then I think her interpretation of using that to measure differential amounts of cheating is right. If you say, well, it's not clear what we're learning from the audit unless-- because that is related to where you are in the VAT chain, it's a little bit less clear.

She's trying to infer something about the level from-- when she measures the derivative, she measures what is D taxes, D audit for the three different places. She's trying to infer something about the level of preperiod evasion from that derivative. And I guess the point I'm making is that only is valid if you basically assume that the shock is getting you to the same point in all those places.

And I think that's not totally obvious. So that's why I like the spillover one better. I think the spillover one's cleaner. Questions about the experimental design?

OK. So what does she find? Oops. She finds-- this is the spillover stuff.

So this says do you pay more VAT if you're going to be audited, right? Then yes, you do. OK. That just says that's the own effect.

And the question is are there upstream or downstream effects? And the answer is there are pretty strong upstream effects, OK? So what this basically says is I audit these people over here, right? And these people also pay more taxes.

OK? And that's it. And I think that's the key point. Why-- the retail people not pay more taxes? Well, because they have a lot more flexibility at the bottom end, right?

So people who are downstream, they can adjust in the three chain version. In the real world, it's more complicated. But certainly, the more downstream you are-- the retail people are the ones that have the most ability to evade things.

So that's the key point over here is that it seems to spill over upstream. And I think that's the best evidence that we have that actually, these chains seem to be mattering. In the deterrence paper-- this is too small, sorry.

So let me zoom in on it since you don't have the handout. The key point here is just a heterogeneous treatment effect. So where does the deterrence effect have the largest effects? It has the largest effects for the people who are more downstream.

So that basically is arguing that there was more cheating happening-- under the assumption that you can learn about the level from the derivative that I talked about before, she's arguing there's more cheating down there. OK, that's what I wanted to say about this paper. But do you have any comments or questions from having read it? Kadesh, you look like you have formulated a question.

AUDIENCE: I just wanted to [INAUDIBLE] that the deterrence experiment-- do you think that was designed with this question in mind? Or do you think there was-- they learned something. And then was like, oh, maybe we should test the [INAUDIBLE]?

BEN OLKEN: Oh, did they go in-- did they go into this--

AUDIENCE: With that test, with this theory in mind.

BEN OLKEN: I can't speak to what she was thinking when she designed the experiment. But I suspect they did. I think this is the key question in VAT those people were trying to understand, is understanding the spillover.

So I think they were trying to go into it with the two things. I think the reason for the two experiments is the deterrence experiment-- she can measure the outcomes. She doesn't need to know who the suppliers and the clients are.

You see what I'm saying? So the advantage of the deterrence experiment is you get this enormous sample. And you don't need to know the whole VAT chain.

So I think there, she can try to say, well, look. I'm trying to understand something in a really big sample. What can I possibly learn?

And she doesn't know who the suppliers and the clients are. The reason for the audit experiment and why she had to do that on a small sample is she only can learn who the suppliers and the clients are once they do that manual audit process. At least in this context, it's not sufficiently automated.

So if you had a fully automated system, you could run the spillovers in this experiment, too. But she can't do that because they don't have that data.

AUDIENCE: So the deterrence thing is just adding this kind of [INAUDIBLE] validity?

BEN OLKEN: Yeah. I think it's adding a different piece. Under the assumption that I was saying before that you can learn about the preperiod level from the derivative, it is testing a different aspect of this question also, which is to say, look. There's two questions you might believe.

One is you'd say-- two questions about the VAT from the simple theory. One says, look. This thing spills over. The other says there should be way more cheating at the very bottom, at the very downstream.

And so she's arguing that by learning about this derivative, you're learning that maybe there is much more cheating there. It's certainly consistent with that idea. Other questions? Yeah, Ahmed.

AUDIENCE: Signing returns to the letter is quite cheap, right? It does feel like there's a simple way of increasing her complexity on this--

BEN OLKEN: Well, the question is what is-- you can't deter people-- you have to have rational-- you ultimately have to deal with rational expectations. So you can't tell everyone they're going to be audited with 100% probability, and not audited with 100% probability. Eventually, they'll it out.

AUDIENCE: But for the short term, it's free money. [INAUDIBLE]

BEN OLKEN: But in the long run, it has a lot of-- I think this was worded in a way to be honest and truthful. But I think that you have to-- and that suggests that's true. But it's not a free money machine.

I think you probably have to actually increase the audit rate, too, ultimately. But there is this basic fact that many governments are not spending enough resources on enforcement, which I will argue in the next paper, which I'll probably teach you at beginning next class, at least in Indonesia, and is actually a very active debate in the US right now. Should we spend more money on tax enforcement?

And many estimates conclude that governments are not spending nearly enough money on tax enforcement. It's a very, very high return activity. Yeah?

AUDIENCE: In a case where a supplier and a buyer are bargaining over the tax, why do we expect it-- why do you expect them to-- for a solution to marketing game to be the true amount, the correct amount versus just some solution to some marketing--

BEN OLKEN: Well, the tax code doesn't care.

AUDIENCE: Oh. Why is that?

BEN OLKEN: Well, imagine I move this line over here. As long as they agree, the government's still getting taxed on this whole amount.

AUDIENCE: So they're probably going to say [INAUDIBLE].

BEN OLKEN: For as long as they agree on what they're going to report-- that's the beauty of the VAT. If we're poor-- if we say-- for everyone except the bottom layer, my sales to you are how much-- what I sell to you is how much you buy from me. So if I say I'm selling to you at a lower price, I pay less tax.

But you pay more tax because your inputs go down. So at some level, the government doesn't care. You see what I'm saying?

AUDIENCE: It was meant for firm size stuff, right?

BEN OLKEN: Yes, it might have implications for the economy. I think also, I think maybe for the deterrence reasons, maybe for the denouncement reasons we talked about before, there's a focal point on the truth. It is an interesting question, actually.

There may be some-- I suspect there's some good microtheory bargaining answer that I don't know the answer to for why that might be true. But I don't know it off the top of my head. Yeah.

AUDIENCE: Is there any real literature, if I'm really into this kind of systems, to organized crime--

BEN OLKEN: So what, sorry?

AUDIENCE: To organized crime and money laundering?

BEN OLKEN: Tell me more.

AUDIENCE: So the general idea is that maybe you're organized crime. You use small retail stores to just laundering money through there. And then if you set up a system where you have double reporting like this one--

BEN OLKEN: You make it harder.

AUDIENCE: It would make it harder to money-- uh, yeah.

BEN OLKEN: Yeah. So sorry. So why is money laundering relevant here? So money laundering is almost the inverse problem, right? The money laundering problem is the inverse problem. I have a bunch of cash it's illegal that I want to make look legal.

And so the way that they do that is you basically, you have a business that has a lot of where the actual income level is not reportable. And you overstate your income, essentially. So you buy a cash business. The money launderer runs a cash business and overstates the amount of actual business going through as a way of mixing the cash they want to make legal with the legal cash.

And because the government can't observe the retail side, you can overdo it. That's the idea of money laundering. It's the inverse problem of this.

People want to pay more taxes, right? So yes, there could be implications to that. I haven't thought them through.

I think as long as-- that's why the money laundering guys work at the retail level. It's that same point that the government really doesn't know what's happening at that very last piece. And that actually leads me to the very last thing I want to say on this topic.

Actually, no. Sorry there's two things I want to say on this topic. Number one is one cool thing about the VAT is that in principle, governments are collecting huge amounts of data on who's trading with who, at least in the formal network.

And I think a bunch of researchers have realized this and that this is a cool data set, very hard to obtain because it's tax data-- but a cool data set in principle on the whole network of supplier-supplier supply chain relationships. And I think Dave Donaldson and Dina Pomeranz and several co-authors actually have paper using this data, as do, I think, a few other people.

So it's an exciting area people are starting to think through. And the second thing I wanted to mention is-- the last thing I would say is, well, what can you do at the very end of the chain? And so Joanna Naritomi has a paper on this, which basically says can we find a way to incentivize people at the very bottom of the chain to be honest?

And I'll just note, if you've ever been to a store where it says if we don't give you a receipt, your purchase is free, what are they doing there? Why are they doing that? It's the same idea.

They're trying to incentivize the consumer to make sure that the transection is reported through the cash machine, right? In that case, usually, it's about the employer worrying about the clerk stealing the resources, which is why they want it to be reported. But the same idea is, how can we incentivize the very retail consumer to make sure things are properly reported?

And so she studies a policy in Brazil, where the VAT receipt issued at the final thing was also a lottery ticket. OK? So for every dollar of actual legally reported VAT, you get to enter some lottery. And the government gives you prizes, OK?

And it's very small. So the value is 2% of the total purchase price. So I don't know what the VAT rate in Brazil is, but let's say it was 10%.

And even if the retail sector is, I don't know, 30%, it's still a small share of the total VAT the government is collecting even in this last piece. But it's at least some reason that if you were at least totally indifferent-- you shouldn't be totally indifferent. And she finds that actually, this leads to a huge impact, a 24% impact in reported revenues from the very bottom firms.

And so that suggests that maybe there's something you can do to incentivize people down here as well. OK, I'm out of time. I'll stop here. I'm going to finish up by talking about tax administration next time.