

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

ROBERT

So we've had a mix of lectures in terms of focusing on specific topics and the details of models with notation.

TOWNSEND:

Sometimes that's coupled with an overview of the literature to put the paper in context. Sometimes-- and today will be an example-- we're just going to try to do one thing and do it reasonably well, although there is a lot of material today.

There are other things that were listed on the reading list-- Cynthia Kinnan's job market paper, for example. As we go through this, I'll point out what she was doing and how that compares to what we're doing today. I deliberately decided not to go through her paper on the front end of this because it just chews up more time. And there's been that macro lecture, and maybe the labor lecture was-- filled a lot of material. And I'm sort of in the mood of doing a smaller thing well today rather than trying to cover too much.

So there is a background, though, and that is constraints. We've been talking about constraints all semester. Consumption smoothing literature, including not just full insurance, but permanent income, buffer stock, those. We did the standard incomplete markets literature. There's private information stuff. We have mentioned it in class from time to time, have not done too much explicitly, although the Ben Moll paper and the others that lecture were macro models based on assumptions about information structures.

Limited commitment we've actually done quite a lot-- not just how it affects consumption, but also in those macro models. So anyway, there's a consumption literature out there. There's a bunch of investment literature out there which includes adjustment cost, sensitivity of investment to cash flow, structural modeling, and just outright reduced-form empirical papers. And I've already mentioned some of this macro literature, like incomplete markets. Some of these papers we covered in class, not all of them.

And finally we get to this, which is small but growing literature trying to test across different models. Most of the above pick one thing to do and try to do it really well. Some compare across maybe two models and some even more than two. So I could probably add two or three more References. There just isn't that much out there that systematically are agnostic about what the underlying constraint is and set out to discover the best fit against the data, and hence what the obstacles really are.

So the point is we're not going to just look at investment, not just look at consumption. Largely we will look at both together. We're not going to look at just incomplete markets or endogenous information-constrained markets. We're going to test for them, both classes within and across, and talk about how to do it and what kind of data we need.

The tools-- we're solving these dynamic models. We can allow any number of financial information regimes. We're going to use maximum likelihood to estimate the parameters, which allows us to be more general in a couple of ways. First of all, we can back out all the structural parameters, not just the parameters that are in a particular Euler equation. And likewise, we can have more than one equation, so to speak. We can look not only at consumption Euler equations, but investment rate of return equations.

So again, I think in my mind, at least, these are familiar themes that we've been covering in bits and pieces in each of the various lectures. What was Cynthia doing in her job market paper? She was looking at various financial information regimes, but focusing on the Euler equations.

And as you go from full risk sharing to limited information about output to moral hazard, constrained insurance, the form of the Euler equation varied from one to the other, including what variable should or should not show up as lags. And the basis of her task was to see whether lagged inverse marginal utility was a sufficient statistic. I'll say more about that as we go through.

So there's a long list of regimes. I think part of the point of this is that the technology is available to test almost any regime subject to computational constraints. So under incomplete markets, we get out of autarky, which is the worst, savings only as in buffer stock, maybe borrowing up to a limit. We've talked about that in the natural borrowing limit and other limits, and then a single risk-free asset, which is like permanent income, unlimited borrowing and lending.

And then we have the endogenously determined incomplete regimes, namely moral hazard, limited commitment, hidden output, unobserved investment, and the least constrained regime, the full information regime. So there's six or seven of them here. And we're going to go through tools that allow you to test one against the other pairwise for any pair depending on what data you want to use.

So there is a mechanism design contract theory part which we've been putting off until today, largely. There is dynamic programming as in value functions. We have been seeing versions of that through various other lectures. There's linear programming. I'm going to say why momentarily, although we were already starting to do that when we did Rogerson's paper on labor supply. And although it was probably hard to figure it out, that's what Victor Zorn was doing in the TA session last Thursday. And maximum likelihood, which sounds familiar, anyway.

So we compute, we estimate, and we test, basically. We can do it on actual data. I'm going to focus on the contrast between the urban data and the rural data. We can actually use it on simulated data. So I vote for this technique, which is generate the data from the model itself. Then you know for sure what's generating the data and see whether you get back what you put in in terms of the financial regime and the underlying parameters. And the results are reassuring subject to measurement error.

One comment-- that you have to be patient to get to the-- get through the next 45 minutes or so-- the criticism of maximum likelihood is it's kind of black boxy. You don't really know, other than trying to fit histograms-- it's not like you're focusing on ROA and how it varies with wealth or transitions in the capital stock, or for that matter, the time series that the model is generating. But I'll come back at the end and show you the pictures of the actual data, and you'll get-- you've seen it, actually. But I'll remind you parts of it, and you'll see why-- what it is that makes it difficult for some of these financial regimes to fit the data and hence what the obstacles seem to be out there.

We use the Thai data because we have both the consumption and the asset and income data, and using both is helpful. But that doesn't mean these techniques are limited to rather special databases. Mostly if you have surveys of firms, they would not ask about the consumption of the owner.

Fortunately, the reverse is not so-- is less constraining, which is household-level surveys done by the World Bank, Living Standards Measurement and the FLS, Family Life Cycle in Mexico and Indonesia, and so on. They do typically ask the household a lot about their enterprises. But we've done this in Spain with just data on investment. So the techniques work even when you don't have the consumption data itself.

And one of the main findings, reassuringly-- it's not like we always get the same thing back. For one thing, we don't get full risk sharing much. Sometimes we do. It will not surprise you when we get it given the other papers we've discovered in class. The big interesting thing is there is a difference between the rural data and the urban data. And the rural data, fairly limited financial regimes, like savings only or limited borrowing, they fit the best when you use the investment and income data.

But in the urban areas, even when you use the investment and income data, you get something less constraining like moral hazard. So arguably, the information problem-- or if you believe in missing markets, they're more missing in the rural areas than they are in the urban areas. All right.

So here's the model, utility over consumption and effort. Output is stochastic. Instead of writing output as a function of effort and capital plus a shock, this is a more general histogram. You've seen it at least once before, the probability of any given output given effort and capital. Households are either on their own or entering into contracts with a financial intermediary. It's partial equilibrium facing some exogenous outside return.

And you can think about financial intermediaries as being competitive if you want. You could also think about it as a stand-in for the community as a whole. I'll show you where this stuff matters. And you're going to solve this contracting problem for many dates, even potentially over an infinite horizon.

OK. And finally, how many people do we have? You could actually think about this as a risk-neutral household running a business-- sorry-- risk-averse household running a business facing a risk-neutral intermediary as if they're only two people. But a lot of what we do is easier to interpret when there is a continuum-- many, many, a continuum of household enterprises. Because then we can talk about the fractions of households who took effort, had certain capital, and experienced a certain output. That actually eliminates uncertainty from the point of view of the intermediary because all these things average out. Yes, Matt.

AUDIENCE: Can they re-contract each period [? for the ?] [? consumer? ?]

PROFESSOR: Here we're largely ruling it out. There is some limited commitment in the sense that they can walk away and go into autarky. That's about as close as we get in this paper. I kind of know something about how to extend it, but it's not going to be in the lecture today.

AUDIENCE: So how do you get-- just, do you get time variations in those data, or is it just a cross-section?

PROFESSOR: It's both a cross-section and panel. You use both aspects. The best way to think about having multiple intermediaries is there's ex-ante competition among them to service households, but there's something committing households once they sign up to a long-term agreement.

So what's the initial state for a household saying? Well, for sure at least the initial capital stock when we visit them in the baseline survey. And then the second argument depends on the financial regime. If it's like borrowing and lending, then it's their current assets or their net indebtedness.

Or if it's one of these contract regimes, it's essentially some utility constraint, some reservation utility as if it had been promised in the past. And I'll say more about that. That's the key to handling the dynamic incentive problem.

Clearly, something like unobserved utility is not seen, so we're going to have to parameterize it with the mean and the variance and estimate this unobserved distribution of debt or promises in the population. Timing-- OK, so those are initial states. Then capital is used in production along with effort, which may or may not be observed depending on the regime. Output is realized.

There is this pre-existing financial contract which determines the debt, say, that they can take on, or savings for tomorrow, or transfers as an insurance. And finally, after that financial stuff, they eat some and invest the rest. So it's kind of like a standard neoclassical setup. Capital today, you produce with labor, then you can either eat or invest.

The fancy stuff is happening with the financial part, which is the difference between what you have available and what you either give up or get depending on the financial contract. Capital depreciates. That's pretty standard. And of course, then you go to the next prime-- next period with all these prime variables on them, so the transition from the state today to the state tomorrow with a financial contract in between.

OK. In particular, we're going to talk about what goes on within the period in terms of the effort, which is induced or assigned, the output, which happens through Mother Nature, the decision for capital tomorrow-- and I'll go more into this promised utility for tomorrow. This pi thing is like the probability of this whole quadruple. And it looks daunting.

Now, from a statistical point of view, it's like a histogram. If you had a finite number of values that q , z , k prime and w prime can take on, then you just have a bunch of points in dimensional space. And then you can talk about what mass, what height of a bar, of a histogram bar that those values take on. So statistically, it's an easy way to summarize the data.

From the point of view of a financial contract, there's a lot of degeneracy. So it's not true, for example, typically that effort is random. For a given set of parameter values given incentives, there's one effort that's going to happen and none of the other happen. That's not inconsistent with this. It's just an extreme case where almost everything is zero except for one point of effort, and that's one.

It's also true that there might be a functional relationship between consumption and output, maybe a nice, smooth function. Or if you want to grid it up, you've got a series of dots that lie on a line. So then we talk about the probability of Q and C as if know it was a shotgun. But in practice, all the mass is going to lie on that line. So you're used to thinking about, let c be a function of q . Let effort be assigned. We'll take care of how much capital is invested, et cetera.

OK. It is true also that capital might be indivisible. We've talked about this. You get project ideas. You may or may not want to do it. The equipment is really chunky, building that warehouse for the chickens. So that's actually more realistic. And then the probability is kind of serious, which is, what is the probability you're going to do it, or what fraction of the population are going to do it?

When we talk about fractions of the population, you should be reminded of Rogerson. There it was you work overtime or you work or you don't work at all, and the issue was what fraction of people are working. So that was the first experience in this class where we had a probability number. And when we have non-convexities, this is the generalization of it to [INAUDIBLE] probabilities on-- potentially on more than one thing.

And there's a reason, and I'll show you when we write it down. It turns everything into a linear programming problem. I'm going to talk about particular utility functions, particular production functions. But really, we don't need to assume-- that's both the strength and limitation of this. I'm not going to show you a lot of closed-form analytic solutions. But on the other hand, we can solve anything numerically.

And Cynthia, again, just to alert you, was backing out from Euler equations with first-order conditions, which Lagrange multipliers are binding and why and so on. So it's not like you can't-- you can, in principle, do both, actually. You don't necessarily have to-- OK.

So let's just think about a standard problem without the lotteries. This is the autarky problem. Household enterprise has a realized output q , depreciated capital from last period. Actually, it's to q_i it's already realized. In autarky, there's nothing left other than deciding on what to eat and what to invest. So this is the capital stock for tomorrow. Oh, sorry. And then there's this effort z . z is entering in this utility. It's also entering in this production function.

You want expected utility? Fine. Just sum up over all possible outputs, taking into account Mother Nature's way of determining stochastically what those outputs are. And whatever you decide to take over to tomorrow enters in the value function tomorrow. OK? So we're looking for this infinite horizon solution. We're looking for a value function that's like a fixed point that solves the functional equation like this.

Now, again, when q_i is realized, you still have to decide on investment. That's why i is on both k' and q . But z is determined before output is realized. So that's just one number. You with me? OK. It looks-- this is an equivalent problem. And this looks familiar, except the i 's are missing.

So it's output plus depreciated capital less investment for tomorrow, deciding on [INAUDIBLE], but now we've taken this deterministic-looking problem and just replaced it with this probability object. But it allows all these special cases. Anything that can solve this solves this subject to grids, subject to approximating with a finite number of outputs, et cetera. All right.

Now, it looks here as if the choice object is this probability number including the probability of output, but how can that be? Because Mother Nature plays a role. So we kind of have to constrain these histograms to respect the relationship between k' and q , and namely, this object. We don't want to lose this object.

Well, it's a fancy way of saying the probability of event A conditioned on event b times the probability event b is the same thing as a joint event a, b . So this was the probability of q' and z' because I summed over everything else. This is a probability of c' given z' , and then this is, summing over everything else, the probability of z' . So it's a cute trick. So we're not going to cheat on Mother Nature. We're constrained to follow Mother Nature.

Another thing to-- say what's the dimensionality? Well, the state variable is k , so the question is, how many little k 's are there? If we grid up k in the small, medium, and large, there's three of them, for example. It could be 10. That's our choice in terms of the grid.

But whatever it is, we've got to solve this functional equation v and find the fixed point by iterating. OK? And the larger is the dimension of this, the harder it is to do that. Jan.

AUDIENCE: [INAUDIBLE] formulations [INAUDIBLE]? In the second case, as in the agents can do some randomization [INAUDIBLE].

PROFESSOR: Yeah, I should have--

AUDIENCE: [INAUDIBLE] is the first case [INAUDIBLE]

PROFESSOR: Yeah. If we don't get the grid right-- for example, if we have a continuum of values of z , and so on, then we have to make sure that the solution here is on a grid point down here. Or to put it the other way around, if the grid is serious then the household may want to optimize in terms of choosing a probability. And in that case, this is not just going to replicate. This allows more.

So if it wants to, we can find a solution, but it may choose to find something else. There are grid lotteries, but it's too easy to detect them in the output. You've got medium effort and high effort, and then the thing is putting probability on both of those. They're like adjacent points. So it's clear where the program's trying to get something in the middle that's just not available.

And if the grid is serious, then, as in Rogerson, that's fine. If we think the grid ought to be not there, and it's just meant to be an approximation, then you probably want to slice and dice it a bit more and search more intensively between those adjacent points.

AUDIENCE: For the second problem, is it possible that the solution [INAUDIBLE] random contract?

PROFESSOR: It is possible, yes.

AUDIENCE: But in reality, why the agent can randomly choose a [INAUDIBLE] So [INAUDIBLE] the first formulation, a agent can not choose a random contract. And in the second one, you can choose.

PROFESSOR: Yeah, yeah. I agree with your point. You think it's infeasible to randomize?

AUDIENCE: So I think if your assumption is that the [INAUDIBLE] agent can do randomization, then I agree that these [INAUDIBLE] formulations are equivalent. But if you don't assume that--

PROFESSOR: I am assuming they can do the randomization. And again, as in Rogerson, it's a way to smooth out the non-convexities. That's the main-- you know, we've turned non-convex problems into a linear program.

Oh, why is it a linear programming problem? Well, these p_i s are the choice variables. It's tempting to say, oh, no, it's supposed to be k and z . No, no, no. It's the probability of $[? \text{ } q \text{ } ?]$ So this is just a number, this utility, if you were to do a certain quadruple. And this is the probability of that quadruple.

So these are the fundamental policy variables. And what is the dimensionality of them? Number q cross number z cross number k prime. So this can get pretty big pretty fast.

And how many constraints there are? Well, in this case, just the one, although probabilities do have to add up. But essentially, one constraint. Well, not quite. It's for every q bar and z bar. So actually, there's number q cross number z constraints here.

I'm not going to belabor this. But when you go through the other regimes, you're going to keep an eye on how many constraints there are and so on. And eventually, we'll get to regimes that are actually nice and challenging to compute.

Here's the--

AUDIENCE: I still don't understand exactly what you mean when you say that the p_i 's are the choice variables and the q and the k 's and the z 's are just numbers. I mean, they still ought to be consistent. So when you choose a p_i , then you're doing--

PROFESSOR: Suppose you had a continuous consumption schedule, so c is a function of q . And then the paradigm is either you work hard or you work very little, so there's just two. So if you choose not to work hard, that's going to be reflected in output. But you're still facing that consumption schedule. Or you could work hard or potentially randomize across working hard or not. So the randomization could be in the effort.

AUDIENCE: So the p_i 's are functions? I'm choosing a function.

PROFESSOR: Hmm?

AUDIENCE: I'm choosing a function.

PROFESSOR: Actually, here, it's so general that you don't see any functions. You're just choosing mass points over a finite number of q 's, z 's, and k primes without constraints.

AUDIENCE: I am, quote unquote, [INAUDIBLE]

PROFESSOR: Yeah.

AUDIENCE: Isn't it more like choosing a distribution, though, instead of saying--

PROFESSOR: If these were-- yeah-- if these were a continuum, then this would be like a density, a multi-dimensional density. But I don't know how to compute those.

AUDIENCE: It's not-- so correct me if I'm wrong, but it's not that I can put anything next to it. I have to plug in to the u next to it the corresponding stuff in there. But then I just compute every single thing inside there and choose the best one after I compute everything.

PROFESSOR: That's right. So we specify the grids. So we know the set of feasible choices for all the quadruples qzk prime. For any particular one, we know what this real number is. You use Matlab code to generate these weights. In this case, it's a weight on it on the objective function, or you use codes to generate this guy, and then there's weights on the constraint set. So

AUDIENCE: So now I understand that. But then, in some sense, I don't understand what the difference from the first case is-- like, for on the first problem. That the probability of q given k and z could, in principle, be non-linear. And so we're just--

PROFESSOR: It's the extra gain from randomization. I mean, suppose-- take the point of view that there are a continuum of-- well, so far we don't have the social planner. I mean, this is just an individual optimization. But I will try to answer your question.

When we have a village-wide resource constraint, we've got to decide how many resources to use up in investment and how much to leave over for consumption. And it may be that you don't want everyone doing a large, chunky project. You want to just take some people to do it and some people not. And then you know whatever is produced as output, use that some of that for consumption and some to save for tomorrow.

So from a social perspective, you rarely want everyone doing the same thing. You want to choose the fractions or people doing one thing or the other. Does that help?

All right. So this is with borrowing and lending, so it's essentially almost the same, except you can add to your current resources by borrowing, except you've got to pay back your loan from yesterday. So now you've got two ways to intertemporally reallocate consumption, one through the capital stock, the equipment, and the other through financial borrowing and lending. And it's already written in lotteries. By the way, it enhances the state vector from the capital stock to the current debt as well.

Now we can knock off savings only. That's where a basically, if B means borrowing, then basically, it can't be positive if you're not going to allow it. Or you could allow some small amount and have a positive but low amount that you could borrow as b_{max} . So we can do savings only, et cetera. If you don't want to limit borrowing at all, that's fine. Then you just let it take on any value subject to [? grade ?] issues. Yep?

AUDIENCE: When you do this kind of stuff in Matlab, to do these different scenarios, is that just [INAUDIBLE] one line of code?

PROFESSOR: Which?

AUDIENCE: Like, to have [INAUDIBLE].

PROFESSOR: Yeah. Well, actually, you can almost generate it from the grid, because you might have a large set of possible values for borrowing. That would be the unrestricted problem. If you want savings only, then you just cut off all the positive borrowings or anything in between. So this one you can handle in terms of generating the underlying grid.

Now we get to these mechanism design models including full information. Now we're going to have to go back to the households as a group dealing with this financial intermediary. The main thing is this-- let's look at the household because we were looking at households a minute ago. Now it's as if they surrendered all their output to the bank-- bear with me-- but they get some of it back in terms of transfers.

Actually, a more natural way to write it would have been give them q and let them pay back loans, state-contingent loans. It's equivalent. The transfer can be positive or negative. And it's kind of the basis of consumption, but it's adjusted, as usual, by investment. So then the contract has to solve for these transfers, and the transfers are a function of q .

Now, as I said, the way this is written, the q ends up with the bank, but the bank is paying the transfer. So this is basically a surplus generated from households. What fraction of households have output q greater than τ , so the bank is getting money from them? Well, that's this π . That's the fraction.

Well, you know, it's infinite horizon. So this is the surplus generated today-- or if negative, the loss-- summed up over all surpluses and losses depending on who's at what states and what the contract assigns. And then you have tomorrow's profits. This is a small, open economy. There's an interest r , 1 plus a little r . So this is just the present value of tomorrow. So this is surplus today plus profits today plus profits tomorrow. So it's as if the bank is trying to maximize overall profits.

Now, what's the constraint? Why not just screw the households? Well, the answer is-- it's not quite right, but you can think about this as a reservation utility. So you can't take too much away. Otherwise, the households will cry foul and walk away. OK?

Actually, technically, this promise was predetermined from the previous period. Or equivalently, part of the control variable is the promise from tomorrow on, w prime. It's easy and yet amazingly powerful. So just think about incentives. You have long-term contracts.

Should I work today? I can get rewarded or penalized depending on what my output is today. But it's not just a static contract. There's tomorrow too. So maybe my history of observed outputs will be used tomorrow in terms of the insurance and credit contract I'm going to get, which in turn could be used in the third period. And we're like, oh, my god. That's a really big object.

But no. What do the households care about? They only care about their expected utility. If I'm looking forward to tomorrow, I don't have to solve tomorrow's problem as long as I know what the utility consequences will be. So this is kind of a reduced form way of handling the multi-period incentive problem.

What I'm saying is the household's expected utility is the utility outcome from today plus the expected utility tomorrow. It's not only-- it's the probability of being assigned w prime tomorrow jointly with output today. So now, not only is consumption moving up and down with output today-- and it will move when there's limited insurance-- but also, tomorrow's utility will vary up and down.

And because you have concave utility and this force for inter-temporal smoothing, you won't just do one, and you won't just do the other. Actually, you'll load as much into the future as possible. Front loading is a bad thing. The longer the horizon, the more powerful the incentives, because you've got the whole future. So there's a lot of economics behind this sort of innocuous-looking two-period problem where w prime captures tomorrow.

All right. That's actually the full insurance problem, and then we can add moral hazard. So basically, what this says is if z bar is assigned today in the contract and the guy actually does it, he has to be wanting to do it relative to-- and the bank wouldn't know even though z bar is being recommended, the guy is doing z hat, shirking. So this is a standard moral hazard constraint.

Of course, the inequality makes them want to do the recommendation. But the program has to evaluate any possible deviation a shirking agent might do and make sure the utility consequences for the agent are worse than following the recommended plan. Now, again, because we've embedded Mother Nature in this π thing, we have to re-normalize the probability. I actually showed you this once.

But I'm sure with discounting, you can't remember backwards. More on that next Tuesday. But we just adjust the likelihood to reflect the fact that now z hat is taken rather than z bar. And it's not maybe obvious where this is coming from. It is nevertheless true. You just have to write out all the conditional probabilities and start summing up. I can give you a reference if you want to see where it's written out. Questions? Yes.

AUDIENCE: I just want to make sure that I understand this formulation. So I think the [? language ?] of this formulation is that if you don't have this formulation, first of all, you need to solve the incentive-compatible contracts one by one. And in this case, the randomizing contract is like randomize your wealth. So you may need to solve one problem. So I [INAUDIBLE] you randomize--

PROFESSOR: Well, actually, you're suggesting a simplification that we're not using, which is any random-- oh, you mean the w, the promise?

AUDIENCE: Yeah. I think the advantage of the second formulation compared to the first line-- in the first line, you need to solve first all the incentive-compatible contracts, and then you'll do an optimal randomization among these contracts. Right? Here you can just-- x and t, you can randomize over this z and the k, and that will give you a similar result. So essentially, you only need to solve one linear programming problems instead of a lot. Is that--
[INTERPOSING VOICES]

PROFESSOR: Well, I guess. We're searching jointly over all the possibilities. Now, it's true on the one hand, it's just a linear programming problem. So we can just get the best code available and use it. On the other hand, there's lots of variables, lots of states, lots of constraints.

So it's not like it comes for free. If I knew something analytically about the underlying contract, it would be good to use it. But if I don't know, then I could just guess wrong and put some functional form which is incorrect. And that's why-- that's what's good about this.

I don't-- now, what was Victor talking about? Victor was talking about going back to deterministic contracts and then solving them with a nonlinear optimization problem and then comparing that solution to these linear codes. And I'm not sure if he showed you at the end, but when the grids are really coarse, the linear program seems to work, but it can be a really bad approximation. So there's still trade-offs. It's not a miracle. Other questions?

All right. So now we can do limited commitment. So this is what I mentioned. You can go into autarky. You can decide to do that after your output is realized. It's like, I'm not paying into the risk-sharing group.

Great. You financed my project. I'm the big boss now, and I'd just as soon be on my own. So what we say is-- first of all, we compute the value function for autarky. Well, that's cool. We already did that. That was the first financial regime. So we already know this guy.

And then the contemporary situation is you've got output plus depreciated capital, and then you could walk away. No transfers. They're gone. You just keep it all and decide maybe what you want to carry into tomorrow.

So we can call the solution to this thing ω . By the way, the z is already foregone. You've already made that decision. It was in disutility part. You've already got capital. You already got funding. Now you have output. Can't go backwards, but you can walk away.

So this is basically the maximizing utility. And we make sure that if you follow the plan, you're not going to be tempted to do that. So compute v, then compute this. Then we have this, and then you impose this as a constraint. So this is a limited commitment constraint.

We've talked about collateral constraints. It's very related to that. You can walk away, but you might-- with collateral, you have to sacrifice something. Not here. They actually keep all of their capital.

If we had financial savings, then they would lose that. They would lose the stuff they had in the bank.

AUDIENCE: So what keeps me from doing that is that the continued valuation is going to be on that side?

PROFESSOR: Yeah. This is playing ball, and this is being tempted to pull out. It doesn't mean that it isn't binding on the solution. This can do damage. It could be a binding constraint and have a big Lagrange multiplier. So the solution will look different.

But you never see the out-of-equilibrium event that they walk away. But the damage can be done. Not getting a whole lot of capital, not having a whole lot of insurance-- those things happen. I mean, again, it's rich guys who would prefer, say, not to pay into the system. They're going to be tempted. So you start eliminating the people that pay in, and that starts limiting the insurance, et cetera, et cetera.

And then finally, we have this hidden output. So here, the idea is income is produced, all right, but the outsiders don't see what it is. OK. But then you say, well, tell me anyway. Send me a message.

So the idea here is q is actually-- q bar is actually realized, and the business says so. Believe me. My profits are low. Or-- and they are low. Or q bar is realized, but we have this counterfactual where he says something else about q .

Maybe in that particular period, it was advantageous to say profits are really, really high, or it could be vice versa. When they're high, you say high. When they're low, you're tempted-- sorry. When they're low, you say low. And when they're high, you're tempted to say low.

All right? This allows for any q that they will, quote, tell the truth. Well, great. It's just another inequality. No problemo. We know how to generate constraints. But again, it will do different damage. There's going to be consequences for the underlying contract. And remember, the goal here is to get to the data.

You can assume these parametric utility functions-- constant relative risk aversion, power sigma, disutility of effort, a power of Frisch elasticity. OK? So 3-- we're going to be limited, actually. We're going to estimate sigma and theta, α equal 1. Here's a production function. We actually load in an observed histogram. But you can do constant elasticity of substitution if you want. Yep?

AUDIENCE: So you [? test ?] one where you can verify a state? Because now the output is hidden, right?

PROFESSOR: Yeah. I didn't do it.

AUDIENCE: Why not?

PROFESSOR: I forgot about costly state application. But anyway, it's doable. Hong's working on it. And these other things, like the discount rate, the depreciation rate, the outside interest rate, and so on. So this is very much in the spirit of calibration. These numbers are similar to numbers you've seen in various papers in the classes trying to be somewhat realistic. Yep.

AUDIENCE: When we're [INAUDIBLE] looking at the matrix p of q given c and d_k . Is that just for setting up the moral hazard constraint?

PROFESSOR: No. We need that in general.

AUDIENCE: [INAUDIBLE]

PROFESSOR: So we're going to say that we see effort even though, in the models, sometimes it's unobserved. I mean, the question is what an outside lender would see, not what the households tell us when we interview them.

So anyway, we take the stand that we see a good version of it, and capital as well, and we see output. And I think that-- let's see where that is. Oh, my god. It's way-- there it was. So there is an empirical histogram of capital, labor, and output. I don't know if you really get a three-dimensional view of that thing.

I can kind of see it's bending over here, and then [INAUDIBLE] flip on me. But-- so we can load that in, or we can actually say, no, no, no, it's CES, some elasticity of substitution, and estimate whether it's [INAUDIBLE] for linear or stuck in between. Yep?

AUDIENCE: [INAUDIBLE] question. On graphic views, is there a way to understand [INAUDIBLE] whether those--

PROFESSOR: Yeah, but it's kind of hard because I don't know what this thing is doing. It doesn't look visually like it kept going down. It may actually come back.

By the way, it makes my point. We don't have to assume any kind of concavity in the underlying primitives. If it has this sort of scallop shape, wonderful. Bring on the lotteries. Then it will span the arc line. We can even allow risk-loving households here. I mean, there's really no restriction in the underlying perimeters. We have parameterized utilities to make people risk-averse, but in principle, we don't have to put restrictions on what we load in.

OK. So preferences, technology-- I'll just say a few words about dimensionality. As we go through autarky savings, full information, moral hazard, et cetera, you can start counting the number of linear programs that need to be solved, the number of variables in each program, the numbers of constraints.

I didn't even dare show you this unobserved investment one, but the hidden output is getting up there too. We actually have a technology to compute these things with hundreds of thousands of constraints. I mean, the commercial code is CPLEX, and there's open freeware that's comparable to it. It's the latest Princeton interior point. It's not just a simplex algorithm.

You can use it as a student, actually. It cost me a couple thousand dollars. But anyway, so we can handle fairly large numbers. Now, that said, what's the tension here? Well, you're going to see not only do you have to iterate off the value functions and solve the linear program at each iteration and solve-- this is just one step.

But then we don't even know what the parameters are. So then we got to do it for all the set of possible parameters and generate a likelihood. So these guys start to get demanding not because you can't solve it once in 20 or 30 seconds, but because you've got to do it hundreds and hundreds of times.

So you can understand why there's a big interest in having relatively efficient code. It starts to constrain you in terms of the kinds of problems you really want to consider. All right. So how does it work computationally? Well, once you solve for the optimizing policy π^* , you have a transition function, basically. You start with promised utility and capital today, integrate out over everything else other than capital and utility tomorrow, and you get the probability of that.

So you get this Markov object, right? It's simple enough conceptually-- the probability of states tomorrow given states today. So that's kind of the underlying dynamic engine that's chugging along for each one of these financial regimes. And then any time you pick up certain w and k in the solution, you can generate the contract, because that's the stuff that we just integrated out. But it's still there. You can still use it.

So we can generate histograms, CQ configurations. We can have two cross-sections. We can have panel. There are some limits in terms of the length of the panel that we can actually use. You still with me?

We've got to estimate some parameters, so we're going to have these underlying structural parameters. We don't see, say, the distribution of promised utility, so we're going to imagine that's generated maybe like a normal distribution with a certain mean and a certain variance, and we have to estimate those.

Better yet would be a mixture of normals which can approximate any old thing you want, but that raises the dimensions. And we're trying to be lean in terms of numbers of parameters. So what do I mean by likelihood? The probability of getting observables y , which could be in a static cross-section, values of consumption, output, investment, and capital. What's the probability of seeing a particular configuration like that given the observed capital today as a function of these underlying parameter values?

So because the model is already using lotteries, you already get the probability of these objects. It just sort of comes for free as part of the optimizing solution. Now, the next thing is, do we see things perfectly? No, not necessarily.

We can put in-- let me jump a second. We can put in measurement error and say if c^* at j were the true value, we don't see it. We see some measured version with error. So this is classic econometric contamination, classical measurement error. We can put that on everything that you like.

And then the program would say, what is the probability the underlying output and consumption would be c^* and q^* ? That we generate from the code. And then we can say, well, what we see is c having q hat, so that could come from any c^* and q^* . So you get a new sort of distribution of observables. Well, that's what I just said in words.

Well, what's the point of the likelihood? We actually have the data, and we see people with a certain capital stock getting a certain output and having a certain consumption, investing a certain amount. These are our observables. And so we see the histogram in the data, and now we have a histogram in the model. So we can say does it-- what parameters would best rationalize the data if the data came from that financial regime, a particular one?

So let me-- so I said we have these data. This is Kansas and the Rocky Mountains for Thailand. A mixed metaphor, but it's also the investment. Consumption is, again, pretty flat, so it's going to suggest a regime where there's a lot of consumption smoothing against income fluctuations. On the other hand, investment isn't flat. And I haven't even shown you the transitions in the capital stock.

So when you-- you pick the rural data, and you use, say-- say choose consumption and output only. There's a tie. If you believed it was moral hazard regime-- for example, it's estimating the degree of risk aversion at 1.02, the Frisch elasticity at 1.6. It's got a certain mean and variance of this underlying unseen distribution of promised utilities and an estimate of the measurement error.

And it does this for each financial regime against the same data. And then we use this sort of information criterion, which Vong created, which allows testing across non-nested regimes. In this case, it's a tie. But if you went to the investment data alone, it would be savings, savings only. So in the rural data, we get a limited regime.

I probably won't have time to show you the Monte Carlos. The Monte Carlos are like this. You pick a set of parameter values, maybe the ones we estimate in the data, for example, and then generate the data from the model, and then go through all of this. And depending on the degree of measurement area-- error that you use to contaminate the data, you can-- especially when you use joint consumption and investment data, you pretty much get back what you put in, which is reassuring. But it's a bit black boxy. We don't have an analytic proof.

So if you use consumption data alone, two cross-sections, you can use two-year panels. Then there's-- in the rural data, it's kind of hard to pin down too much. In terms of the regime, moral hazard is in there. Sometimes full information is in there. Limited commitment is in there.

But again, when you use the investment data alone or use the joint data, it's pretty clear that savings only, no borrowing, buffer stock model is the one that the data like the best. But when we-- if we use the network alone-- you're saying, well, you've contradicted-- no. Like, what Cynthia and I do, and so on, we can actually get the full information regime out of the consumption data. So that's very reassuring because in those other papers, we didn't use this method. There's a lot of consumption smoothing among the networks.

And if we go to the urban data, we get this result that even when we use the investment data, it's the moral hazard regime, not the limited savings only regime. So they're different across the two specifications, although it is true that if you used the investment data alone, it would still like savings.

So there's still this-- and we've talked about this-- money doesn't flow from unproductive to productive people the way these relatively unconstrained regimes would imply. However, if you looked at the consumption data jointly, then it's one likelihood. And it decides in the urban area that, oh, well, moral hazard actually fits better than savings only, but the reverse is true in the rural data. So although the investment data doesn't fit perfectly well in the urban data, the verdict is we do a ton of robustness checks. And I'll skip it.

So let me just-- I really want to say something about heterogeneity, but I'd rather show you three slides before I quit-- namely, in the actual rural data, this sort of diagonal here represents the persistence of the capital stock. So it doesn't move much. It moves very slowly. You can hardly see mass off the diagonal.

In the urban data, it's still sluggish, but at least you can move away more. If you have this wonderful financial regime, you're going to adjust almost instantaneously to the observed productivity. So this speed of adjustment of the capital stock is really the thing that's kind of pushing the likelihood toward limited financial regimes. By the way, these are the best-fitting regimes through the lens of the model, which still have an excessive amount of smoothing.

These are the time series-- levels of consumption over time, levels of the capital stock, levels of income. We did not use all of this data. This is like setting aside certain things in the data and then looking after the fact as a new criterion. And we're doing pretty well. The standard deviations, we do less well, but for obvious reasons that there are stuff in the tails that the model doesn't like.

But if we smooth off 1% of the tails and kind of get rid of it, we actually bring the standard deviations generated by the model much, much closer to the data. And we can actually use the level on borrowing, which we didn't use [? an ?] [? estimate, ?] and got really, really close to that. And finally, here is this data on the return on assets with these low-wealth households having high returns. And the savings only regime which fits best is trying to do that.

And this is the urban data, which is already a little more dispersed, but it still has this low-wealth, high-return guise. And it can generate that somewhat with this moral hazard. So these are the familiar objects that we've been looking at. And finally, let me just end with this.

Sorry. I guess that discussion at the beginning of class took a chunk of time. So if you believe the structural model, you can do policy evaluation. And here we're doing something simple like just changing the interest rate. Now, the point is some people win, and some people lose. And also, who are the winners and the losers depends on the financial regime.

So this is the consumption equivalent welfare gain of lowering the interest rate. Negative and positive-- this is as we vary the level of current savings and the level of assets. But this is what it would look like if it were not the savings regime, but the borrowing regime at the estimated parameters we get. And then this is the difference of those two.

So it's another difference and difference, but it's a difference in the welfare criteria across two financial regimes. And the point is it matters quite a lot, if you did this interest rate subsidy, what the financial regime is in terms of who wins and who loses. These are non-trivial gains, and so on. All right. I'll quit there.