

[SQUEAKING]

[RUSTLING]

[CLICKING]

PROFESSOR: So we have a lot to talk about in the last week in terms of credit. I'm not quite done with the labor section, so what I'm going to do is just spend about 10 minutes just giving a couple of highlights of some of the last-- a little overview of some of the topics that are in the labor notes. You're welcome to review them in detail, and then I want to get into the credit stuff in detail.

So there were two other topics that I just wanted to give you. The first is, there's been some additional attention-- and these are all what's happening in kind of the urban labor markets. So there's been some work recently on understanding matching frictions. And in particular, a big issue in labor markets is the fact that not all jobs are the same, not all people are the same, so we want to match people to jobs.

A big issue that differentiates labor markets from other markets is kind of the rich heterogeneity, both in the employer side and on the other side. And so job search is a big issue. And this question is a big one in developed country labor markets, and I think people are just starting to understand how those issues may be different in developing country contexts.

And one issue that people have started to think about is that part of how we help solve these problems is, we have various signals that help us signal in the market what our skills are. So you get a PhD from MIT. That is a signal that you can take the future employers and say, I know some economics, for example. At least we hope that's a useful signal. But if those signals are weak and the quality of the signals is not as good, there may be some challenges.

So one paper is trying to understand these issues, looking in the South African labor market. And they do some interesting things on this. They look at giving people information about-- and they note that these skill assessments tell you multiple things. In the MIT example, if you get a PhD from MIT, that's a valuable signal in the labor market. It's also a signal to you that actually you're pretty good at economics because you managed to get a PhD in Economics.

So those information, they provide information both to the job seeker and to potential employers. To get at this idea-- so not with PhDs, for example-- they basically help people-- they do a skill assessments of people. And they give them a certificate with a World Bank brand, like a credible brand saying, here's your skills assessment or whatever.

Then they also do a second arm where they give the same information but kind of privately to job seekers. They show it to them, but they don't give them a certificate they can bring out to people. And they show, in order to decompose, what is the role of the job seekers from the alleviation to the job seekers from the information to potential employers. I'm sorry. This is about certificates. So they basically decompose this.

So I'm not going to go through this in a lot of detail. I just wanted to highlight that this area of, how do we think about the job search in general, is an area people are starting to think through and this question about understanding the role of credible signals for employers. And that is one way to look at this. And it's a very nice paper. So I'm not going to go through it in detail, but I encourage you to look at it.

The other thing people are starting to look at is job search and online platforms. So this has been a big move in the developed country labor markets. It's the world of these online job matching programs. Is that actually relevant in these lower middle-income countries? I think people are starting to think maybe it is, so that's something people are working on.

Job training, there's a lot of work on this question of, how do we teach people jobs, et cetera. Is there a general purpose training, specific training? Again, I'm not going to go through this in detail, but I wanted to give you a pointer that some of this work exists. And if you're interested in these questions, you should have a look at it.

The second issue in urban labor markets that I also want to highlight is the rule of working conditions. And this is an issue, I think, if we think both about the development of the Industrial Revolution and the area post the Industrial Revolution. In a lot of countries, like the US or the UK or other kind of richer countries today, or in a lot of lower middle-income countries today, how do you think about, what are working conditions like in a lot of these industrial factories?

There still not perfect everywhere, of course, but there's a lot of regulation, for example, in more developed countries. It helps make you think about worker safety and things like that. If we have large informal sectors that aren't subject to those regulations, what's happening in these lower middle-income countries? I mean, to me, this is a really interesting and really important issue and one that I feel is not particularly studied to the extent it probably should be in the economic literature. So I think it's a really important area.

So as I mentioned, there were huge uproars about this kind of 100 years ago in the US, for example, about some of these workplace tragedies. There have been similar tragic incidents. There were a couple of very famous ones in Bangladesh about 10 years ago. There was this collapse of this factory that had many textile shops, hundreds- or actually, over 1,000 people were killed. So how do we think about some of these kind of workplace issues?

Again, particularly when you have a large-- and why is this kind of a development issue? It's because we think that regulation actually is what helps solve some of these problems in more developed countries. But with these large informal sectors, the regulatory issues, maybe not so. Regulatory solutions may not work so well.

So I just think I wanted to flag that as something I think as a really important area. There was one paper recently by Laura Boudreau, which I thought was a very nice one about looking at these worker management safety committees, and actually randomizing the creation of these into Bangladeshi factories and looking at some issues here. More generally, I think, what is the role of multinationals actually? And are they playing a role in disciplining the working conditions in the factories in lower middle-income countries that are producing their stuff, I think, is a really interesting area.

Again, I'm not going to go through this in detail, given the limited time I have. But I kind wanted to just flag this paper if you're interested in it and just more generally say that I think that this issue about multinational-- issue of working conditions in general, and multinationals in particular, and what role they play is an important one and one that's a good area for future research.

OK, so that's all I'm going to say, and then I'm going to switch gears. Do people have any other comments or questions before we switch to the last topic? OK. So sorry I didn't get to go through those papers in detail, but there's a lot to talk about in credit too. I want to make sure I get to it.

OK, so credit is obviously a hugely important issue. If you think about what we've talked about, if you think about the overall production function, there's like-- we think about a production function that looks like this. We have labor, which I just talked about, human capital, which Ezra talked about in extensive detail. There are some other components of this. One is productivity, and the other is capital.

So understanding what are the challenges and overall production in lower middle-income countries compared to the developed world, you want to understand all four of these components. So we've talked a lot about labor. We've talked a lot about human capital. We're going to talk next semester much more about productivity. There's going to be a lot of discussion about, how do we think about productivity, and how do we think about heterogeneity in productivity across firms.

And the final piece is, we should think about capital. And we're going to split our discussion of capital across what I'm going to talk about today and next time and what we're going to talk about in 14.772 next semester. And in particular, what we're going to talk about is credit constraints.

So normally, we're maximizing this production function, you know, minus-- if you were maximizing something like this, over l , h , and k minus whatever-- wL minus whatever. I'll call it wh of H . You would choose the amount of capital, such that a prime-- sorry-- AF prime sub k is equal to r . You would just maximize. Choose your level of capital, such that the marginal return to the capital was equal to the interest rate or the rate of return to capital, which we think of as the interest rate.

So what we're going to talk about today are reasons why we may not be able to do that. And in particular, we want to understand how does the market for capital or the market for credit work, and are there reasons why firms might not be able to get this much capital. Firms, businesses, anybody who needs capital might not be able to get this much capital. It might be limited in their capital.

So today, I want to talk about the basic theory of credit constraints, and that's going to basically take me all today. So how do we think about why might people not be able to get this much capital? And I'll talk about moral hazard reasons, adverse selection reasons, and monitoring reasons, and we'll think about what does that all mean.

And then I'm going to talk on Wednesday about the role of microfinance and how we think about microfinance as one way of solving these kind of credit-- or not solving, but ameliorating some of these credit constraints for really poor-- not poor-- but typically poor, very small micro enterprises. And there's been an explosion of work understanding both whether or not microfinance does or does not ameliorate credit constraints, and understanding all the many different features that happen in microfinance, and understanding whether those are important.

So that's what I'm going to do this week. What we're, then, going to think about is in 14.772, we'll think about the role of banks as intermediaries. What are banks doing? How are they allocating capital? What are they doing in credit constraints? And more generally, understanding a firm's demand for credit and the impact of credit on firms. You can think of it in some ways as estimating that F_k thing. What is the impact of additional capital on the firm sector?

So I split this topic, either the stuff that we think of as a little more-- this is mostly about people and the basics of understanding these models. I'm going to do now. The stuff about really getting into how this affect firms, we'll do next semester. That's the plan.

OK, so here's the neoclassical model of the capital market. Imagine that everyone faces the same interest rate, adjusted for risk. Now, what that means, by the way, is if there's a d percent risk of default, then-- so we're going to work with what's called the gross interest rate, which is how much-- so you repay the gross, which is often more convenient. The interest rate we think of as like-- so imagine I have a 3% interest rate, so I would have to pay 3% of interest.

The gross interest rate is like 1.03. It's the total amount you have to pay back, and that just turns out to be convenient for the math sometimes. So imagine I default. And different people differ in their default risks. Then what is the bank going to get back? Well, the bank's going to get back-- well, with probability d , it's going to get nothing. With probability $1 - d$, it's going to get the gross interest rate.

So in a basic kind of market, if we have different chances of default and that chance is exogenous or whatever and observable, then I'll just be happy to offer you all different loans, and the gross interest rate I'm going to charge for your loan is $1 - d$ times r . So I'm going to get back my r in equilibrium. OK. And basically, that's kind of what you get. That's what you get, and then the expected marginal product of capital should be equal to $1 - d$ times r because that's the interest that you have to pay. Fine.

So here are some stylized facts, and I'm not going to go into the detail between these, but there's some work that suggests these are true. So we suggest that maybe that simple neoclassical model is not everything that's going on. So then, by the way, one of the implications of this, by the way, is that what's happening to depositors and what's happening to the people who are getting the capital should be small. Imagine there's free entry into banks. Then basically, that gap will be kind of competed to zero because depositors get the highest interest rate, borrowers will get the lowest interest rate. That gap will be very small.

So what do we see? So we see, first of all, there's a big gap sometimes between the deposit rates and loan rates. It seems like there's a lot of money going in that intermediary sector. What's that about? Talk about that. A lot of variability in interest rate within the same subeconomy, and that's true even with very low levels of default. So we're going to see heterogeneity in interest rates that seems to be maybe something else going on beyond just defaults. So we'll see if there's something about that, even though there is ex-ante competition in markets.

And another thing we're going to find is that richer people seem to be able to borrow more and pay lower rates of interest than poorer people. And more generally, it seems like people who borrow more are paying lower interest rates, and that seems a little weird, actually, when you think about it. You might have thought that maybe the more you borrow, the riskier you are, and maybe the default rate should go up. But actually, it doesn't seem that way. It seems like there's actually a positive correlation between the amount you borrow-- sorry-- a negative correlation between the amount you borrow and the interest rate. So that also seems kind of interesting.

So it seems like there's something a little more complicated going on in the credit markets. And so what might that be? So here's a model, a simple model, that of moral hazard in investment choices. So imagine that there is a menu of possible investments that I as a firm can choose. And they each differ in their probability of success, p . And investment p is going to yield some return, r of p , with probability p and return zero, with probability $1 - p$.

So my expected return is going to be e of p , and that's going to be our-- that's going to be the probability of getting the positive return times the return, if you get it. That's my expected return. We're going to assume that r prime is negative, that the product with lower-- the product with higher probabilities of success are going to return lower amounts as they succeed. And I think that's a pretty general assumption because, otherwise, there would be dominated projects. So we can choose the best project for everybody, and then it would happen. Otherwise, not very interesting.

And we're going to assume that that's going to yield. And then we'll assume that this e function is concave, so there'll be an interior level of the optimal project for maximizing expected returns. OK. And we're going to note by p^* the maximizing project, the efficient project. That's the one that maximizes the efficient returns. So that's just maximize. p^* is just the max of this thing. OK, so we would like people to choose p^* in some sense.

In addition to the project choice, the p , there's a return to scale parameter, which is given by f of k . And maybe that's going to be diminishing returns of scale or whatever. So your choice is, I'm going to invest capital k in project p . So I'm going to make two choices, a k and a p , and my return is going to be the-- the p is going to choose the expected p , and then I'm going to scale it up by this f of k .

OK, is that clear? So I could go into making-- I don't know-- belts or t-shirts, that's p . And then I can choose how many of them I want to do. That's f of k . Clear?

Finally, suppose that the investor has some wealth, w . So if the investor wants to borrow-- wants to invest k , she has to borrow $k - w$. So that's the credit. That's credit. It's how much they have to borrow. The gross industry is r . And we'll have limited liability, which is, if you can't pay, what can they do? There's some collateral. Another way to think about is, you put up your wealth as collateral, they can take your wealth but nothing else.

So I'll just tell you, that is important because you can see why this is going to matter. So if I am investing k , and my wealth, w , is approximately equal to k so I'm only borrowing a little bit, then if my project kind of fails, I'm going to lose a lot. They're going to take almost my entire wealth. Whereas if my wealth is, say, zero, then if my product fails, I get nothing in that state so I'm only interested in the upside. So the amount of wealth, the amount of collateral I have, is determining what share of the downside of this project I have to bear versus the bank has to bear.

And this idea is, you can immediately see why it's going to be nice for banks to want you to have some collateral. And particularly why they're going to want you to put some of your own money into projects is because they want make sure that you're bearing some of the downside risks, that you're not taking-- otherwise, you take super risky projects because heads, you win, tails, they lose. That's not a very good deal for the bank.

And so lots of loans have this feature. So for example, if you want to buy a house, you need to put 20% down on the house or something. That's basically saying, they want the loan to be no more than 80% of the total because they want you to have some stake in what happens if things go wrong.

There are more complicated situations in mortgage markets actually because in addition, often when the loan goes bad, it starts from the top. And so if you have 20% down, the first 20% is your loss, and the bank only loses it more than that. So actually, the ordering of it is important too, but we're not going to get into that in this model. OK, is this clear so far?

OK, so suppose that we can contract over k the amount of capital, but we cannot contract over the project choice, p . So k is going to be fixed, but then the borrower is going to-- for a given k , once we agree on k , the borrower gets to say, OK, now I know what my k is, I'm going to choose my p . And they can't contract over that.

So in particular, what are they going to solve? Well, with probability, p , they're going to get f of k times the return. So the f of k times r of p -- that's how good the product is-- times the scale. And then they have to pay the interest on the amount they borrowed. With probability $1 - p$, the bank takes their wealth and nothing else. So another way, you can also equivalently rewrite this as saying, they paid the gross interest rate on k in all periods, and then they lose their wealth in this other period.

So the first order condition is that the f of k times e prime-- e prime is the derivative of the expected value with respect to p -- has to be equal to the effective r times k minus w . This part doesn't matter. And so we're just going to get-- so p times r , that's e , so the derivative. So p times r was e , the expected return, so it's just f prime-- I'm sorry-- f of k times e prime of p is equal to-- there's a p in here, so that's just equal to r times k minus w . OK, so that's the first order condition.

So recall that the optimal p had e prime of p equal to zero. That was a socially efficient level of p . So what you can see is that whenever w is less than k , whenever you're borrowing something-- so if w is equal to k , by the way, if it's all your own capital, you pick the socially efficient level of a socially efficient project. Why? Because the cost and the benefits you totally internalize. Whenever we have a wedge where you're gaining more of the upside than downside, that's going to lead you to take riskier projects.

And how do we see that? We see that basically that's because e prime is greater than zero. Because in order to make this thing positive, the-- oh, no, sorry. You're going to take more risk. Sorry. Because e double prime is less than zero, we're going to take a little bit more risk in order to make this thing positive. So the optimal project is going to be a little bit riskier in order to make e prime positive.

And that's what I was intuitively saying before. Whenever you're not perfect, whenever you're getting more of the upside than the downside, that's going to lead you to take too much risk than is socially efficient. Is that clear? OK. That's basically what I said. OK, clear so far? Yeah? All right.

So this already gives us some comparative statics. So this is the condition that I just derived on the last slide, that f of k times e prime of p was equal to r k minus w . So we can just divide through by a -- if we can pull out a k and make this $1 - w$ over k , divide through by f of k , and we get the second expression here. e prime of p is equal to k over f of k times r times $1 - w$ over k .

So why is writing things in that way helpful? It's because we can immediately start to see a few comparative statics. So the first one is, we can see how does p prime depend on interest rate. That's our p hat. How does the equilibrium choice of project depend on the interest rate?

So when interest rate goes up, e prime goes up, therefore p goes down. So when the interest rate is higher, we're going to be choosing riskier projects, again, assuming that we're in a world where you're borrowing something. Obviously, if w is equal to k , then we're just at the socially efficient level. But assuming w is less than k , we have that higher interest rates are going to lead to riskier projects.

Clear? By the way, stop me if I'm going too fast. Sometimes when you go through a model on the slides, it's too fast. But if anything, just slow this down.

STUDENT: Higher interest is called higher r ?

PROFESSOR: Higher r .

STUDENT: And higher p is higher risk?

PROFESSOR: No, lower p is higher risk. So higher r means higher e prime. But we have e is concave, so we're going to go we're going to go down to get a positive.

STUDENT: [INAUDIBLE].

PROFESSOR: Yeah, it's the same thing because f prime is concave, and the higher marginal product means you have less capital. Same idea. Basically, we have higher levels of e prime. We're going to get higher levels of e prime with lower levels of p because e prime was concave. Sorry, e was concave.

Similarly, we can say, how does p depend on the returns. Returns is f of k . F of k divided by k is, what is the return on capital right. Like you invest k , and you get back f of k , so this is like the average return on capital. This is going to be increasing for the same reason because this is 1 over that. So it all goes backwards. So higher returns leads you to choose higher probability projects.

And this one, I think, is-- these two are certainly going to be the most relevant for us. The third one is, how does p hat depend on k over w , where k over w is leverage. k over w is-- obviously, this is 1 over k . How does it depend on what fraction you're borrowing? Well, we know that it's going to be decreasing. Again, we know that when k is equal to w , we're at the optimal level. And as we start to increase k , we're going to be choosing risk here. Holding w at face, we're going to be choosing riskier and riskier projects. Clear?

STUDENT: What's the intuition for the second point?

PROFESSOR: So I think it's a similar one, that basically the-- actually, I don't have a good intuition for the second point. Does anyone else want to give it a shot? Sorry, I take it back. I don't have a good one. It's a little hard to think about because you're holding w over k fixed. Hold on one sec. Let me see if I can give you a good intuition.

I'm sorry. No, I don't have a good intuition. Let me think about it. Anyone else want to try? Ed, do you have a good intuition for it? No? Sorry, I should think about it. Sorry I don't have a good answer for it. OK. But luckily, I'm going to mostly talk about r and the other ones, which I do have a good intuition for. Yeah, Hazel?

STUDENT: So we're assuming that the borrower is risk neutral, but if we assume risk aversion, would we-- so like taking too much risk-- so if we assume that they were risk averse, could we get rid of that in the model?

PROFESSOR: If we assume they're risk averse, it's going to lead them to take less risk. Hold on. Let me think. What? Sorry.

STUDENT: Sorry.

PROFESSOR: No, go ahead, Hazel.

STUDENT: Because they were taking too much risk before, right? And then if you introduce risk aversion, they might take less.

PROFESSOR: Yes. That clearly sounds right. I think it's still going to be the case they're going to take-- so yes. For sure, if they're risk averse, they're going to take less risk. That's going to shift the whole thing in the direction of less riskiness. And you could get that maybe the optimal thing is to get e^* , but e^* may no longer be-- the same p^* may no longer be optimal because it's no longer optimal for them.

So I think it's still going to be the case that if you think about socially optimal including the risk aversion, like the riskiness to them, it's still going to be the case they're going to take too much risk relative to that. But yes, you're right. Certainly, if we add risk aversion, they're going to take less risk, and that's going to shift the whole thing up.

STUDENT: Intuition is just to care about f times e . And if f goes up--

PROFESSOR: You care about what?

STUDENT: You care about f times e to [INAUDIBLE] objective function. If f goes up, then you're more invested in the return once your repayment [INAUDIBLE].

PROFESSOR: Oh, I see. This relative to that?

STUDENT: Yeah.

PROFESSOR: OK, there you go. Yeah, I think you're right. OK, thanks, Ed.

The other thing I want to note here that comes out of this model, which is out of many of these models, is there's a positive probability-- there's a positive correlation between the default probability, which is $1 - p$, and the interest rate. So in general, we had higher interest rates lead to riskier projects, and we get that in this model. That was this prediction over here. And that positive correlation between the interest rate and the riskiness is a classic prediction of many of these kinds of asymmetric information models, including more hazard models like this one. Clear so far?

So the final thing that we need to do is, we need to add an exogenous supply of funds for the lender. So the lender can go out to the capital markets-- call it Wall Street or whatever-- and they can borrow at rate, ρ . So they're only getting repaid, of course, with probability, p . So their equilibrium condition is that their cost of funds, ρ , has to be equal to their expected return. With probability p , they get r . With probability $1 - p$, they don't. And so that's going to determine their return.

So what is it going to look like for equilibrium? So there's two different things going on. The first would be this market clearing thing, which is that ρ is equal to p times r . So this is r , and this is p . And then we also derive from the previous model this moral hazard kind of equilibrium, which also gives you p as a function of r . So as I've drawn here, there's a single equilibrium. This is going to be the equilibrium in this market where we're going to have projects, p , down here and this interest rate. OK.

But this thing is-- I mean, obviously this is not drawn exactly right-- but this function here, I guess this hyperbola, given by this ρ equals pr , that's fixed. But the shape of this particular p hat of r is not obvious what it's going to be. It's going to depend on the underlying things that went into that previous model. In particular, it's going to depend on this e function, which could be-- we haven't specified what that's going to be.

So we could draw it like this with a single equilibrium. It could look like this, which has-- I mean, we know that it has to be decreasing, but the shape of this thing is not clear. So if I've drawn it like this, there are three equilibria, where these are both, I think, the stable equilibria. And here, there could be two equi-- so here, there could be two equilibria in this credit market. We could have one where we have really high interest rates and very risky projects, or we can have very low interest rates and not very risky projects. Both of those could be equilibria, but it's a multiequilibrium model.

Or you can maybe have a no lend equilibrium where, again, you have to draw both of these things going out like this. But it's not obvious these things are going to start to intersect. And so in this model, the credit market could collapse entirely if there is no intersection between these two points. OK.

So this already gives us some important things going on. It can already be an explanation, for example, why you might have multiple different things going on. You could have really high interest rates and really terrible projects, for example. Or sometimes you can have credit markets that collapse entirely. And it just depends on the shape of that p hat of our function. OK.

What's happening with capital? Well, you can think of us as having derived that optimal p hat function, which is a function of the returns, the leverage ratio, and the interest rate, which is in itself an endogenous function of k . So now, the bank and the lender know that conditional on a level k and an interest rate, they're going to pick project, p . They can't contract over p , but they know what that p function is. And so they can make a choice collectively like, how much capital would you like to borrow.

So then I can solve. I can maximize over capital the return here, where I'm going to maximize the amount I get times the expected return minus the amount I have to pay back. And here, I've just substituted in for the fact that in equilibrium, I have to pay back ρ . So with probability, p , I have to pay back r times-- I have to pay back r , so I know this is just pr , so I just substitute it for ρ , the market amount.

OK. What? Sorry.

STUDENT: Is it still the same borrower making this decision, or is the bank deciding on it?

PROFESSOR: This is contractable, so they can decide together. The borrower can decide it, but the bank will be happy to do this in equilibrium. The key thing we had to keep track of is, once they choose k together or whatever-- this is a contract-- then the borrower goes off and chooses p . So the key point is, we're choosing the level of k , taking the borrower's second period choice of p as that. That function is given.

And that's what I just said. We're only maximizing this over k . We can't choose k and p jointly because we know that once we choose k , I'm going to go do whatever p I want. And we might do better if we commit to a level of k and p together, but we can't. OK.

So that gives us a first order condition. So what's the first order condition? Well, it's just F' of k times e of p plus the other term, F of k times e prime of p times p prime, so dp/dk . And all that has to be equal to ρ . Clear? By the way, thanks to Ed for helping me clean up some my algebra on this slide, which I had some issues before. But I'm pretty sure this is correct. So that's what we just had.

What would be the first best? The first best would be that the overall return, f' prime of k times the expected return, should just be equal to the market interest rate. So that would be the optimal level of capital to invest. And we can see what's going to happen to the level of capital. So in particular, we can rewrite this equation as f' prime e of p is equal to ρ minus this term over here. OK.

f' of k is positive. e prime is positive. We had that from one of our assumptions before. And we also just showed earlier that dp/dk is less than zero. OK, because of, say, for example, the leverage channel. So what this means is that if we have a concave function, a concave production function, then the level of investment we're going to have is less than k^* , the optimal level of investment.

Intuitively what's going on, what's going on is, the interest rate is kind of high, given the moral hazard, and there's going to be more of this with more capital holding this rate fixed. So the interest rate is already high. That's going to drive down the optimal level of k , and then the additional channel is only going to reinforce that. So the interest rate is high because we have the moral hazard problem.

And then if the higher interest rate I choose a lower level of project, that makes the interest rate even higher, and we end up in a-- the equilibrium of that drives capital even lower. And so, again, that says, look, we prefer to be able to commit to this level of p that wouldn't have this problem. It would give us a lower level of interest rate and a higher level of capital and higher level return, but we can't commit to that. We can also show similarly that dk/dw , is greater than zero through a similar argument.

Just stepping back, what have we shown? This is a moral hazard model where the choice of projects was uncertain-- sorry-- was uncontractable. And that plus limited liability led people to choose these risky projects. The choice of risky projects leads to these higher interest rates, which leads to lower levels of capital in equilibrium.

STUDENT: In the real world, don't people often contract over p when they're like, oh, what do you want the money for? For this business idea. OK, no I don't like your idea. I won't give you the money.

PROFESSOR: Yes. So I think you're exactly right. A whole lot of what banks are doing is trying to-- you can think of a lot of what banks are doing-- well, let me also talk about the adverse selection model too. It's not obvious to me how much they're doing they're screening about the adverse selection. They're definitely trying to figure out p and whether it's the ex-ante. Whether it's the moral hazard part or it's the adverse selection part is less clear to me, so we can come back to what they're doing in the real world.

I think a lot of what they're doing looks more like the adverse selection part. So the selection part is kind of the ex-ante, which we'll talk about in the next model. Ex-ante is saying, hey, Tishana, what are you planning on doing with the money. What's your p ? You can think of that as like-- that's ex-ante, trying to figure out, we have different p 's. Your p is this, Aaron's p is this, and so on and so forth. We're trying to figure out who's got the good p 's and lend to them. That's adverse selection.

This model will basically say, I want you to come in here every month, and I want you to tell me what you're doing. Kind of that very, very regular monitoring. Banks also do that, but that's the closer to this one. So I think you're right. I think that you can think about the whole-- a lot of the relationship between banks and their clients as being about trying to combat these two basic issues. And if they're perfect at combating them, then this goes away. In reality, I think we probably think they can combat them to some extent, and this is capturing the residual.

Other questions or comments?

So question. Would you say there are credit constraints in this model? What do we mean by credit constraints? We're going to talk a lot about credit constraints. Do we think people are credit constrained in this model? It's kind of a little definitional, but I'm curious. What do you-- Yes? No?

STUDENT: I would say, no, because they can borrow as much as they want.

PROFESSOR: OK, no, because they can borrow as much they want. They're just choosing k as much as they want. What would be the argument to say that they are credit constrained?

STUDENT: Because they're getting less than what the borrower gives.

PROFESSOR: Yeah, exactly. So that's exactly right. So I just want to be clear on this. This is a model where I think you're exactly right, Cash, that basically this is a model where they're not constrained, in the sense that they could get more k if they wanted to. The optimal level of k for them to borrow is whether the interest rate will adjust as they borrow more k . And that's going to lead there to be-- it looks like they're not borrowing as much k as they might have wanted to if these information asymmetries hadn't been there. So that's model one. That's the moral hazard model.

Here's another model, which is the adverse selection model. This is an alternative story. And going back to what I was talking about with Tishana, this is kind of like screening heterogeneity among projects. So now, in this model, we're going to suppose that p is a fixed characteristic of an individual. Some people have high p , some people have low p . OK.

So in this model, we say, look, suppose the bank offers an interest rate, r . Who's going to take that loan? Who wants that loan? Well, I'm going to take that loan if, again, same little liability. The project fails, probability $1 - p$, I pay nothing. Probability, p , I get f of k of r of p . That's my return. And I have to pay this gross interest rate, r , on the amount I borrow, which is k minus w . So p is going to drop out, but if r is greater than-- so big R , the return, is greater than r , the amount I have to pay back, r times the amount I borrow, divide it by f of k , then I'm going to want that. OK.

So if you recall from the beginning, we had that r prime was less than zero, that we had the lower-- the riskier projects were the ones that were paying the most. The higher return projects were paying less. OK. What that means is, there's going to be a cutoff value where people with p less than some cutoff value are going to be the ones who are going to pick-- going to take up the loan.

r of p less than some cutoff is the same as saying that p is going to be less than some cutoff, so the super risky projects up to some cutoff are the ones that are going to take the loan. OK. So we can define \tilde{p} of r as exactly that cutoff, so we can just define \tilde{p} , such that r of \tilde{p} is exactly equal to this. And everybody with a p less than \tilde{p} will borrow. OK.

And so what happens? Well, as r -- so I increase the interest rate-- the \tilde{p} is going to shift to the left. Because if I increase the interest rate, the return's got to be higher. Higher return projects were the ones with the lower p . So again, higher interest rates are going to shift that threshold to the left, and we're going to get riskier projects.

So we're also going to get that same correlation that we had from the moral hazard model, which is a correlation between the interest rate and the riskiness of the projects, although it's happening through this different channel. One was this ex-ante, who wants the projects. The other is the ex-post, who is the interest rates.

And intuitively what's going on in this model, what's going on here is like, look, imagine the return is the-- in the good state, I have to pay back a lot. Who's going to want that? Only people for whom the good state is really, really good, and those are the people who were risky. So again, we have this negative correlation here.

Now, what's the market equilibrium here? Well, again, we have that ρ , that r times the amount you get-- so here, we have a slightly different condition. What is the p ? For a given r , who borrows? Well, it's everybody with interest rate with p less than that cutoff value. So what's the average probability of success? It's just the expectation-- this is e meaning expectation, not e meaning the expected return. So sorry for the imitation.

It's the expectation of p conditional on p less than \tilde{p} . We know that as \tilde{p} goes down, the expectation of p goes down. But beyond that, it's the shape of that thing depends on just what the distribution of p 's is in the population. OK. That also is going to mean that r , the interest rate charge, is going to be equal to ρ divided by the expectation thing. If this thing is less than 1, that means that the r 's that going to be charged to people are going to be higher than ρ . So we're also going to have higher interest rates than equilibrium. It's in the first best or whatever, then that ρ thing.

But the reasoning is a bit different. It's all about this ex-ante screening. Going back to your point, Tishana, this is where the banks are going to want to invest in doing everything that we can say to screen people, kind of ex-ante and get rid of the people who are going to be kind of lower ρ . Because for a given interest rate that the bank's charging, conditional on charging r , if it can weed out some people who are the worst, it's going to do better. That's not necessarily equilibrium, but that's going to be the bank's incentive function.

So I announce r . I know that everyone with p less and \tilde{p} would like it. How can I make even more money? I'm going to be able to break even, but I can make even more money if I can find the really low people and knock them out. I want to get people who are just as close as possible to \tilde{p} . And you can think, a lot of what banks are doing in that screening process is trying to figure out who are the good borrowers, who are the bad borrowers, can I remove the really bad ones.

So again, we can draw that same kind of equilibrium. Here, this thing is the same market clearing condition, but now, we have a different-- instead of a moral hazard condition, we have an adverse selection condition, which says that here is the expectation of p , given that p is less than \tilde{p} , where \tilde{p} is a function of r . And so this curve here doesn't come from project choices.

This is just a question of, what is the distribution of p 's in the population. This is just like a property of the p function of the population. What is the distribution of p 's And again, you can draw that in funky different ways. It can look like this. It can intersect. It can go on with equilibria and so on and so forth. It's got to be downward sloping, but beyond that, unclear. OK. Clear so far? OK.

So both of those models, the moral hazard model and the adverse selection model, are going to feature high interest rates, or at least both of those are reasons why the interest rate will be higher than the market clearance. This is why the market for, say, microfinance, the interest rate would be higher than, say, the--

Imagine you have a simple model where microfinance institutions are raising capital on Wall Street and lending it to small lenders. This already tells you there should be reasons why the interest rate is paid by small borrowers, or it tells you why the interest rate paid by the borrower should be a lot higher than the level that the bank is charging on interest rates. But it comes from the default channel. All right. So those models predict high interest rates, but the reason for the high interest rate is to cover the expected defaults.

So one unfortunate fact for this model is that, at least in microfinance, at least in some cases, interest rates are high. But often repayment rates are actually high too. OK. So then that doesn't work in this model because the only reason in both those models that the rates were high was to cover the defaults. So to the extent this is also a force, if there are reasons for that wedge beyond the default rates, we'd add something else to the model.

One simple way to do it, by the way, is if you thought that the lenders were not competitive, then they're just charging markups. It's not in the slide, but if the lenders are not competitive, then, of course, they're going to jack up the interest rate and make more profits. So fine. That could be part of it.

And by the way, there are reasons to think that there may be noncompetition in lending markets. In particular, this whole question that we were talking about getting to know people's p 's and monitoring them, there's an important rule for information. So in particular, imagine that I'm Tishana's lender, and she and I have been working together for years. I may have information about her p from my years of interaction with her that, if Rabajit chooses another lender, Rabajit doesn't have.

So that may mean I have some advantage of lending to Tishana over Rabajit, and I'd be able to use that and extract some rents from that. So in particular, there may be good reasons to think that the model of perfect competition is not necessarily the only model. It's not necessarily a perfect model for this. That's one.

But here's another model that also could be relevant, which is monitoring costs. And in particular, this monitoring cost idea captures the-- this actually is not going to quite capture it the way it's going to be written down here, but it captures the idea that banks are spending a lot of money doing stuff to make sure people repay. OK.

So one version of monitoring is, we have to spend a lot of money upfront figuring out who the good p 's and who are the bad p 's. That's one version of it. A second version is the moral hazard version, I think, that we were talking about earlier, which is to say, look, I'm going to meet with you every single month and make sure you're doing what you said you were going to do and not deviating and doing a worse project.

But a third one is, maybe I can actually just spend some effort to get you to repay. So that's what this model is about, so let's think about monitoring here. So let's abstract from the project choices for the moment, just to make-- and the project's failing. But instead, we're going to think about repayment.

So as before in the previous model, suppose that if I invest k dollars, I get a gross return, f of k , the gross interest rate, r . And so if I have wealth, w , and I want to invest k , I need to borrow k minus w . And the amount I'm supposed to pay is k minus w times r . That's the same as before. OK.

But now, suppose that it's not the good state or the state anymore. Suppose that basically I can just choose to stiff you and not repay by paying a cost η . Right, η ? Yes. We don't have an η variant yet, do we?

STUDENT: [INAUDIBLE].

PROFESSOR: Huh?

STUDENT: It's already [INAUDIBLE] η , so presumably it just wasn't--

PROFESSOR: It wasn't a big one.

STUDENT: --a big deal.

PROFESSOR: Yeah, they skipped a couple recently. They skipped η apparently because everyone thought it was going to be called the "new variant," and they thought that would be confusing. So anyway, η , OK. By cost, η , that's proportional to the amount invested.

So that's what's going to happen? So lenders are going to provide financing up to the point where the borrower is going to repay. So this is a slightly different model, but in this case in particular, they'd like to get paid back. What's the borrower's choice? Well, the borrower could pay back and get f of k . And the borrower's net return if they pay back is f of k minus the amount they have to pay. Or the borrower could pay this default cost and get f of k minus η of k . And you can think of this as like, they come and seize my factory and take some share of it. So it's a slightly different model.

So if we rearrange terms here, the kind of key quality condition is that $r k$ minus w has to equal ηk , or that k over w is equal to r over r minus η . OK. So in this model, this is actually going to give a maximum amount the bank is going to want to lend you. With amount of borrowing, it's going to be increasing in w . The more w you have, the more you can borrow, and decreasing in r . OK.

So now, let's suppose there's a monitoring technology. So the lender needs to spend some resources in order to make the borrower want to repay. Basically, in other words, η is equal to zero if the borrower doesn't spend any money. But if the borrower does some stuff, they can get you to repay. Maybe that's like finding out where you live. If you think about a lot of the microfinance institutions we'll talk about next time, they might be clustered together in groups of people. And the basic idea is, if you defaulted, other people in your group can impose social sanctions on you if you to default.

There's some cost we have to create to get you to repay. And in this model, they're going to say, suppose those are a fixed cost. It's basically a cost per borrower. And there are some estimates from this old paper by Aleem, which says that maybe they're pretty high, particularly for very small loans. And this cost would be a reason why basically microfinance has really high interest rates, because if I have to find out, yes--

Fixed cost may be a little bit much, like it's probably cheaper to find out if Joe's dry cleaner is a better investment. Or Joe's tailor is probably easier to figure out if they're worth lending to than the Gap or something. It's probably not totally fixed. But the idea that it has this large fixed component, or it's definitely not increasing very much, seems kind of reasonable. And if that's true, that's going to give us a reasoning why microfinance is making lots and lots of small loans. It's going to have to have really high interest rates because they have to cover this fixed monitoring cost.

So basically, if you're making a small loan, you have this fixed cost. You have to recoup it. In terms of interest rate, it's going to be a large percent of a small base, whereas if it's a really big loan, that fixed cost is going to be much smaller. That's basically what's going on in this model. OK.

So how is that going to work? So if it's a fixed cost, then the lender's going to just break even, so their return has to equal the market return plus the fixed cost. And so that's kind of it basically. If the amount they're borrowing is really small, then we only have to raise the interest rate a little bit to cover this fixed cost. If that borrowing is really big-- sorry, if the borrowing is really small, if you're raising interest rate a lot to cover it, if the fixed cost of the amount is really small, we don't. OK.

So if we had this condition from before of what the optimal level of the credit constraint is, we can just substitute that in and get the return. It's going to be equal to the market return plus this markup term, which is going to be a function of their wealth. That's just combining these two points, and they're basically for really-- and so what's the point-- for really rich borrowers.

For really rich borrowers, they can borrow a lot, and therefore, we can spread that fixed cost over a small amount. For really poor borrowers, that credit isn't really binding. We're not going to be willing to lend them very much, and therefore, that interest rate we're going to charge them is going to be a lot. And then you just solve. And then, of course, that has this additional thing, which is the interest is in an equilibrium condition, so you have to do a little algebra to solve that out. But that's the basic logic of what's going on.

So in this model, some people will never be able to borrow. And this condition over here, if η is greater than η^* , then there's no equilibrium. Either there's no solution where you can cover the monitoring cost with the interest rate because the industry gets too high, and then they won't repay and so on and so forth. There's no solution to that. If it's less than that, then there is a solution, and then you have this condition I said before that r is going down with wealth going up.

So this is a model where we have people who are going to be actually credit constrained. Some people just cannot borrow at all. It's not this, you can borrow a little bit. Here, some people are just totally cut out of the market. And we also have this thing now where the people who are borrowing a little are the ones who are paying the high interest rates, which is, if you think back to one of the motivating factors I kind of gave you at the beginning of the talk, this is now something that explains that. And that also is consistent why we often-- consistent with finding the fact that we have all of these very small loans having very high rates.

We also get-- maybe I won't make so much of this-- but we have this multiplier property, which is basically that ρ , $d\rho$, how the rate charged depends on underlying market rate. This thing is going to be less than-- this thing is going to be a multiplier that's going to make the rate be very sensitive to the underlying rate, so it also kind of moves around a lot. But I'm not going to focus on that as much.

So what did this model add to these other things? Well, a few things. So it added a wedge between the cost of capital in the interest rate. That was given by high monitoring cost. It gave you a sensitivity, which I talked about in the last slide. It also implies that subsidizing the cost of capital can lead to welfare gains because r will go down, and this kind of allows for us to borrow more.

It also suggests that these monitoring costs are really important because these monitoring costs are what's excluding people from credit. And this question of, how do we think about what kinds of monitoring costs are going to be really effective is going to be something we're going to talk about a lot more next time in the microcredit lectures. OK. Yes?

STUDENT: Two slides ago--

PROFESSOR: This slide?

STUDENT: Before. What is cost for?

PROFESSOR: Oh, these lenders are not wildly profitable. That's all it means. So we're seeing these big spreads, but it's not like the microfinance organizations are making tons and tons of money. So that was rejected. Like I said, the other thing is, maybe they're not competitive in certain events. This bullet point says, well, maybe it's not that they're totally making lots of rounds.

I think it could be part of it, but I think in general these organizations are not wildly profitable. A lot of them actually are losing money, or certainly kind of breaking even, which suggests that it's not-- although I do think the rent thing could be part of it, it's not mostly about rents, and that's why we kind of need these real costs. OK.

So the final or penultimate thing I want to talk about today is this paper by-- I want to talk about the theory. So basically, we had three models that you should keep in mind-- moral hazard model, adverse selection model, and monitoring cost model. So those are three, I would say, frameworks or tools that you should keep in mind as you think through some of these credit issues.

So this Karlan and Zinman paper was basically trying to understand, how do we think about, can we differentiate between the first two models, adverse selection models and moral hazard models. And prior to this paper, there was a bunch of previous papers who said, look, can we test for asymmetric information by looking for, say, a positive correlation between the interest rate and the default rate risk. That was kind of that was coming out before, but they didn't really say, was it adverse selection or hazard because both those models kind of produced the same thing.

So they will say, well, can we actually do something to differentiate between those two. So again, adverse selection is the idea that some people are riskier than others, and it's a characteristic about you. And moral hazard is, people can choose how much risk they want to take, and that's an ex-post choice. And so these are going to look similar in the cross-section. We basically talked about that.

But the idea of their paper is to say, well, can experimentally we do something to tease this apart. And this paper pioneered this experimental design, but it's now been used for many, many, many different kinds of things. And I think we may have already seen it already, but this is the original paper, so I wanted to go through it a little bit.

So their idea was, adverse selection is driven by the ex-ante interest rate. Moral hazard is driven by the ex-post interest rate. And in the world, these are the same. The interest rate that you get when you-- like, I advertise and say, hey, I'm offering a loan at whatever, 10%. If you sign up for that loan, the rate is 10%. So those things in general, in the world, are the same.

But maybe in an experiment, we can differentiate between them, and we can do that by surprising people with an interest rate, with a change in the interest rate after they've accepted the loan. So you advertise one rate, you get take up on one thing. You change it ex-post, and you get something else. And in particular, you can't, I think, make people worse off. I can't advertise 10% and surprise you to 20% because, number one, people would be really annoyed at you, and number two, you still get a selection effect because people would drop out again.

But on the other hand, almost everyone is going to be happy with an improvement in their terms. So if I advertise a rate of 20%, and then surprise some people with a rate of 10%, there's going to be essentially no selection difference because everyone who wanted the thing at 20% will be happy to take it at 10%. So you can hold your selection constant but get a different kind of interest rate ex-post. So that's basically their idea. They have this other dynamic incentive thing, but that's, I think, neither here nor there.

The basic idea is this. We just say, look, we offer people either a high rate or a low rate ex-ante. If you sign up for the loan with the low rate, you get the low rate. If you sign up for the loan with a high rate, with some probability, we surprise you and give you the low rates. So now, we have three boxes, and if we compare the high offers to the high offer-- and this middle box is the key box.

Because if we compare the high offer, if we hold fixed the high offer rate and look at the actual contract rate, this box to this box tells us the moral hazard effect, because you both selected in at, say, 20%, but only some of you experienced 20%, and some of you experience 10%. And if we compare this box to this box, we get the pure selection effect because both of you guys experienced 10% ex-post, but here, you were selected in as having a 20% rate, and here, you're selected as having a 10% rate.

And this design is useful for distinguishing between selection and moral hazard not just in credit, but kind of everywhere basically. And so people have done this thing in all other settings, trying to use this surprise improvement to differentiate between selection and moral hazard. It's a very general purpose experimental design. It's also, I think, a nice example of how an experiment can let you do something that you could not do in the real world. It's only because you have this ability to manipulate and ex-post change people's-- I mean, you could find a natural experiment to mimic this, but in general, it'll be hard to do.

OK, is this clear? Yeah?

STUDENT: One quick question. In the design that they're doing, the information of the adverse selection, is there some bundle of the selection type, but also potentially selection on moral hazard basically? Like, you wouldn't be able to distinguish that just from this, but you can--

PROFESSOR: Sorry, say again what you mean, Aaron.

STUDENT: So with adverse selection, there's both I'm selecting on my p, but also I know how I would respond. There's also, I guess, this notion that you can select on how you would respond to different incentives.

PROFESSOR: Oh, I think I understand. So you're saying, in a more general model, in addition to selecting on just having fixed p 's or totally endogenous p 's, I have another thing, which is, some people have different p of r s. Is that what you're saying?

STUDENT: Right. Exactly. Yeah, so it's not anything we've discussed previously, but there's this other dimension that you could model that isn't captured by the comparisons that we're doing here.

PROFESSOR: So OK, let's think that through together. So what's going to be picked up here? So you're saying, we have three dimensions. We have heterogeneity in p ; ex-ante, your choice of p ; and then you have this other thing where people may differ in their p of r function.

STUDENT: Right.

PROFESSOR: So how does that affect your interpretation here?

STUDENT: I guess my interpretation is that the hidden information effect that they're picking up is some composite of the two different types of selection, that it's both selection on types and selection on slips basically.

PROFESSOR: Hold on. Let me think for a minute. So it's clearly the case-- yes, I think that's right. So it's clearly the case, if we compare these people who had the same selection, some of them got the lower interest rate, that's a pure moral hazard effect. And if they compare this versus this, then-- yeah, you're right. I mean, yes, we don't know what-- they're experiencing the same p ex-post also.

So what we're getting is like, in my example, we're both lowered to 10%. We're getting group number one experience at 10% and group number two experience at 10%. So you're right, we don't know what was the slope of their selection thing. But we are at least evaluating it at the same point, so it's still clearly picking up, I think, an adverse selection point. But yes, exactly, whether the adverse selection is a level or a slope, we are at least evaluating at the same point.

One thing I'll say is that actually if you read the details of this paper, it's not the perfect setting to map into the model. It's consumer credit, and so it's a little harder to think about. I mean, of course, we can think about the same issues in consumer credit as in firm credit, but I don't think it necessarily maps to that model that we had in mind as nicely as maybe it could.

But I think this paper was really innovative for creating this kind of general purpose. I mean, there may be another paper like it. This is the one I certainly know of that did it first. There may be others beforehand, but it certainly-- this design that they use here, I think, became very popular, and I think people have used it in a wide variety of settings.

So in the end, the results are not super duper strong. They find a little bit of evidence of the contract rate, not a huge amount on the offer rate. But I want to focus on this paper less for the particular results they have, but more for that design of how you might think about distinguishing between adverse selection and moral hazard.

OK, so I'll stop here. And then I will talk about microfinance, including the reading you guys did on Wednesday. OK, thanks.