

[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT
TOWNSEND: Thanks, everyone, for coming today. So first thing I wanted to do is just review where we are, as usual, in terms of the calendar. So we did the first part of contracts and mechanism design as kind of an introduction last time. And now we're going to do the second, which is very explicit about another application beyond medieval villages. And we'll definitely have real, live data and so on.

And then by Thursday, we will get into general Equilibrium, again, Walraisian worlds, and begin to look at trade applications both in emerging markets and then later in the US. I'll say a little bit about the reading list, especially since you might miss this. The paper application that we're going to feature today is this paper with my coauthors on five villages. And we're going to both stress the application using the tools you've already learned in a different context but also some of the techniques.

Remember, the class is trying to balance both learning tools and how to use them as well as the economics so please don't miss this Karaivanov starred article, including moral hazard programs with lotteries because I will say a lot about those techniques today but not enough. So there's two starred articles.

Again, I have earmarked a few of these from what we did last time. For mechanism design, we did a single period. We worked through the so-called revelation principle. And finally, we got to a two-period problem. And the review question is written here is, now maintaining the lotteries, write down the programming problem for the two periods, dates one and two.

I'm not asking you to literally write it down. But could someone volunteer to tell me about the truth-telling constraints at the final date two, in words, what they look like? And then also in words, what the true telling constraint of date one looks like.

AUDIENCE: I remember the truth-telling constraint at the final entry is that no matter what you already announce at the date 1, based on the [INAUDIBLE] how would you announce at the date 1, you want to choose-- The truth-telling bring you the maximized utility. [INAUDIBLE] function is like the utility at the date two is the [INAUDIBLE] and which modify the probability function. The probability function was depending on what you announce in the [INAUDIBLE] one. And then the whole function would be bigger than any announcement at the [INAUDIBLE].

ROBERT
TOWNSEND: OK, and what about at date one?

AUDIENCE: The date one is that you do not only care about what you will get at this moment. But you also care about the future. I mean, that for date one, the maximize the utility. Not maximize, the you're telling the truth that bring you the utility, which is larger than any date 1.

ROBERT
TOWNSEND: Yeah, and what is it assuming about date two, even though it's a date one constraint?

AUDIENCE: Assuming about the date two is that it-- The date two not care about what-- Yeah. They do not care.

ROBERT TOWNSEND: Yeah, you're very close. So the point of working backwards, at date two, we get the truth-telling constraint. So the utility is greater for telling the truth and announcing something counterfactual. That's now imposed as part of the solution as a constraint.

So working back to date one, it assumes truth telling at date two. So we only have to worry about whether to be truthful at date one or lie at date one. Because whatever they do, truth or lie at date one, they will tell the truth that date two.

So related to all of this is the message, the announcement of date one. You said something not quite right. But I may have heard it wrong.

There's a probability transition function that has to do with a probability of states at date two, conditioned on dates at state one. And that never changes. That's just part of the environment.

What changes, until the one that you mentioned, is the message at date one, depending on whether or not the total the truth. So the very last part of this review question asks to describe the way in which past history matters or contemporary outcomes.

AUDIENCE: So the key is that the history matters, the probability function, right? Because if you're telling the-- Whatever you tell at the date one, that will decide the probability because the probability was conditioning on what are you are telling, what do you announce.

ROBERT TOWNSEND: So I don't think I have that lecture open anymore. We'd have to go back and look at it. No, it's buried back deeper in them. Go back and look at the Markov Process. The probability P of θ_2 , condition on θ_1 , there's never-- That part never changes.

AUDIENCE: Oh, sorry. I'm mixing it. I should say the lottery. I always say the probability. Yeah.

ROBERT TOWNSEND: OK, OK, yeah, lottery. Yeah, they're both probabilities. I understand. OK, good. Thank you. That's perfect.

OK, and then relatedly, there's this question right afterwards, what does the solution to that two-period problem look like? Is it full insurance? Is it borrowing and lending? And if it's not either, one what is the logic of why we know it can't be either? So let me take--

AUDIENCE: Let me take a stab at this. But yeah, so I remember that we said that the full insurance wasn't optimal because if you had full insurance, then it wouldn't be truth-telling. And so you would never be able to reach that optimal equilibrium there. For the borrowing and lending, I remember we had some kind of argument where we said that the incentive constraints won't bind.

And so because those didn't bind, it's equivalent to just not having the incentive constraints at all. Which means you might as well just do full risk sharing. But we already showed that we could do that. But I don't think I quite recall why the incentive constraints don't bind in the first place.

ROBERT TOWNSEND: Well, it's a perfect answer. Thank you. And to answer your question, the presumption is somehow, if your income is low, you want to borrow.

And you strictly prefer that to lending. So that strict preference is a strict inequality in the constraint and hence, not binding. And likewise, if your income is high, you strictly prefer to lend rather than borrow, OK?

AUDIENCE: That makes sense. Thanks.

ROBERT Yep. Your answer was really good. Thank you. OK, so let's go to maybe one more thing here.

TOWNSEND:

I did feel that I was rushing a bit toward the end of the lecture. So I'm happy to take questions rather than going through these. Let's try a couple of them to get the conversation going.

I drew a distinction between encryption, hashing, and cryptographic puzzles. Does anyone remember?

AUDIENCE: Yeah, I think I know this one. So encryption is where you have some specific key secret key shared between two parties. And then you can use that key to encrypt a message and then decrypt it at the other end.

Hashing is some function that it changes pretty randomly, depending on the input. So it's distorts things that encapsulates the whole message. And if you change a message, it's very hard to change a message to keep the same hash.

A cryptographic puzzle, I think, it's something that is hard to do. And it's known to be hard in some way. And I think homomorphic encryption, I think that's the one where you have basically what amounts to a linear function.

So you just encrypt something. Add together all the encrypted things, and then you can decrypt the sum. And then you can get the correct decrypted sum without having to know any of the individual components, or something like that.

ROBERT Very good. Excellent. Just a couple words of comment. The encryption part, it's private and public keys. And the
TOWNSEND: private keys are never shared. It's the public key that's shared between the parties.

AUDIENCE: I just described symmetric encryption. Yeah, it's also private, public keys.

ROBERT OK, you have the background for this. OK, excellent. Yeah, and hashing is what you said. It's kind of a one-way
TOWNSEND: function. You know what goes in. But you cannot interpret what goes out.

And the puzzles are fine. And homomorphic encryption is fine. So thank you, that's great.

Sounds like you might have known the answers to these things before the class. But I mean, from other work that you're doing. But that's fine.

Let's see, one last thing. Consensus Protocol. Is a consensus protocol necessarily incentive compatible? So this is like begging for a discussion of how computer science use the word trust and economics use the word trust. And again, I'm just taking volunteers today.

So again, lots of information on those slides just kind of appeared on one slide and I didn't belabor it. The basic idea in computer science is that most nodes are trustworthy. They're going to follow the protocol.

But a handful of nodes or actors are either faulty, in the sense of the computer doesn't work well, or nefarious in the sense that they're determined to undercut the system. So the basic premise is, most nodes are honest in the sense that they're going to follow the algorithm. In economics, it's a dismal science. We don't trust anybody.

So everyone, on the other hand, can be given incentives. So we assume maximizing utility or profit-maximizing behavior. And then we load things up in a way that even though there are secrets, like preference shocks or unobserved states and so on, the actors have an incentive to reveal what those states are. So the message is, as in mechanism design that we did earlier in the lecture, are part of the contract.

And they're going to be sent honestly. But it's due to the incentive constraints. And without the incentive constraints, we kind of assume everyone's going to cheat.

So it is kind of a philosophical difference. Anyway, that's why, in the lecture I put the discussion of bitcoin and other algorithms in the context of mechanism design to draw this contrast. OK, so that's the review sheet. And let's get to don't lecture here.

OK, so today part two, more on contracts. So the application is going to be financial markets. And we're going to try to figure out what kind of contracts a small business would write with a lender. In particular, what the underlying obstacles to trade are.

We're going to want to identify, from data using the theory, what those obstacles are that limit the ability to go into business or limit the ability to finance the business once you're already in business. So that's the economics part of it. The technique part of it is, we're going to do a structural estimation of the model. But we will also look at some reduced form techniques that are data summaries that are consistent with the structural model.

Essentially, we use reduced-form techniques to identify stylized facts different across different regions. And then we use the theory to try to make sense of what we're seeing. It's kind of an experimental laboratory.

And there are going to be two sources of constraints. So one is for short, a collateral constraint. Agents can borrow, but potentially not much. In order to borrow, they have to post collateral so that the amount they can borrow is limited by their wealth.

That means low-wealth households may not borrow very much. And it may not make sense to go into a business at all because they can only operate at a very small scale. Or if they do go into business, because it's the best alternative notwithstanding, they are not investing very much because of that collateral constraint.

Now when there's a collateral constraint that's binding, the borrowing is a proportion of wealth. So if you're thinking about going across different firms in the cross section, as net worth increases, borrowing would increase, again, for constrained entrepreneurs. On the other hand, it doesn't necessarily mean that with this collateral constraint everyone's going to be borrowing. They may have enough wealth to entirely self-finance.

Or the borrowing constraint isn't binding in the sense that they don't go up to the limit. They have some wealth. They borrow additional resources. But they could have borrowed more and they don't do it.

On the other hand, a different constraint is a moral hazard. Now moral hazard means that, it's a funny terminology. It's like you buy an insurance policy that's supposed to pay off if your house burns down. But because you're insured, then you don't pay attention and take the appropriate precautionary measures.

So the house is likely to burn down. For some reason, in the insurance industry, this is referred to as moral hazard as if it were, quote, "immoral." But again, from our dismal science, economics point of view. It's not immoral, per se. It's just that it didn't have the right incentives.

And in this moral hazard world, effort, diligence, taking care is not observed. So the only way the insurance company or the lender can get the money back is when there's not an adverse event or the firm is successful in making profit. Now that means that poor borrowers who owe a lot of money don't have an incentive to work very hard at being diligent because they're going to pay the bulk of their returns back if they're successful to the bank or to the insurance company. So they don't take care. And the probability of failure is high. It's a constraint.

Because the probability of failure could be high, especially for low-wealth agents, the banks, in order to break even on those loans, have to charge high interest rates. And in fact, the rate could be so high, it would be high enough for the lender to break even but the borrower is no longer interested because it's just too high of a rate. And likewise, high rates, they may not want to borrow very much, even if they do go into business.

Another difference with the collateral constraint is in the moral hazard world, the incentive constraint is always binding. That's a version of what we were just discussing, full insurance is never optimal because under full insurance, you exert minimal effort. So you don't offer full insurance. But likewise, then effort it is lower than it would have been in a full information world. So the incentive constraint, to be diligent, is always binding.

And the way out of that problem, if in a cross section, entrepreneurs have more and more wealth comparing one to the other, then they can self-finance more. And the more you self-finance, the less you have to pay back. So in this case, as wealth goes up, borrowing goes down.

That's exactly the opposite from what we got in the limited liability case for constrained entrepreneurs. When wealth went up, borrowing went up. So I alluded to these reduced-form statistics. We will see again at some point a summary of the data that, depending on the region that they're in, we will see either this one result in the cross section or the other one.

And the goal then is to take the full-blown model and do structural estimates to determine which constraint is binding. It could be one. It could be the other. Actually, technically, it could be both.

And I'm going to show you that once we review the tools we already have, pending on the appropriate constraints, it's actually easy to test one versus the other versus multiple. It's relatively easy to do. But we are learning some of those new tools today. That's half of the point of the lecture.

So I've made reference twice to differences across regions. I think I may have shown you this picture before. But I'm not sure. We certainly looked at the Thai data.

We looked at the monthly data when we did risk and return in village Thailand, which was the second application of the risk-sharing sections. So anyway, Bangkok's over here. This is central region. These yellow provinces were ones that I had surveyed in annual data since 1997.

And over here in the East, but they call it the Northeast, we have two other provinces. These are predominantly rural, agricultural. The ones near Bangkok are much more industrialized.

And here's a map of Sisaket, just to show you where the data are coming from. We've got a lot of little villages here, these little dots, as well as parts of the road network. And then the colored circles correspond to, if they're purple, these clusters are where we gather the annual data that we're going to be using today.

In the previous lectures, Risk and Return in Village Thailand, we were using the monthly data. And those are stars. And if you look really closely, you'll eventually find a cluster with the stars that correspond to the monthly data.

The annual data are spread out much more. Those stars are up here. They're just not visible. I know these provinces like the back of my hand because I've been there so often.

OK, so again, Thai data, 1997 onwards, almost 3,000 households. They're running their businesses. As you know, livestock, fish, shrimp, farm, and small business like vendors. And we have those two regions of the country.

The cool thing is that in 1997, we had the foresight to ask about some of the things that had happened in the last five years. So we have their wealth retrospectively, what it was five years ago before 1997, and also what kind of occupation they were engaged in five years before 1997. So that those are two of the key variables, namely, how much wealth did you have in 1992? And what business were you in 1992?

And then, did you change your occupation, say, going from being a wage earner to going into business over those five year periods? And again, we're going to find the bottom line is the data on wealth and occupation transitions as well as the cross-sectional data is going to favor moral hazard as being the constraint in the wealthier central region. And in the Northeast, it's going to be that limited liability, although frequently, both limited liability and moral hazard in combination.

So hopefully, you're wondering how on Earth we could ever figure this all out. Here, by the way, is the picture of increasing wealth and increasing probability of making a transition into business. So this line here is upward sloping, the higher the wealth, the higher the transition.

But this could, in theory, be a very steep profile or a very flat one. So part of it is estimating the profile suggested by the theory and then comparing it to the actual data. Here's the model.

So agents, all these ingredients should be very familiar by now. You've seen them at least twice and maybe more. Households that could run a firm or could just be a wage earner cares about the utility of consumption and effort.

So the consumption part is constant relative risk aversion. And the degree of risk aversion is gamma when gamma is zero. This thing is linear. That's risk neutral so not risk averse at all.

As gamma gets higher and higher, you have more and more curvature. And then entering inseparably, linearly is this utility of effort, Z being effort, taking care, due diligence, how hard you're working, again, raised to a power like constant relative risk aversion to a different degree. Gamma 2, rather than gamma 1 and kappa, reflects the utility trade-off between consumption and effort.

So not to say you need to memorize this. But we can pay attention to the parameters. We got three of them here, gamma 1, gamma 2, and kappa. So we would like to estimate those.

And we will start looking at the model when gamma 1 is zero, hence, risk neutral, and then let them be risk averse and estimate the model generally. The estimation will use linear programming. You've seen this trade-off between consumption and effort.

When we did the lecture on utility and preferences, there was a slide at the very end that had the agent choosing between work and wheat as an output. And we actually did that in the right-hand side with respect to consumption and leisure or the left-hand side with respect to consumption and effort. So this is a parametric version of those same preferences.

More on the model, households differ with respect to initial wealth. As I said, call that A for assets. They differ with respect to talent, θ . Some people are better things than other people like running a business.

And S as is the level of schooling, formal education. All are observed by the agents and even by the banks. That's a bit of a stretch. But we are imagining the bank can see wealth and talent and education.

We however, only see what we have in the data, which is wealth and schooling. We do not see talent as modelers. But we can make some assumption about it, namely, it's log linear. Talent log θ is log linear in assets and schooling.

And we assume that error term is normally distributed with a mean 0 and variance of 1. Why are we doing this? Well, we want to distinguish the constraint. But the worry is that talented people in the past have been more successful and hence, already have higher assets.

We want to be able to distinguish the set part, which is wealth, which plays a role in the financing constraints, versus talent, which is productivity, which is entry in the production function. If we didn't allow talent to depend on wealth, then we could incorrectly infer that financing constraints are playing a role because higher wealth allows people to do better. But actually, it could be playing a role through this third variable, namely talent, that higher wealth means higher talent, reflects higher talent. And it's actually higher talent that is the reason some people are in business and operating at a higher scale.

So this is kind of a shorthand way to allow things to be correlated and not incorrectly infer something about causality. Other parts of the model, q is output, profits. It's a function of effort and capital if you're a firm. Or wage-earners, if you're not a firm, this output can take on two values for simplicity, namely success and failure.

So when they're successful in the project, the value of q is θ . So more talented entrepreneurs have higher output conditioned on being successful. And the other simplification is the failure, which means they have zero output. And if you have zero output, then you're not going to repay a loan. And there's no output with which to repay the loan.

So this whole branch of q equals 0 becomes simplistic. Anyway, yeah, two extremes here, three really. Talent only enters through success. There's only two values of profitability, success or failure. And failure corresponds with zero output.

So let's look then at the probability of success. The probability of q is equal to θ , the high one, conditioned on effort and the level of capitalization of the firm. And that's essentially in the numerator like a Cobb-Douglas function. K is capital. z is effort to the power is α and $1 - \alpha$, a function you've seen before.

Now why is it being divided by something? Well, this whole thing is supposed to be a probability. Probability of success is a number between 0 and 1.

So if k and z were 0, the numerator is 0. 0 divided by 1 is 0. So that's the left endpoint.

And likewise, if k and z are getting really, really big and there's no natural bound here yet, then this numerator is getting really big. But so is the same term in the denominator. So this whole term is going to 1 because the numerator and denominator are getting very, very similar.

So this P function is a probability between 0 and 1, no matter what k and z are put into the firm. Output is always going to be observed in the models that follow. The issue is going to be whether this effort z is observed or not and whether k is restricted by the collateral constraint.

And I will show you that in a minute. And finally, the other branch, not to lose sight of it, they could choose not to be a firm, in which case k is equal to 0. And when k is equal to 0, the firm is not capitalized. And effectively, the household is a wage earner.

So earnings in the wage sector are also stochastic and probabilistic. You can work hard to find a job. That doesn't mean you will find one. The higher effort you devote toward finding a job, the higher would be the probability of getting the wage. And again, this is bounded between 0 and 1.

Now if you are a firm, you may or may not borrow. But if you borrow, there's going to be an interest rate schedule that determines how much you pay back. And that interest rate schedule is a function of two of the things the lending firm observes, the wealth of the household and the household's talent.

And this schedule is going to be such that the banks, although maximizing profits, can at best break even. And I'll show you that constraint in a minute. But I already alluded to it in the sense of moral hazard being a problem because repayment going down and banks having to raise interest rates to cover what otherwise would be losses. They have to get enough from the successful firms to cover the fact that the unsuccessful ones are not paying off the loans.

When you're a worker or a firm that is not using all their resources to borrow, you have some residual assets left. They put it in the bank. And they earn the riskless rate of return. So in some respects, there's both borrowing and saving going on here, depending on what occupation you choose and your underlying characteristics.

So how are we going to solve this problem? We're going to solve for a constrained optimum. We are going to write down a quote, unquote "planner's problem," which I sometimes refer to the math problem, which is maximizing lambda weighted sums of utilities subject to resource, truth-telling, and liability constraints.

Now to restate that, we're going to maximize the utility, a weighted sum of the utility of borrowers and the utility of the banks. The banks, however, could make profits. So we're just distinguishing, in their case, profits as a version of utility.

The second thing is, and I didn't reload this picture, if you go back to the lecture where we introduced the notion of Pareto optimality. We had the so-called Pareto frontier. And we decided the points on the frontier would be Pareto optimal.

And you could achieve those points by maximizing a particular lambda-weighted sum of utilities, or in this case, a particular lambda-weighted sum of profits and utility of the borrower. But there is another way to do it, which is to maximize the utility of the borrower, subject to a fixed profit level of the bank. Its equivalent, normally, with these concave frontiers.

Then the next step is, OK, there are many, many banks. And they're competing. If one representative bank were to make a profit with respect to the borrowing customers, another bank could step in and undercut the first one and do slightly better.

So this competitive process is going to drive the profits of the bank down to 0. So we're going to be looking at particular Pareto optimum where the utilities of the borrowers are maximized, subject to a zero profit constraint on the part of the lenders. There are other optima we're not going to look at today. Although, it's going to be clear from the equations how we could have done it more generally.

So this is the utility of the borrower. There are three possible branches. This is, again, the risk-neutral borrower.

To get warmed up, if k is 0, which is a choice. But if the solution is k equal to 0, then, in fact, the household is a wage-earner, not a firm, deciding on effort, which determines the probability of getting the wage, although it comes at a disutility cost of working. And they put all their wealth in the bank. So that risk free rate of interest times assets is something they get at the end of the period, the savings account.

The second branch is they do go into business and there's capitalization is positive, but it's less than their assets, so they can basically take what's left, assets minus capitalization, and put it in the bank. So these guys also get a return on savings. But part of their wealth, the k part, is stuck into the production function. So the return on that comes from the possibility of profits. And you still have this trade-off of probability of success versus the disutility of effort.

And then the final branch here is, these guys have a k , which is relatively large relative to their assets. So they got to borrow. And this sign here is a bit confusing. But if k is greater than a then a is less than k . So this term is negative.

So naturally, they're quote, "negative" paying off the principal and the interest on the loan. And they only pay it off if they're successful. So again, this modified Cobb-Douglas probability expression is entering in there as it does for the probability of success.

So those are the three branches. This is the break-even constraint of the bank, which as I've said in words a couple of times already, this is the risk-free rate. So a bank has basically got to get funds from somewhere, from depositors, actually. And they've got to break even because of all that competition.

So they may get a higher return, depending on the borrower's wealth and talent. But they only get that back when that borrower is successful. So this is the expected return given the interest rate schedule. And that is driven down to the risk-free rate. So we have this 0 profit constraint to impose as part of the solution.

So again, to repeat myself, instead of maximizing a lambda-weighted sum of utilities, we're going to maximize the utility of the borrower subject to a particular lambda for the bank, a lambda, which is associated with making zero profits. And we just impose the zero-profit constraint from the get go. It's also kind of interesting in terms of where we are in the class because when we first introduced utility maximization, we had these partial equilibrium exercises in which prices and so on were taken as given, although we varied them to see what the household will do.

Likewise, when we did the firm, we had input prices as taken as given. And we did some experiments. And here, likewise, we're in quote, unquote "partial equilibrium" in the sense that wage and the interest rate are given to us.

We're not trying to explain them. They're going to be part of the data. On the other hand, we're definitely doing a contracting a problem here. And to do it, we have to draw on those tools about Pareto optimality and so on and how to determine Pareto optimal outcomes via programming problems. So we kind of are drawing on to different segments of the class that we've covered already in this application.

So let's talk about limited liability. So limited liability, let's just jump down here. It's easier here. k , that they can use the capitalization of the firm, is bounded above by their assets. Actually, it's λ , some proportionality times assets.

So this is referred to as a borrowing constraint. Now λ could be greater than 1. They may be able to borrow more than their wealth. But it is limited.

There's a little bit of a derivation here. I was thinking about it this morning. It's not really all that compelling.

You can derive this constraint by saying there's a maximum amount they can invest in the firm. Then you write down the equation for borrowing, which is the difference between capital and assets. And then substitute in that maximum amount for capital and derive this equation.

This now has a λ on both sides of the equation, which you could drop. And you get the equation 6 as a result. But I'm perfectly happy if you wanted to start with 6. The capital is bounded from above due to a collateral constraint related to wealth.

That's the limited liability version. Here's the moral hazard version. In this case, the bank does not see the borrower's effort.

So whatever the contract is in terms of that repayment schedule, and whatever the level of capitalization that the borrower can manage, the choice of how much to work is done by the borrower alone and cannot be dictated. And it's going to be summarized by a first-order condition, which is like a solution to the subproblem. Which subproblem? This subproblem here.

So this is the branch where the household as firm is borrowing. And its payoff in curly brackets is, in part, a function of z . So if you differentiate this function with respect to effort, you're going to get the way that the probability of success is increasing with effort. You're going to get a part with a negative sign reflecting the utility of effort.

And this probability of success is also entering into this-- It's like the winner's curse. You win. Your profits are successful. But you've got to pay off the bank.

So the harder you work, the more you repay. So take the derivative of this. And you get this first-order condition over here, equation seven.

So that's the formality of it that it captures the notion, as I've said, that no matter what contract the firm is entering into, the bank cannot see the borrower's effort but knows how the borrower is determining it. So we get this equation 7 as another constraint on the program, if we believe moral hazard is strained.

So we did limited liability. We did moral hazard. We could have both. And now I'll show you a very cool picture.

So this is utility of a household as a firm. But on the x and y-axis are the inputs, the capital level and the effort level. If we did what we would call first best, meaning that we ignore all the limited liability and moral hazard constraints, then the borrower would maximize these choices of capital in effort, choose then to have maximal utility.

And where is the utility maximal utility? It's in this bull's eye over here. So in other lectures, we've talked about bliss points, meaning, if you're below bliss, utility is increasing. Or you could go above bliss, which is kind of weird, in which case you want to dump stuff.

This is an endogenous bliss point. This came from solving an underlying problem. But we can, nevertheless, plot the solution this way.

And where do they want to be? As close to the bull's eye as possible but not northeast of it. OK, so now what are the constraints on the problem?

There's this liability constraint, which is capital times lambda. And if that's the constraint in the world and hence, in the program, it just simply says, capital cannot be to the right of it. It can be to the left of it or on it.

So all these solutions on the vertical line and left of the vertical line correspond to satisfying this capitalisation constraint for a given value of lambda, which is listed down here at the bottom of the slide. And where is the maximal solution? At the tangency right here. That's where the slope of these concentric rings goes vertical.

OK, let's do the moral hazard. So ignore limited liability constraint and trace out this guy. I did not a priori what this constraint was going to look like.

But if you plot it in the space of k and z, it looks like this. It's like this parabola. Alex was my student and RA for this project. And Anna and I, former student, were writing this paper.

We shared the problem with Alex. And Alex came up with this graph. And Alex became a coauthor of the paper and subsequent collaborator in many things. Anyway, I love this graph.

So if you want to satisfy the moral hazard constraint, you're going to be on the parabola. And where is the maximal point? It's again, a tangency with a concentric circle. And it's right here.

So interestingly, if you compare these two points, the limited-liability constrained solution and the moral hazard constrained solution, limited liability involves less capital and more effort relative to the moral hazard solution which kind of makes sense because with moral hazard, the problem is inducing effort, not capital. With limited liability, the problem is financing that level of capital. and effort's fully observed.

Now it turns out, if you have both constraints, in this world with risk neutrality, there's nothing to optimize. It's like two lines cross.

Actually, technically, it's a vertical line crossing with a parabola. But it completely nails down the solution. So k and z are at the crossing over here. So what we're doing conceptually then is loading in these parameter values for, in this case, a particular talent level, particular asset level, α in the production function, κ , which is the disutility of effort. And the power is λ , $\lambda = 1$, et cetera, solving the model and delineating what happens with these various constraints, one, the other, or both.

OK, so now we're going to generalize this. And in particular, we're going to let there be some risk. There is already risk. But we're going to let the household be risk-averse. So their payoff can depend on q , which they now care about.

The dispersion, because they're risk-averse, so there's a consumption, this profit-sharing schedule. And we're going to want to solve for these probabilities. This kind of came up in the beginning of class, this π is a probability number, an endogenous number that reflects the solution to a programming problem. So we're going to want to determine the probability of seeing this quadruple, a particular value c , particular value of q , a particular value of effort z , and a particular capitalization level, as a function of observed and unobserved characteristics of the borrower, namely, θ and s , talent, assets, and schooling.

So again, I don't remember what slide it's on. But we will determine the solution by these probabilities. There are lotteries. They will make our program and turn it into a linear program.

And again, this is a bit of a review, but now in an interesting applied context. Remember, we had discrete values for consumption when we introduced the notion of risk. And we talked about the probability π , of choosing a certain bundle, discrete bundle, c , c_s , where s delineated all the possible discrete values of c .

So that was one place that we did it. And there's been others. But I don't have it right off the top of my head. Oh sorry, in the mechanism design literature in the previous lecture we did it. We had deterministic programs and programs with the lotteries.

So that's the second reason to have the lotteries. And I made the point there that they were linear programs. So despite putting in truth-telling constraints, when we have the lotteries, the constraints are linear in the π and so it was the objective function.

So here, we're going to have these moral-hazard constraints, et cetera, which could turn this into a very nonlinear problem because we have the original maximum, the objective function subject to a derivative of part of the maximized function as a constraint. And Lord knows what that looks like. But we turn it all into a linear programming problem with these lotteries.

Not only that, we're going to imagine discrete values. So it's not a semi-infinite program. It's a finite, discrete linear program. And that allows capitalization to be chunky.

Maybe you get a machine or you don't, for example, literally, a discrete value. That also causes a problem, just the way chunky consumption causes a problem. But never mind. Lotteries are the cure-all. It just always works to turn it into this.

So don't pass out. Each of these lines will make sense to you. This is the objective function. It's the utility of the household as a function of consumption and effort.

This is expected utility because part of the probability involves consumption and effort, which may potentially be chosen at random, according to π . And then we sum up over all these discrete values of c , q , z , k , summing them up, weighting them by π . So this whole thing is just expected utility. In this case, the borrower's risk-averse. So we're allowing for that.

Now this constraint is what I call a Mother-Nature constraint. What Mother Nature is giving us is the technology that maps z and k along with θ into the probability of success, q . So this \tilde{p} function is given. It's part of the underlying environment.

But you might think that, in part, we're solving for the relationship, the probabilistic relationship between q , z , and k because everything looks endogenous up here. So we need to impose another constraint to make sure we respect the Mother Nature. And if little statistics, this is like the probability of a conditioned on event b times the probability of event b is therefore the probability of the joint event a and b . And on the left-hand side, it's already the joint event, a and b .

So this is the probability of q , conditioned on z and k . This is the probability of z and k because we've integrated out the other objects, z and q . So this is the probability of q , z , k .

And likewise, the left-hand side is the probability of q , z , k , where we somehow integrate out the only other object, the c . So this is a Mother-Nature constraint. Again, it's all linear in the π s. So we just code it in to Matlab.

This thing down here is the break-even constraint for the bank. And you can think of it as maybe the bank has to pay out more in consumption than output. In this case, by the way, when the firm is not successful, the output is 0. But the borrower is risk-averse and it's going to view it as a disaster to get 0 consumption, maybe even minus infinity utility.

So typically, in the failure branch, consumption is greater than q . That means the bank, which is like taking all of the q and giving some of it back in terms of c is suffering a loss on the consumption part. On the other hand, on the right-hand side, it's like the firm, the household as a firm is surrendering all of its assets a to the bank to manage. And only some of it gets used to finance.

So the firms make a profit on the right-hand side, times the difference the bank got that it didn't give back, times the external outside interest rate. So this is a zero break-even constraint with the lotteries. This is moral hazard. It's akin to truth-telling.

So what this says is, the truth is that not only does the borrower characterize by θ a and s . The borrower contemplates doing the recommended action z . How much effort is supposed to be applied under the contract?

But you can't make the borrower do it. The borrower has to be given incentive to work, especially if this is a high level of z . But the borrower says, nah, I might not do it.

So even though z was recommended by the bank, the borrower does something else. The borrower does z' . Note that z' is entering into the utility function here. So he really is going to do that or is thinking about doing it.

And this would be the utility consequence. This stupid thing here, we have to readjust the probabilities because when we impose that Mother-Nature constraint, we embedded the probability of success into the endogenous object. We would now be making an error if we did not adjust for the fact that the effort being taken is z' rather than z .

And this is derived in one of my papers. Hopefully, it's somewhat intuitive that at least one needs to do an adjustment. And finally, this just says the probabilities all add up to 1.

So the programming problem is to maximize the expected utility of the borrower, subject to all these constraints, Mother-Nature constraints. There's a lot of them. For every q , z , and k , this is the break-even constraint. There's only one of them.

This is the moral-hazard constraint to be obedient rather than shirk. And there's one of these for every k . That could be assigned and then for every z , an alternative value z' . And there's only one of these constraints 11.

So I'm actually enumerating the set of all constraints. You can solve these things for a large number of control variables π_i , like thousands. And you can have hundreds of constraints. And it's still solvable in things like CPLEX, which is a version of which Gurobi is publicly available software.

So we did this before Gurobi was on the market. So we paid \$2,000 to get a hold of CPLEX. And we solved the linear programming problem. I'm going to call the solution the probability of being an entrepreneur.

It's a pretty brute force, namely, what is the probability the capital level is positive? If it's positive, they must be a firm, the way we composed the problem. If capital were 0, they're a wage earner.

So ignoring everything else by summing up over everything else, we only care about capital being positive. And we call that the probability of being a firm. It's still conditioned on talent, assets, and schooling. So we have to solve this problem for all those stratifications in the data.

So let me go back again. Is it a moral-hazard problem? Is it a limited-liability problem? Is it both? Is it neither?

So we have the master program here that allows everything. All we have to do is comment out the relevant constraint. If we want to look at the-- It should have been written here.

If we want to delete the moral-hazard constraint, we comment out this thing, equation 10. And there's a limited-liability constraint, which somehow I'm not seeing in the notation. But it should have been there. And we can comment it out.

We can allow both moral hazard and limited liability. Or we could have neither. Once you go to all the trouble to code this thing up, it's not hard to do all these different combinations and take it to the data.

So how do we go to the data? We compute. So again, I'm not being shy about this because we advertise that numerical methods are important at the beginning of the class. So I do want occasionally to show you how numerical methods can get used.

And also, I know you guys are a great class. So I'm not worried that this is necessarily too hard. But it may be the first time that you've seen it, which is kind of cool because you'll learn about maximum likelihood.

So what do we do? Let's fix all the parameters, alpha, kappa, all of them, a whole array of them-- I have that on the slide coming-- and solve that programming problem for the stratification of ability, education, and wealth. And say, do that for the moral-hazard regime.

Let's keep moral hazard in mind now. Drop limited liability. Now we want to construct the likelihood function, which is for given parameter values, the model predicts the probability of being a firm for households of this type.

The model tells us that. And in the data, we see the number, the fraction of the sample that are firms in that wealth, talent, schooling, category. So there may be a bit of a mismatch in the sense that for given parameter values, the probability is low in the model and very high on the data or vice versa.

So then logically, oh well, we can play around with the parameter values. We don't know what they are either. So we vary, kappa and alpha and all of those things, to make the likelihood to maximize the probability that we would see in the data what is predicted by the model.

Maximize the probability, we would see in the data, which is a given. But it's predicted by the model. And again, the choice is varying over all possible parameter values.

Now that doesn't mean the likelihood will match exactly. It just means the likelihood is as high as it can possibly be after searching over all parameter values. Here's a bit more of the math of it. The string of parameter values would be risk aversion, work aversion, the weight on disutility in the work part, alpha in the production function, delta 0, 1, 2, which are the weights on theta A, theta and S, I think, or the other way around, lambda being the limited-liability constraint.

So we're going to maximize over all these parameters and get something like H of A to be the probability of being a firm, conditioned on wealth A . And it should have had schooling in here but not talent. We had to integrate out over talent.

So this object comes from the model conditioning on the stratification in the data and a string of parameter values. So when we say maximize the likelihood, L , likelihood, is a function of the parameter values, which we will vary, to maximize what? E_i is a statement about the data. It's 0 or 1. Household i is either a firm or it isn't in the data.

So we count up the number. Every time a household is in the data as a firm, we look at the associated probability of predicting that from the model. This is a probability, likewise, logs make it the same function. It's just a monotone function of the probability.

And logs are cool because they allow you to write everything is separable instead of multiplicative. Here's the other event, a given household in the data is not a firm. It's a wage earner. So one minus 0 is 1. So this branch is wage earner.

And what's the probability of being a wage earner? It's 1 minus the probability of being a firm. So this is the likelihood function, the log likelihood function, a function of the parameters.

And this is where the code is maximized. This is not linear code. This is nonlinear optimization code. But there are many routines out there to do this.

And if you don't want to think about maximizing its nonlinear function, you can just evaluate it at thousands of these parameter configurations and pick the max, which actually gets close to the maximum. So too tiny to read, but conceptually, we have a moral hazard column, a limited liability column, and both constraints column. And what we're doing with various specifications, including this one, estimating talent, is reporting the maximized parameter values.

So again, that vector string was γ_1 , γ_2 , κ , and so on. And you can see the point estimates. For risk aversion, which was γ_1 , we're getting the point estimates here of 0.06 or 0.02. So almost risk neutral, actually.

But in parentheses is a standard error, which is expressed a degree of confidence about the point estimates. So those are the parameter estimates. So what do we find? We find at the end of the day, a couple of years of research, that the central region favors moral hazard. And likewise, in the cross section, as we increased the wealth going from one household to the next, an increase in wealth is associated with an increase in net saving, which we also measured.

So this is consistent with moral hazard in the slide I showed you at the beginning. That if it's moral hazard and you have higher wealth, you avoid the damage caused by the moral hazard constraint by self-financing. So you should be saving more as wealth goes up.

In the Northeast, the story is different. Being constrained, like below the constraint or at the constraint, is not related to borrowing particularly. Now it is true that if you were at a constraint, the liability constraint, and wealth went up, then they'd all be constrained by definition.

But there's a branch of entrepreneurs, depending on talent, like if talent is low, their wealth is more than adequate. So they may not be borrowing at all. And likewise here, we get the opposite that when wealth goes up for constrained entrepreneurs, borrowing goes up. But again, there's this branch of entrepreneurs that are not necessarily borrowing. So that wording on that bottom bullet point is a bit difficult.

So we have succeeded to solve the problem. We've identified the constraints that limit small businesses. Interestingly, the constraints are different in the different regions, in the industrialized area near Bangkok, which are very urban-- one of these villages is across the road from a Ford Motor Company plant, for example-- it looks like moral hazard's the problem.

Out in the Northeast, the problem, which is agrarian, mostly rural, the problem seems to be there's a very limited financing constraint, limited liability. And a lot of the households don't have a lot of wealth so they're either constrained or don't go into business at all. Now we're not doing a policy thing here. But it would be good for policymakers to know this because if they want to alleviate the constraints on small businesses, they should direct their policy toward maybe in the Northeast, allowing more assets to qualify as collateral, not just land but their motorcycles, which they finally did, and other things, putting it in escrow to allow or to have the bank to have a claim on it and allow to increase borrowing.

In the central, area it's an information problem. So there, if you're determined to try to solve it, it must be better monitoring. You need better signals of the underlying effort by, well, let's try costly state verification, for example, which is something I talked about last time. Not with respect to verifying output. But with respect to getting reliable signals of the underlying effort of the borrower.

OK so this is the economic science part of it. We have models. We have data. We're trying to get as close to the data as possible.

We're also solving an interesting financing, getting understanding of the financing constraint. And thinking about how policy would change, depending on what we find out rather than one size fits all. So that's all I have for today. OK, thank you very much.