

[SQUEAKING]

[RUSTLING]

[CLICKING]

BEN OLKEN: Hi, everyone. Let's get started. So I think it's now me for the rest of the semester. So for whatever reason, yesterday I swapped one lecture last week. Now I won't do that this week. Now I'm back.

So what I want to talk about today is land markets, including the paper you guys read for today. And broadly speaking, we're going to cover, I think, two broad themes in that market. So one is why the allocation of land may matter. So that's moral-hazard and sharecropping issues in particular, related to the paper you write today. But I'll start with some basic theory on this stuff.

And then the other question is, how do we think about property rights and why do secure rights over property matter, how do we think about land titling and things like that, and the implications of those for investment decisions. So mostly, today is going to be moral hazard and sharecropping. But I guess today is Wednesday. So Monday we'll be primarily on investment decisions.

So stepping way back, sharecropping is a very, very, very, very old subject. It's been around for forever, essentially, or for a very, very long time. And that's when laborers pay owners a share of the output. So sharecropping is a ubiquitous subject.

And it has been thought about for a very long time, going back even to economists like Smith and Marshall were talking about this issue, and not only just talking about it, but recognizing that there may be an inefficiency from sharecropping. And the inefficiency is you don't get the full returns to your labor because some share of it is going to the landlord. So that could lead you to not work as hard as you would otherwise. That's the basic issue.

So here's just a super-duper-duper simple, simple model to make this very clear. So imagine that, basically, there's some output. The output is f of i , that's where that's some input. Tenant pays some cost, c , could be their own labor, whatever. Owner receives a share, α , of the output.

So what's going to happen, the tenant is going to solve some problem. They get taxed, essentially, $1 - \alpha$ on the output. They pay the full costs. The first-order condition of that is just that the marginal product is equal to the marginal cost, the module of a tax rate.

So if they're being taxed, that's going to be-- the marginal product is going to be greater than c . And therefore, if f is concave, they're going to put less input than they would otherwise. That's super-duper-duper basic tax-- if you tax things, people do less of them.

So one thing that's worth noting, actually, that's very important is-- this is the key thing that's driving this distortion is that this is a tax on output, not a tax on profits. So if I wrote this exact same-- if I changed this so that instead of having a tax on the output, here, I taxed your profits, which is your f of i minus c , then in that case, it wouldn't matter.

You can just check that the solution, here, is the same as if there was no-- I didn't have to share the output here. And that's this-- general fact is that profit taxes are nondistortionary. And that's actually related to broader issues of why corporate income taxes, for example, taxes on profits, not taxes on revenues. It has to do with these kinds of ideas.

Of course, the problem is that what we can observe is your output. Although as you saw on the paper, even observing the output is hard. But it's at least more plausible that you can observe the output, rather than observe your profits, especially when we think of an agricultural-- like a self-employed farmer or something, where a lot of their inputs are their own labor or things like that, that's very difficult to actually observe and measure. So the distortion comes from the fact that we're taxing revenue, not taxing profits. OK? Clear?

All right, so the solution-- and actually, I think one of you-- at least one of you talked about this in your reading response, is the solution to this problem is a rental contract. So faced with this, what do they do? Rather than have this distortionary contract, what should happen is, the tenant should rent the land from the landlord for a fixed rent, r , and keep all the output.

And so that contract would look like this. It would look-- the tenant would maximize i minus f -- sorry, maximize over the inputs, f of I minus cl minus r . The fact that this r -- r is a lump sum here. So it doesn't affect the decisions. So that gets us the first best.

So in some sense, that's the benchmark solution to all these problems, is we should just-- the tenant should rent the land for a fixed fee. So why does that not happen? There's a couple of reasons why this doesn't happen. And that's what I'm going to talk about for the next half an hour or so.

So one comes from Stiglitz. And the idea is that there's a trade off between incentives and risk sharing. So the problem with this contract is that if we imagine that the output is risky, the tenant is bearing all the risk. And the landlord is getting a fixed fee regardless.

You might imagine that the tenant is typically a poor farmer and the landlord might be a rich landlord. In that case, you might imagine that the rich landlord might be less risk averse than the poor tenant. But even, in fact, if they're equally risk averse, the optimal thing is not for one of them to bear all the risk and the other to be totally insured.

The optimal thing is to share some of the-- share of the risk between them. And once you share the risk between them, then you end up in some kinds of contracts where you end up-- that look more like sharecropping contracts.

And so that's the basic idea of this Stiglitz model, which is actually one of the original moral-hazard models, actually developed with sharecropping as an example, is to say, well, we'd like to-- we end up in these contracts that are productively inefficient because we want to do some risk sharing. So let me go through a little bit of that in a little-- go through that a little bit [INAUDIBLE]. Any questions? No? Clear?

OK, so that's basically what I-- so just to be clear, the idea of this, that farming is going to be risky, so we're going to put some uncertainty on this, we know risk-averse agents want to be insured against this risk. But the problem is, if this effort choice is uncontractable, that leads to inefficiencies.

So here is a very simple-- like the simplest possible kind of moral-- simplest version of this model. So let's consider a two-state version, high and low output. There's effort, which is denoted by e . The farmer chooses some effort, e . But the landlord doesn't see e . They only see output.

So we can't write-- and the key point is we can't, therefore, write contracts over the level of effort, e . Effort's going to be costly to the tenant with some costs $\frac{1}{2} c^2$. And the output is going to be stochastic, with the effort determining the probability of the high state.

So with probability e , the output is high. And with probability $1 - e$, the output is going to be 0. And the way we think about these models is the farmer and the landlord are going to write a contract which is going to specify a payment to the farmer, from the landlord-- I think that's the way I did it.

Doesn't matter which way you write the payment-- which is going to be a payment h , little h , if the output is big H , and a payment, little l , if the output is zero. So we're going to be interested in what those little h 's and little l 's are. OK, so it's set up clear?

So just to be clear, the key thing in this model is that the effort choice is going to drive the probability of the high state or the low state. OK, so what's the first best? Yeah, Davey?

AUDIENCE: The solution-- like here, the output is [INAUDIBLE].

BEN OLKEN: Output is observable. What's that?

AUDIENCE: And contractable.

BEN OLKEN: Yes, exactly. The output is observable. We're going to write contracts on output. But we can't write contracts on effort.

So basically-- so the way this model is going to work is, the landlord is going to set up the contract-- here's the little h , here's the little l -- which are functions of what the realized output is. And then the tenant gets to observe h and l and make their decisions about how much effort to put in. And the landlord gets to set-- and the landlord, in turn, chooses the little h and little l , knowing how the tenant will respond and make their contract decisions.

So this is a standard moral-hazard set up. But not everyone has seen these before. So I want to go through this. OK, clear?

All right, so what's the first best? Well, the first best would be that-- so remember in the low state, we get 0, just by normalization. So that's going to make our-- just make the math easier.

So what's the first best? Well, the social optimum is to choose the level of effort such that the return-- the expected return, which is the probability of getting the high return times the value of the high return, you balance that against the cost of effort. So we just take the derivative here, and e is equal to capital H over c . That's the optimal first-best level of effort.

So that's the benchmark. The question is, are we going to get that level of effort or not? So what's going to happen? So the question is, can the land-- by the way, if I'm going too fast, you should let me-- slow me down. That's the problem with teaching on slides, teaching models on slides, is you tend to go faster than people can write. So if I'm going too fast, just slow me down.

So can the landlord implement this first-best contract? What's going to actually happen? So how do we solve this? So the problem for the landlord is written as the following.

So the landlord has two choices, what is little h and what is little l . And what is the landlord care about? Well, with probability e , which will be determined endogenously, the landlord-- the output is big H . And the landlord has to make a payment, little h . So that's the landlord's return in the high state, was probably 1 minus e . We're in the low state. And then the landlord has to pay l and gets no output.

So that's the landlord's objective function. By the way, this is all going to be risk neutral, as you'll see in a second-- but we can introduce risk aversion later. But-- so OK, so that's the landlord's problem-- and has two constraints, the incentive-compatibility constraint in the individual-rationality constraint.

The incentive-compatibility constraint says the landlord is going to solve this problem, knowing that the tenant will in turn take the contract as fixed, and choose her level of effort accordingly. So the tenant, here, is going to solve for a level-- so the level of effort, e , is going to be chosen by the tenant, to solve the tenant's own objective function, which is, well, with probability e , I get payment, little h . With probability 1 minus e , I get payment, little l . And I have to pay my cost of effort.

And there's a third constraint, which is the individual-rationality constraint. And the way to think about the individual-rationality constraint is that the tenant has some outside option, w bar. And the tenant's net utility has to be at least equal to w bar, to be low bar.

And that's to say, the tenant could not take this contract, and could walk away and just get a job as a day laborer at some wage rate, w . So the net utility from the tenant-- so that's this term, here, evaluated at the chosen e , has to be at least as high as her outside option.

So this is the problem. Clear? OK. And those of you have taken contract theory already, this should be very familiar.

So how do we solve this problem? Well, we begin by solving for the-- taking the farmer's solution, taking the contract-- the ic is given. So we basically say, let's first solve the farmer's problem. Given the little h and little l , what is the farmer going to do? And then we'll plug that back in to the landlord's problem.

So what's the farmer's decision? Well, the farmer is going to solve-- is going to maximize the same with this objective function over here. Remember she gets payment, little h , with probability e , payment little l , with probability 1 minus e . She pays the cost of effort.

So her effort decision, as you can see-- so we can just take the derivative of this e , the difference is, she's going to get h minus-- what additional e does, it increases the probability she gets h minus l , the difference in her-- the difference of the high state is different than her payment in the state minus the difference in the payment in the low state, and just is equal to the marginal cost of effort.

So her choice is going to be-- e^* is going to be equal little h minus little l over c . So what you can immediately see here is, this is a model where, if there's no difference between the payment in the high state and the low state, if she's fully insured, she puts in no effort.

So there's going to have to be some-- this is a model where, if there's full insurance, there's going to be no work. So they have to provide some incentives in order to get her to put in some effort, otherwise why would you do anything, because there's some cost of effort. Clear?

All right, so now, the landlord-- so we can also see-- we can all see immediately, by the way, that to get the first best-- remember, the first-best effort-- doo-doo-doo-doo-- first best effort is we'd like the landlord to-- I'm sorry-- we'd like the tenant to put an effort equal to $\frac{H}{c}$.

So what does that mean? Well, the contract-- any contract that's going to get the tenant to put in the full amount of effort, that's going to imply that $h - l$ is equal to H . Which is another way of saying it, is that the gap between what she gets in the high state and what she gets in the low state has to be equal to the social return of being in the high state versus the lowest state.

That's the only way to get her to put in full effort. So you see that by saying, look, if we want this to equal to $\frac{H}{c}$, is we have to make the difference between h and l equal to H .

So that is rental. What is that? That's a rental contract. We've just derived a rental contract. Why is that a rental contract? Because the tenant will make some fixed payment to the landlord-- another way of viewing it, any contract where the difference between h and l is equal to H is a contract where she makes some fixed payment to the landlord.

Or maybe goes the other way, depending on what her incentive for that-- what her IR constraint is. There's some fixed payment that goes one way or the other and then she keeps all the output that is produced. So that's a rental contract.

So we've just derived the result, that basically-- another way of saying it is, in this model, the only way to get the first-best effort is to have a rental contract, which is basically what I just said on the slide. So why is that? We need the worker to face the socially-optimal return to effort. So we need e stars-- we need this, her choice of effort, to be equal to the socially-optimal level. The only way to do that is to set $h - l$ equals to H , which basically means she keeps all the return on the market.

Now in this contract-- now we've pinned down $h - l$. But we haven't pinned down-- but we can slide the level of those things up and down. That's the rent.

And so how do we choose the rent? We choose the rent so that the farmer is going to obtain a lower bar and expectation. That's the IR constraint. So the farmer's utility-- remember the farmer-- what was the farmer's utility? The farmer's utility was-- so now, in the first-best contract, with probability e , the high state happens.

And she gets h . The rent gets paid either way. No matter what, you pay the rent. And you have to put in the cost of effort.

And we know, by the way, that the effort, here, is going to be equal to $\frac{H}{c}$. So we can sub in $\frac{H}{c}$. So substituting in $\frac{H}{c}$, here, her utility is going to be $\frac{H^2}{c}$, minus $\frac{1}{2} \frac{H^2}{c}$. So we get half an $\frac{H^2}{c}$ -- sorry-- yeah, so we get a $\frac{1}{2} \frac{H^2}{c}$, minus the rent.

So that's her utility under this contract, her expected utility under this contract. And that has to be equal to the outside option. So that we're going to set that equal to w , lower bar. And so that's the rent. That determines the rent. So the rent is going to be the level where she's exactly indifferent ex-ante.

And that's a general feature of the way that you solve a lot of these problems, is you figure out-- you first kind of solve for what the IC constraint and what is the difference between the-- you solve the problem to say, what's the difference between the highest and the lowest payoff. And then you slide the whole thing up or down so that the agent is exactly indifferent to their outside option in their IR constraint.

So that's the general way you solve these moral-hazard problems. So that's the rent. That pins down the rent. OK. Clear?

OK, so that's the final contract. Final contract, now-- so now we know both the little h and little l . So little h -- sorry, so our r , here, was h^2 over $2c$ minus w bar. So we can substitute in. So we get-- so the payment is 1 minus the r because it goes the other way. So this is minus r .

That happens in either state of the world. And the difference is she gets to keep the output. So that's just another way of writing-- there's a rent equal to H^2 over $2c$, minus w bar. That's the rent. And then she keeps the return, the h . That's just the way of expressing it in this contract space.

Is that all clear so far? All right. So this is a nice contract because the farmer is going to-- the farmer gets no rents. It's nice from the perspective of the landlord. The landlord likes his contract. The farmer is getting no rents. She's getting exactly equal to her outside option in expectation. So the landlord is capturing all the surplus. And she puts it in full effort. So that's great. Landlord loves this contract.

Problem is, there's two challenges with this contract. The first challenge with this contract is the farmer is bearing all of the risk. So that's the first problem. Second problem is, the farmer, with positive probability, is going to earn this lower-- this l . And that may be negative. So she may get less than 0.

So what happens if she has no money and can't pay? So then we have a second issue, which is what we can think of what are called limited-liability contracts, where you're capped at 0 from below. You can't ever go below zero. If you have no money, you can't pay the rent.

And so this one is like a risk-- an insurance issue. And this is a limited-liability issue. And both of these can give rise to what looks like sharecropping arrangements. But the implications of them are going to be a little bit different. So that's what I want to go through next, is to say, well, how do we think about each of those? And how could each of those different reasons lead us to a sharecropping arrangement?

Questions? All clear? Yeah?

AUDIENCE: In practice, limited liability just means that I can't go-- I can only take-- I can't force her to take out [INAUDIBLE]?

BEN OLKEN: Yeah, so correct. So when so when might limited liability be a thing or might not be a thing? So yes, right. So you can't-- so you're right. It means we can't force you to go take out a loan to pay the rent, for example.

Although even in debt contracts, why might that be? Well, you could just leave, for example. It can be hard to-- yeah, so that's basically the-- and we'll talk about those contracts, limited liability, also, when we get to the credit lectures at the end of the semester, where a lot of these same issues are going to come up.

Yeah. But you could leave. You could be hard to enforce. There's a variety of reasons. Or you might-- right. So you might have asked her that. Yeah. Other questions?

So let's explore both of these. So what happens when the farmer is risk averse? So now let's assume that the landlord is risk neutral, but the farmer has some utility function, which is concave over consumption in both states.

And actually, you can do this with, as I said, with both parties being race neutral. Actually, it doesn't-- just the math gets more complicated. Actually, all these moral-hazard models, actually, the math can get a little unpleasant. Ed, you're nodding. I can see you.

They're surprisingly like not pretty models for-- I think there's nothing-- I don't have a good reason why that is. But that's true.

So OK. And if anyone wants any proof of that, go back and read Stiglitz 74. And you'll see it's not pretty. OK. I think it's just like carrying around-- solving all these contracts, with all these concave utility functions, just gets a little difficult.

Anyway, so now what happens? But the key result, actually, is very easy to see. So this is exactly the same as before. Now the farmer has-- this is concave. We're maximizing e times u of h and 1 minus e times u of l , as opposed to just h and l directly.

And so one thing to point out is, if the landlord implemented the contract from before, the farmer's utility is going to be strictly less than w lower bar. OK? Because of the four-- I mean, this is actually not totally precise. Oh, sorry, this is precise, the u of w lower bar. Exactly right. That's precise.

And why is that? That's just from concavity. That's just because the-- so u of w , lower bar is the utility of the expected return. And what it means to be risk averse, or what concave utility functions give you, is that the utility from-- the expected utility is less than the utility of the certainty equivalent. That's just what we get from concavity.

So that means that-- so what does that mean practically? It means, practically, if we took the contract from before, the worker is now below her IR constraint. So we've got to do something, otherwise you'll walk away.

And what can we do? Well, we could lower the rent. We could raise the fixed fee. Or we could provide some insurance and reduce the gap between h minus l .

And the question is, which should the landlord do? Should they reduce -- the should they do it on the fixed-fee margin or should they do some of it on the other margin?

And one simple way to see that they should do at least some of it, at least a little bit of it, on reducing the h minus l gap, is as follows. So it's basically just this argument here, which is starting from the first, best level of effort, the efficiency loss from reducing h minus l just a little bit, is second order. But the utilities-- because we maximized that, so the loss there is second order. But the utility gain for the agent is going to be first order.

So that's a simple argument, which says, the solution to this problem can't have h minus l be equal to the full-- equal to cap H . We're going to do at least a little bit of this reduction on the dimension of providing a little bit of insurance because that's second order, in terms of the total kind of productivity loss, but first order gain in terms of the insurance value. Yeah, Erin?

AUDIENCE: Is that true when the landlord is also [INAUDIBLE]?

BEN OLKEN: Yes, because another way of thinking about it is the landlord is bearing-- let me think for one sec. Yes, it's true. But let me just double-check. Yeah, I'm pretty sure it's true because it's going to be--

AUDIENCE: [INAUDIBLE]

BEN OLKEN: Because I think it's going to be a similar kind of argument that the landlord-- I was trying to think of whether you need any other conditions. I don't think so because I think that the landlord is going to be totally, fully-- it's a similar argument. Going from fully insured to an epsilon bit of risk is going to be less costly than reducing the risk a little bit. I'm pretty sure that's right.

So on the other hand-- so that says, look, the optimal contract is not going to be the first best. It's going to have at least a little bit of a sharing feature. On the other hand, we also know that we can't fully insure the agent because then the agent won't work.

So this kind of argument says that the optimal contract is going to look something like a share contract, where there's-- maybe with a fixed fee, but it's going to have some feature where the share of the output that the agent keeps is less than the full output.

So some amount of the output is shared with the landlord, even though that has a productivity loss associated with doing that. But it's not going to go all the way to the other dimension of full insurance either.

So that's, in some sense-- this is the key insight. And this, as I said, is the argument for sharecropping given by Stiglitz, which is that landlords and peasants actually prefer-- so Stiglitz's argument is actually like-- it's almost revolutionary, in a sense.

It's to say, look-- not revolutionary in the sense of like, let's have the revolution, but revolutionary, almost in the opposite sense, actually, to say, look, you might have thought that this sharecropping thing was this total exploitative thing. But if only the tenants could get out from under it, they would be much happier.

But actually, he says, no, no, no. Actually, the sharecropping is preferred by the tenants, in some sense, to the full rental contract because they're getting some insurance value. And otherwise-- because life is really risky. And so it's a way of providing-- some combination of providing insurance and incentives.

AUDIENCE: Ben.

BEN OLKEN: Yes?

AUDIENCE: Half the rents in this model, because you don't want to include landlords to extract all the surplus or something like that. Does that change efficiency at all?

BEN OLKEN: It doesn't change efficiency. No, sorry. So what happens in this-- let me think for a sec. So you're saying, what happens if you put in a restriction saying they have to give more rents?

AUDIENCE: Yeah.

BEN OLKEN: Let me think for a sec. So if I increase the rents, will I-- I don't know the answer off the top of my head.

AUDIENCE: OK.

BEN OLKEN: So if you're saying, if I increase the rents, will the landlord compensate somehow?

AUDIENCE: Yeah. I'm just trying to understand what-- I mean, I understand why--

BEN OLKEN: So I think, actually, in this model-- no, bum, bum, bum, bum, bum.

So I'm not totally sure. And the reason I'm not totally sure is that-- in general, the way we solve these models is, we figure out what the gap is, and then slide up and down to hit the rent floor. And that's just essentially, what is the impact of raising the outside wage, essentially.

And the reason I'm not 100% sure is it also affects the risk aversion. And that's what's confusing. That's why I'm not giving you a quick answer, because-- in a concave model, if I move the overall payment up, it makes me less risk averse. And so I'm not 100% sure, actually. Does anyone else know the answer to that question? That'd be a good problem-set question.

I don't know. Sorry. That's a good question. Yeah, Kyle?

AUDIENCE: It's just meant to-- you can't write this directly. It's like [INAUDIBLE] tax, right? [INAUDIBLE]?

BEN OLKEN: I'm sorry?

AUDIENCE: Like you said, you couldn't write this contract just like the alpha model that we saw [INAUDIBLE]. Or is it exact--

BEN OLKEN: Yes, you could. Sure.

AUDIENCE: OK.

BEN OLKEN: Well, so that model didn't have any uncertainty. So what is different in this model versus that model? So that model that had no uncertainty, is the key thing. And the other thing which is important is that the effort choice, here, affects the probability of the-- is related to the uncertainty. So it is a bit of a different model, in that sense. But conceptually, yes, you can think of this as a bit of an-- the solution looks like an output tax.

AUDIENCE: Well, it's just the percentage of output.

BEN OLKEN: I'm sorry?

AUDIENCE: It's just the percentage of output that they--

BEN OLKEN: Well, this is-- so, sorry. So in the continuous case-- so you can write-- there obviously are continuous moral-hazard models, which are even more difficult to work with. But you can solve them. And if you take a contract theory class, you will solve them.

And in those models, yes, I think it will look like some-- you'll get some payment as a function of the output. And it'll look like an output tax. Whether it will be a linear output tax, I don't know necessarily that that's always true.

Although Bengt Holmstrom has a well-known paper in contract theory that shows under what circumstances linear incentives are optimal. So under-- I think the-- and I'm not remembering the details of these models off the top of my head. But my understanding is the general class of-- in the general answer, it will not necessarily be a linear-incentive function.

However, what Bengt shows is that under some things under some various perturbations, where you can do some gaming or things like that, linear contracts often do emerge. That's my summary of that. And you're nodding. Is-- anyone else want to comment on that? All right. Other comments?

So I think-- I don't if we still have an example in the problem set. We should make sure-- we don't? That's gone?

AUDIENCE: There's no problem set on this topic.

BEN OLKEN: OK, so if you would like to work through an example of this, I do think we have a previous problem set, that we have put together on this problem, that has an example of how to work through these moral-hazard models. Yes, Davey?

AUDIENCE: Just with up here, the agents have sufficient constraints t minus, right?

BEN OLKEN: Yeah.

AUDIENCE: It's no service to [INAUDIBLE]

BEN OLKEN: Right. Yep. They're still going to slide at that-- it's still going to be the case of the IR constraint is still going to buy. There's going to be no surplus here. Yeah? Comment?

AUDIENCE: You're seeing-- you've seen these kind of contracts in-- for instance, the tax got cut. You pay a fixed rent and then you get everything-- say the first best is-- in fact, in the taxi business, we see the first pass.

BEN OLKEN: We see both. But yes. With Uber, you see a sharecropping contract. They don't call it "sharecropping," but you see a share contract, where it looks-- where Uber takes 30% of-- or whatever the percentage is of the thing. And they get a fixed-- some share.

AUDIENCE: Is there a reason why, say, like supr--

BEN OLKEN: But-- sorry. But sorry, one other thing. But in the Uber contract, it's observable. And I think it may depend-- I think-- I don't know the full economics of the taxicab literature. But I think it probably depends on to what degree-- certainly, to what degree effort is observable would affect what these contracts look like. And Davey is nodding since Davey is actively looking at exactly these issues of observability and the endogenous contract choice in Cote d'Ivoire at the moment.

And we just lost my computer. Sorry, go ahead while I get my computer back. Yeah?

AUDIENCE: [INAUDIBLE]

BEN OLKEN: What? Sorry.

AUDIENCE: I was going to ask, in what kind of industries will expect to see, say, [INAUDIBLE]-- by the first of this kind of contract where you--

BEN OLKEN: So you would see-- so for example, if you had no observability of output, then you would see rental contract. So this contract requires the principal to be able to condition a payment on the output, for example.

Or you also might think that taxi-cab-type drivers may not be that risk averse. So if the uncertainty is day by day-- this is over a whole season. It's one giant realization. You're really risk averse about that. Some days it rains. Sometimes it doesn't rain. I may not be right risk averse about that. I might be able to smooth over days. So that's another reason it might be different.

Other questions, comments? Yeah, Shera.

AUDIENCE: [INAUDIBLE]

BEN OLKEN: Sorry, Shera. You're going to have to speak up. I can't really--

AUDIENCE: When the farmers are risk averse, it's first best to share even if the effort was perfect, [INAUDIBLE]--

BEN OLKEN: Yeah, sorry. So when I say first best, I meant first in terms of output. But yes, the socially-optimal contract includes sharing risk. But it does not get the first best-level of output.

AUDIENCE: So here, [INAUDIBLE] doesn't have the first-best [INAUDIBLE].

BEN OLKEN: Doesn't-- what? Sorry?

AUDIENCE: [INAUDIBLE] hasn't combined-- does the fact that can't observe effort, does that change the allocation? Because it seems like it does, like first-best are still coinciding with second best.

BEN OLKEN: No, no, no, so-- no, no, no. So if we could observe effort, what would be the first-best contract? First-best contract would be we write a contract based on effort. And we would observe your-- but we'd provide you full insurance.

So the first-best contract would have-- if we could implement-- if we could contract an effort, we would say put in e equals $\frac{bH}{c}$. And I'll pay you h equals I equals the expected-- equals w lower bar, basically. Yeah, that would be the first-best contract.

Other questions? OK. So that's the risk-aversion version. So what about the limited-liability version? So let's go back, now, to risk neutrality for a second. But let's impose limited liability.

So let's assume that we can't impose I than 0. So now what happens to the optimal contract? Well, now the landlord has lost a degree of freedom. The landlord, before, could choose the little h and little I . Now they can now they have this problem that little I can't go below 0.

So in the first-best kind of contract, we had this. What happens-- suppose we're in the world, where this thing is-- if this thing is positive, then great. We have no problem. But if this thing is negative, then we can't write that contract anymore because in the low state, they just won't pay. Another way of saying is, they won't pay the rent.

So what do we do? What does the landlord do? So instead, little I has to be equal to 0 because it can't be-- it's bounded by a limited liability. So what happens to h ? So now, where we know the optimal choice of effort. So the landlord is going to say, well, I'm just going to choose-- all I'm ever going to get, if the landlord-- is a payment in the high state. I'm never going to get any money in the low state.

So my-- but I know, still, the agent is only going to be incentivized based on the difference between the little h and little l . So now I'm stuck solving this contract, which is, I'm going to maximize my return, which is the probability of the high state, times what I get paid in the high state. And the low state is just a wash. I don't get anything in a low state.

So that's this problem over here. This is h over c . We know that the tenant is going to choose h minus l over c . But their l is equal to 0. So the tenant's going to choose a little h over c times big H over c .

And so the optimal choice, here, for the effort, for the landlord, is to choose little h equal to big H over 2. And therefore, the optimal amount of effort is going to be big H over $2c$. Everyone see that? All right.

So again, we have less than the first-best amount of effort. First-best amount of effort was big H over c . Now we only have big H over $2c$. So we have half of the optimal amount of effort. So we're also in a world, here, where--

Sorry. And-- so that's the point number one. And point number two, this is sharecropping contracts because in the high state, you keep half, I get-- the lower gets half, tenant gets half. In the low state, nobody gets anything. So this looks like a perfect share contract-- in this simple example, it's a share contract with the share equal to $1/2$.

So now, what happens if we have an IR constraint here? Well, so what's the farmer's utility under this contract? Well, we can solve the farmer's utility under the contract. Farmer's utility the contract is their choice of effort-- we can just solve. Their choice of effort is going to be a little h over c . And they're going to get to keep payment h , in the state. Plus they have a cost of effort. So if you solve this out, it's $1/8$ big H square over c , whatever. It's all thought out.

The key point, though, is the following, which is that's the landlord's utility. I'm sorry, that's the tenant's utility under this. If the tenant is getting at least w lower bar, then under-- you see with utility, great, we're done. Fine.

Now in this case, the tenant also is getting rents. If this thing is exactly equal to the w lower bar, fine. But that's like a knife-edge case. Generically, you may have parameter values where you have to pay the tenant above her outside option in order to get her to work the optimal amount.

And what's different here? The difference is in the other thing, you could just slide the whole contract up and down, arbitrarily, so that the expected value was equal to w lower bar. But now you can't because you have this limited-liability constraint from below.

So she may get rents here. And also, the contract doesn't depend on w lower bar anymore, in this piece of it. As long as this thing is well above the lower bar, we move the lower the bar up and down, the counter doesn't change, whereas in the more higher one, it was always moving around so that she got w lower bar and expectation.

Now on the other hand, if this thing is less than the lower bar, now the tenant has to actually-- the landlord has to increase the share that the tenant gets to keep in the high state, so that she gets at least her outside-option expectation.

So what will happen, in that case, we'll just pick h , such that the expectation in the high state is just exactly equal to the-- sorry, the expected value of the contract is exactly equal to the outside wage. And note that, in this case, increasing w lower bar actually increases effort because the contract becomes-- as her outside option gets better, the contract becomes more attractive. Clear?

So both-- key point here is both moral-hazard insurance rationales and limited-liability rationales, both generate what look like sharecropping contracts. But there are very different kind of stories for what's going on. And they have different implications.

And in particular-- so both of these models generate the fact that sharecropping emerges, effort is less than the first best. But one thing that's really, in particular, kind of interesting here, is they differ in their implications for land reform.

So as an aside, I feel-- and I'll talk about this in the next slide-- I feel like land reform is, in my opinion, a hugely understudied topic. I feel like one of the salient issues of the 20th century was should we do land reform. In some sense, that you can think of like the whole-- you could think of the whole communist revolutions were really about this question of, should we do land reform or not?

And they wrote many other things, too, but should we take away land from the landlords and give it to the farmers? It was a huge question in the world in the 20th century. And I feel like we don't have great-- I think the level of evidence we have is not commensurate with the importance of the question, in my view.

But in any case, these models differ a lot, in terms of the implications of land reform. So what do I mean by that? So in the risk-aversion model, we do a land reform.

We say, OK, now, tenant, you're the owner of this land. That does make you wealthier and makes you maybe less risk averse. But on the other hand, your risk is all kind of tied up in your-- your wealth all tied up in the land. So it's not easy-- that's not clear you can get consumption-smoothing benefits from that. Maybe you can if you borrow it. Maybe you can't.

But you're still-- would want some insurance. And so you might, as a tenant, say, well, I don't want to be in this position of being the owner and bearing all this risk. I might prefer, actually, to be in a contract where I'm sharing some of the benefits and some of the costs-- or sorry, sharing some of the returns in return for some kind of insurance.

So you could imagine, in a risk-aversion model, that some kind of these sharecropping arrangements might re-emerge endogenously. You can imagine, for example, a landlord might make you an offer and say, I'll buy the land from you. And then I'll provide you a share contract.

And that might be better-- including the payment they get upfront for the value of the land-- they might be better off for that because then they're getting some insurance values later. And there's a trade there. But because insurance is helpful, they might emerge endogenously.

Under limited liability, no. Under a limited-liability model, the only reason we had the sharecropping was we couldn't make this payment from the tenant to the landlord in the low state. But if I'm the owner, I'm good. I don't pay anybody.

And I'm not risk averse in that model, in the risk-neutral model. So which of these models is important would have implications for what you would do in the context of a land owner. OK, comments, questions? Yeah, Kyle?

AUDIENCE: Do we know about that question, whether it's good to give the land back to the [INAUDIBLE]?

BEN OLKEN: I'm going to talk about that in a few minutes. I'm not going to talk about it much. But I'm going to talk about a little bit. There are a couple of papers-- I'm still waiting for the great land-reform paper. There are some really nice papers that I'll talk about. But I feel like they have their various challenges. I'll talk about a few of them in a couple of slides. Yeah, Becka?

AUDIENCE: If both effects are present at the same time, does this just default to the [INAUDIBLE]?

BEN OLKEN: I'm sorry. I didn't understand.

AUDIENCE: If both effects are present at the same time.

BEN OLKEN: No, I think he would do some-- hold on. Let me think. Let me think. Yeah, it's got to because if you're capped at-- the whole limited-liability thing-- if you add risk aversion to the limited-liability thing, yeah, you're still going have a problem. You can't pay the rent in the low state.

So yeah, I think it's-- I think if the thing is binding, then that limits your ability to trade off a fixed fee for insurance value. So, yes, I think so. Other questions?

OK, so now, that brings us to the paper you guys wrote for today, which is, do we actually think that they're-- so one implication of these models is that there is some-- that the incentive-- that the contract is going to affect kind of the effort, choice.

And you don't actually need the full-- that's true under any of these models that we talked about. You can go back to the very simple one, the two-line one I put up in the very beginning. But that's a key parameter that we need to understand all this.

So before you actually got through an experiment, before we had an experiment, if you didn't have an experiment, which is what we talked about, how would you how would you go about estimating this thing? What would you want to estimate this?

What kind of regression do you want to run? Somebody? Yeah, Paulo?

AUDIENCE: You may want something that's like output based, on what share of output do they have if they had-- can't tell the individual farmers--

BEN OLKEN: What's that?

AUDIENCE: So you want a regression where your output variable is the output of the farm. And [INAUDIBLE] only have an input on one of your regressors, what share of output or profits they are getting, essentially, the terms of the contract.

BEN OLKEN: Sorry. I'm sorry. I don't fully-- say it again. So you-- yes, you won't put output on the left-hand side. What you want put the right-hand side?

AUDIENCE: Basically terms of the contract?

BEN OLKEN: Right, exactly. Right. The same thing you would do-- you'd like to regress what is the output on what are the terms of the contract. But the problem is, of course, that those-- those farmers are going to be different.

So one thing to do, actually, would be if you had multiple farmer-- if you had farmers who had multiple plots-- so the basic problem with that regression is we have farm-- there might be farmer characteristics that might be different. Farmers who enter into those different contracts might be different people.

So one option would be to use farmers who have multiple plots, some share crop and some owned, and run this regression with farmer-fixed effects, and say, are they putting in more effort on the owned plots compared to the shared plots for example, or more inputs, or whatever?

So would that be a good regression or a bad regression? So that regression would basically-- sorry, Erica. Go ahead. You can answer the question.

AUDIENCE: Oh, could there be an issue with some qualities of the land, that cause it to be [INAUDIBLE]?

BEN OLKEN: Absolutely, yes, so there can be a problem with quality of the land. So it may be that different land is owned versus rented, and therefore, that could be related to the issue. And so in fact, the regression you could run is some output as a function of whether it's owned or not, with farmer-fixed effects. And this could be either inputs or outputs.

So the problem is exactly this one, that you mentioned, of quality. So this is from an older paper. And you can basically see that they have-- looking among the owned plot-- this doesn't have the fixed effects but I think the same thing goes through the fixed effects-- if you compare the owned plots, the sharecrop plots and the fixed-rent plots, the owned plots look a lot better, in terms of their, say, average value, than the sharecropped plots. Actually, in this case, the fixed-rent ones, I think, look the worst of all, at least in some of the villages.

So the point, though, is that you may have different-- and by the way, this is too small to read. That's why you have the handout. So that's why this regression may be challenging.

And so the original approach is, before we had this RCT, in the paper that we used to teach here, was to say, let's control for really detailed plot characteristics. And the problem here is that there may be unobservable characteristics of the plot, beyond the measurable plot characteristics.

It's really hard to-- if you end up writing an agricultural paper, you'll go into the weeds, so to speak. And you will find an incredible amount of-- how you describe a plot characteristics and soil quality and irrigation, water, quality and this, that, and the other thing, it's many, many, many dimensional. And so that approach is really hard to get that right.

So what would be cool would be to have plot-fixed effects in this regression, like maybe some plots change. But then, of course, the question is why they're changing. So all of this, basically-- actually, I can't find a paper that does that because plot-fixed contracts don't change. And we don't know why they don't change.

So that all leads us to the paper you guys read for today, which is, I think, why this paper-- even though, as some of you guys noted, the setting is a little unusual, it's these kind of-- whoops-- new farmers who are interested in farming, as part of these women's groups.

And so the external validity of this may not necessarily be 100% there. The reason this paper got so much attention is it is the cleanest paper that we know of, that actually tests this.

So as you guys saw, what are they going to do? They're going to work with tenant farmers in Uganda, randomize them to receive either 50% of their output-- so after they recruit them, so it's going to be surprise randomization. So there's no selection issues.

Have either 50-- so the baseline is 50% of output. Or you can get to keep 75% of your output, or 50% of your output plus an exogenous cash transfer, calibrated to be equal to 25% of output and expectation, either is a fixed-fee or a risky contract. So that's the basic design.

So what do you guys-- I'm curious to know what did you guys think of this. You like this-- you like this paper or you have concerns about it? What do you think? Yeah, Erin.

AUDIENCE: So one thing I had was with the-- because they're trying to test for the fact that these incentive effects do explain what they observe, not income or risk effect. And so they argue that their risky, exogenous cash transfer is a way of making sure that these risk effects don't exist. What I-- I know this thing that's-- they draw really heavily on this more commercial kind of model, but that in that model, what you see is that the probability of having a high output is dependent on your inputs, whereas if their's-- for the risky, exogenous case-- and it's not-- it's-- whether or not get that additional 25%, in the risky case, it's--

BEN OLKEN: Is exogenous, right?

AUDIENCE: Yeah and that was [INAUDIBLE] versus in the 75% of output, that-- actually 25% was what we get. And so I felt that--

BEN OLKEN: So what's the concern, there?

AUDIENCE: Well, given that-- if this additional 25% depends on your inputs, your inputs are going to change, versus if it is not. And so I don't understand how, then, you can say that risk does not account for the [INAUDIBLE].

BEN OLKEN: So does anyone else have-- anyone want to respond to that? Any idea, any thoughts?

So-- no? Does anyone-- yeah?

AUDIENCE: I'm thinking of the big st-- like the 50% was an exogenous cash transfer, where they get the fixed amount. That insures you completely, no?

BEN OLKEN: That is the what?

AUDIENCE: It insures you completely. Let's just say insurance, so maybe 35% [INAUDIBLE].

BEN OLKEN: Yeah, so the point of this one is not about insurance. The point of this is the income effects. So the point of this one is to say, well, maybe they're doing more stuff because they're richer in expectation. Maybe-- somehow maybe they're not quite credit constrained, but maybe they're sort of credit constrained. Maybe they only feel like they can afford to buy fertilizer if they're going to get a higher payment at the end. So that's what this one is for, is the level.

Well, I think I think Erin's concern, if I understand it correctly, Erin, correct me if I'm wrong, was that the issue with this one is that the payment, here, doesn't depend on your inputs. But-- and so it doesn't exactly mirror the riskiness you would have in this contract because you would-- yeah, because here, you would also be putting in more inputs. And therefore your net returns would be different. Is that your point? Yeah.

So I think the way-- at least the way I think about this is to say, this is actually testing a different issue, which is to say, look, let's suppose that the-- there's two thoughts. The first is to say-- so first order-- this contract gives you a lot-- sorry, pluses and minuses. I think what they're trying to say is, yes, this contract is just generally more risky than the other one. And so if you thought that what was happening is I was changing my behavior because I'm exposed to more risk, in some general sense, you would see things that look more similar between this and this.

So I agree with you that it is different. It's not exactly the same in a couple of respects. The first is, this risky contract is actually not as risky as this risky contract because this one, I think, if I remember correctly, is exogenous risk, where this one is-- sorry, independent risk, whereas this one is correlated with the rest of your risk.

So actually, to me, the first-order issue is that one, which is like-- this one is-- and I don't think-- this one-- someone correct me if I'm wrong. This is just an exogenous, risky payment, right? Yeah, it's like a random draw. So one thing that might have been even slightly better to deal with-- this is not your issue, but to deal with this other issue-- might have been to say, well, I'm going to tie the risky contract to, say, average returns in the neighborhood, in your area, so at least it becomes a correlated shock, not an independent shock, because otherwise, this is less risky than this because it's independent. You see what I'm saying? It's not a correlated shock.

To your point, I don't know of a good way-- I see your point that the net risk-- your point is if I write down the utility, under the things, that the net risk is the same because I haven't put in these additional inputs. I don't know of a good way to deal with that, though. Do you?

AUDIENCE: No. [INAUDIBLE].

BEN OLKEN: What are some other concerns-- what other concerns did you guys have about this paper? I know in general, actually, I haven't read all your comments. You all pretty much liked it. But were there other concerns that you guys had, looking over it? Or things you would have wanted them to do differently? Yeah, Paulo?

AUDIENCE: I think that probably one of the things that I was thinking about was, in making this experiment ex ante, and what-- the outcome that you're going to get is either something along the lines of, OK, it's mainly income effects or it's mainly incentive and moral-hazard effects, or some combination of the two, or neither of them.

And so I think for me, just in terms of policy applications, what is one supposed to get from this paper? Had it been income effects, I think there's a very clear answer of, just give these tenants money and then you can be more productive. But given that the answer was this, mainly, a moral hazard--

BEN OLKEN: There's moral hazard--

AUDIENCE: Yeah, so what is going to be the reform that a government or something does, as a result? It seems very hard to just tell everyone to change the terms of contracts if they're tenants or things like that.

BEN OLKEN: That's a good question. Are there comments to that? Did anyone else-- is there an obvious policy implication from this? Yeah?

AUDIENCE: One thing that I was thinking about was, instead of just saying, oh, I'm going to take 50% of your output or only 25% of your output, if you could have some staggered design in your contract, where you would say, I'll initially take 50%, but if you make above this amount in your output, I'm going to take less. So then you have an incentive to keep producing more as you go along. And I would be curious what that would actually look like, implemented, if that would--

BEN OLKEN: Yeah, so this gets back to this whole question of are linear contracts optimal. So I'm going to not go into that in a lot of detail, and refer you to a contract theory class for this question. But I think there's a couple of issues.

The first is-- I think the general issue with non-linear contracts is you want to put the-- you want to put the part where you give people the most incentive at the point where they're most elastic. And figuring that out, with heterogeneity across people, can be tricky.

So it might be-- and if I put it in the wrong place, then I can-- so you want to say, I want to-- that can be a tricky set of issues. And the other thing is, you can also get gaming-- if you can do things across periods, you can sometimes get gaming. I don't know if that's really an issue in sharecropping contracts.

And actually, I don't know of a good-- I'm sure there is a paper. I just can't think of it off the top of my head-- that has taken some of the Holmstrom linear-contracts ideas and said, do these gaming or other issues apply to this sharecropping contract. I don't know the answer to that question. Yeah, Wesley?

AUDIENCE: I only followed the policy division in things because what seems a little bit strange about this particular setting, these would be how sharecropping contracts actually work in practice, is that [INAUDIBLE] does not-- [INAUDIBLE] does not care if it's maximizing his profits, if he's is getting as much out these sharecroppers as it possibly can.

And so even if the result is that if you were to increase the portion of the output that your share properties can keep, there's not really a way to impose that on landlords who are looking to extract as much as they can out of their contract, unless, similar to what you were saying, if there was like some income effect, where a government-- or if your landlord is the government, who doesn't particularly care about taking a loss and is just looking to increase the output on the land.

BEN OLKEN: Yeah, so I think-- my own take on this is, I don't necessarily think there's an obvious policy implication of this paper because it doesn't actually tell us if it's moral haz-- if it's the insurance kind of-- insurance issue. It doesn't tell us-- it tells us there's some inefficiency going on, in terms of the-- it says that there's some inefficiency going on, where the sharecropping is reducing effort. But it doesn't say-- it doesn't address the question we talked about before, about why is this emerging endogenously. And it seems that we know that in order to figure out the policy answer.

So I think-- to me, this is an important piece of the puzzle. It says, first point is, yes, these contracts have efficiency consequences. So that's actually really important because if they didn't have efficiency consequences, we would just not worry about this whole set of issues, from an aggregate-wealth perspective-- aggregate-productivity perspective.

The fact that it does say that there's like efficiency consequences says to me that, well, OK, now we want to figure out what's driving it. And if you really thought that the issue was about insurance value, is there a way of providing insurance that doesn't have the same moral-hazard consequences?

And there is this a small, but growing literature on agricultural insurance issues, which-- and I'll tell you, actually, the main challenge with agricultural insurance is, I don't want to write-- because of these moral-hazard issues, I can't insure your output, otherwise you don't work.

So what I do is, I write insurance based on your predicted output, which is essentially, I write insurance based on the weather, for example. So actually, you can buy-- in the US or in other places, you can buy weather insurance, essentially. So if it's a drought-- you can buy insurance against a drought. That's an exogenous event that's going to reduce your output. So we can ensure you, to some degree, that way.

The problem with those insurance contracts is the thing we're insuring is not perfectly correlated with your shock. And so they don't provide as good insurance as you might like. And that issue is what's called basis risk, that the thing we're insuring is not perfectly correlated with your shocks. And so if you're interested in that, there are some papers that have tried to think about that.

But if you thought that the issue was risk aversion, then you would be heading in this insurance direction. And if you got the issue with limited liability, then you might head in a different direction, of can we basically make some transfers to people so that they're-- don't have this limited-liability constraint.

Or can we write-- can we help them facilitate them writing-- we might be much-- if we thought limited liability was an issue, we might be much more interested in writing long-term things that look more like debt contracts, for example, and facilitating that thing, which we might not be interested if other things are important. Yeah. Oh, [INAUDIBLE]. I'm sorry. Go ahead.

AUDIENCE: I was wondering, is there a benefit to-- if you wanted to do some staggered levels and incentives, from getting a panel, if that was possible, as opposed to a sequence of cross--

BEN OLKEN: Repeated cross sections.

AUDIENCE: --cross sections, if you're trying to back out what certain types of workers might be, and how those elasticities might differ between groups, where, from the perspective of the-- for policy implications, does that not even matter?

BEN OLKEN: So I don't think-- so you're saying, if we had a panel, and we get repeated observations of each. So--

AUDIENCE: You, changing your contract [INAUDIBLE].

BEN OLKEN: Yeah, so I think it would be-- yes, so I think if you basically found that everybody was the same, if there was no heterogeneity in people's responsiveness at different points in the contract, then you might be more willing to do some of those nonlinear, [INAUDIBLE] contracts. So, yes, in that sense, yes.

But my strong suspicion, you would do that, you would find heterogeneity because there's heterogeneity in most things in life. And then-- so I'm not sure, as a practical matter, it would make a huge difference. Yeah, Sean.

AUDIENCE: Is the right way to think about this paper that, in the framework of the model we just looked at, that is basically estimating the high c constraint, or that [INAUDIBLE] problem?

BEN OLKEN: Mm-hmm.

AUDIENCE: But then what is optimal would depend on a bunch other things--

BEN OLKEN: Yep, yep, yep, yes. Yes, I like that. Exactly. It's exactly saying, we're randomly varying in the contract and looking at how the agent responds.

AUDIENCE: And it's just telling us that the fact that the agent responded, [INAUDIBLE] the IC is [INAUDIBLE].

BEN OLKEN: Is what?

AUDIENCE: It's telling us that that--

BEN OLKEN: It's important.

AUDIENCE: It's important. OK.

BEN OLKEN: Yes, yes. Other comments? Yeah, Ray.

AUDIENCE: You knew there existed heterogeneity. But then the person who's giving the contract could just administer some tests to you, by varying your contract at the start of more sharecropping relationship. So then you will [INAUDIBLE] your time.

BEN OLKEN: Yeah.

AUDIENCE: And could you have complicated [INAUDIBLE] contracts?

BEN OLKEN: Oh.

BEN OLKEN: Well, yes, except you would need to find incentive-compatible ways of getting that information. So I'm playing a game with my landlord and I know that the answers to my game are going to affect my contract, may affect how I play the game.

But I do-- I think you're onto something. This question is like, can we find incentive-compatible mechanisms for getting like better contracts? I don't know the answer to that. But maybe, yes. I don't know the answer.

AUDIENCE: I would imagine if it was long-term enough, [INAUDIBLE] that-- I don't think people don't [INAUDIBLE].

BEN OLKEN: Yeah, so let me-- I keep referring you to contract theory classes. No, in all seriousness, these questions are like, how do I design contracts and how do I get you to reveal information that are going to be-- how do I design an incentive-compatible mechanism that's going to have you reveal information that's going to let me give you a better contract? That-- thinking about these things these models in more-- if you're interested in these models in more detail, that's where we explore them.

OK. And I'm not a contract theorist, although I did take that class many years ago. So I apologize for not having all of the answers off the top of my head. But I have some of them.

All right, a few other things I wanted to say about the paper, first point is-- so this is the 75% contract versus the 50% contract. And nothing else really makes a big difference. So also, going back to, I think, Aaron, your comment of, is the risky thing-- it may not be exactly the same, but it's certainly not the case that this contract looks-- these results look anything like those results.

Plus, also, you actually might imagine it would go the other way. I'm making the world-- I might be-- I'm making my self more risky-- you might actually think if anything-- this could be opposite direction. You actually might make me more conservative, in some sense.

But anyway, it's certainly-- to your point, I agree with-- you're right. But I think also, it's a little bit-- not necessarily-- it may not be perfectly right. But to first order, it's not-- like I'd look at this and I'd say, oh, it doesn't look like it's really the riskiness that's resulting in these effects. They're just so different.

Several of you picked up, in your response notes, this question of randomization inference. So how many of you have seen randomization inference before? Yes? All of you? Great.

AUDIENCE: [INAUDIBLE]

BEN OLKEN: What? Sorry?

AUDIENCE: [INAUDIBLE] recitation order.

BEN OLKEN: Great, so I won't go into that-- any more detail. Great. Awesome. OK.

You see some effects on the intensive margin of capital usage of various inputs. So it does look like some input-- so I do use more inputs. You also see some impacts on unpaid labor, like family labor. Some of you had questions of what this is. I think the idea is, we have lots of people in the household who also are working on the farm.

And I'll talk about-- I'm going to talk a lot about unpaid versus paid labor in the labor section. So let me not answer the questions of what do we mean about paid and unpaid labor now. But let me defer that to the-- there's a paper by Dan Benjamin that I'm going to talk about, in a lot of detail, in the first labor lecture, which is going to think about this issue, like why do we have-- why is there a difference, actually, between-- why might this respond and this not respond?

There's an enormous literature on that question. And I'm going to talk about that. I'm just going to push that over to the land lecture. But we do see that.

And the third thing we see is that we see a change in the kind of crop. So we see maize goes up. And tomatoes go up. And I guess peanuts, I think, can go up in some of the specifications. And they argue that beans are a safer crop in this context. The other ones are more risky crops. And so I take these more high-risk, high-return kind of crops, which is not obvious.

But one thing I just wanted to point out is that the moral-hazard model had this unobservable, or at least noncontractible effort thing going on. And I wanted to ask what you thought, in terms of, if you think about the results, is it contract-- is the thing is the thing that's changing contractible or noncontractible?

You're trying to map the results to the theory. What do you think? Yeah, Sondé.

AUDIENCE: I think they're both [INAUDIBLE].

BEN OLKEN: It seems like it. It seems kind of crop is pretty darn observable. Are you planting like corn or are you planting peanuts or are you planting beans? That's pretty easy to figure out. Yeah. How about some of the other stuff.

AUDIENCE: Maybe the amount of fertilizer [INAUDIBLE].

BEN OLKEN: Yeah, maybe those are harder. So to me, in some sense, one of the things that I think is a little surprising, here, is that at least-- and they do this decomposition thing. And about half of it is-- actually, sorry, sorry. What's observable, what's unobservable, let me not answer that. But they say, what are the obvious-- the inputs the production function, like land, labor, capital, is about half. And then crop choice is the other half.

To me, actually, crop choice seems kind of the most observable and contractible of all of this. So in some sense, that's a bit of a puzzle, actually. Why can't we solve some of these problems by contracting on crop choice? From the first model, I think it's an interesting question.

Other comments? Any other comments on that paper before I move on? Yeah, Davey?

AUDIENCE: It's because [INAUDIBLE].

BEN OLKEN: What? Sorry?

AUDIENCE: Just because to observe it. So all of this can be observed in some ways, but just the customers.

BEN OLKEN: Well, I don't know. These sharecropping contracts, sometimes they're literally paid in kind, literally-- certainly, I've seen examples where you literally deliver x bushels of rice. And then it's pretty observable what the crop was.

I think in general, though, we want to monitor what the output is in order to figure out the half. So I think that they got permission. Yeah, Kaj.

AUDIENCE: I think it would have been nice to measure his preferences just to [INAUDIBLE].

BEN OLKEN: I think that would have been interesting. In some sense, the paper-- going to the previous point, the paper is really about the responsiveness to the contract. It's about the ic constraint, which doesn't really depend on your risk preferences as much. Right? So if it was about-- if this was a paper about ex-ante, what would be the optimal contract, then for sure, that seems really important.

It's not obvious-- it's not as obvious to me how important that is, given this is about how the agent is responding. You see what I'm saying? Yeah?

AUDIENCE: Would it be the agent that's risk averse? What if they try to allocate [INAUDIBLE]?

BEN OLKEN: The agent lied?

AUDIENCE: Like the agent is the problem.

BEN OLKEN: Yeah? It's relevant vis-a-vis they're doing it the other way around, by making the contract, obviously, more risky. So you're right. Yes, this contract does change the amount of risk they're exposed to. But-- so I do-- it is kind of relevant. But if we think back to why did we care about risk aversion so much in all the models we went through today, it was actually-- it was more of vis-a-vis the-- ex-ante contracting choice, not about the ic constraint.

And I liked whoever framed as a saying, this is a paper about the ic constraint. And I think that that's right. And risk aversion, yes, it does matter there, but much less.

OK, other questions? Then we're basically out of time. OK, so we'll stop here. And I'll pick up, on Monday, with the second half of this stuff, which is all going to be talking about-- primarily going to talk about land tenure. And I'll talk a little bit about some of the land-reform papers, just very briefly. And then mostly, I'm going to talk about security, property rights, and investment decisions, things like that. OK, I'll see you all on Monday.