[SQUEAKING]

[RUSTLING]

[CLICKING]

ESTHER DUFLO: So let's start with-- so I want to try to do two things today-- start with the neoclassical model and its implication and then move from the neoclassical model and its implication to a more behavioral model. Which-- both because this is pretty essential to understand savings-- as we will see, the neoclassical model fails to fit the empirical facts in a pretty obvious way-- and second, because the model that is useful to think about savings is a model that is useful to think about many problems in development that have some intertemporal flavor to them, including, for example, health which is also an important topic we're going to see soon.

So let's start with how in macro-- I hoped you would have been taught savings already, or will be taught savings already. Certainly I was taught savings when I taught macro, which is the basic intertemporal optimization model. So I'm going to go quickly because I'm going to assume you've seen it, but it is worth making sure we're on the same page. And if it goes too quickly, you have the slides for easy reference.

So suppose that an individual follows a variable income process and an interest rate. How much-- the question of intertemporal maximization issue is, how much they should save-- and borrow if they can borrow. So we start with the utility at time t defined as the expected value over all your lifetime, of all the instantaneous utility of consumption in all periods as viewed from period t.

We can be very generous, but, of course, very quickly we simplify it in some way, or we parameterize it in some way. Otherwise, it's getting difficult to say anything. So, for example, you could say that you're discounting the future at some rate delta.

So a common example becomes this one that you are discounting at rate delta. So viewed from period t, it's the discount rate of u of ct with prob-- some utility function that we assume constant. So there is, of course, a lot nested in that model. In particular, the idea here that is intertemporally consistent. So you are valuing tomorrow's consumption exactly like today's except with some discount rate that is constant over time. So this is-- a lot of what we are going to today is going to basically say, this is actually not very accurate-- maybe not the best representation of how people think about the future. And that's what we're going to change.

Now assuming that it's the case for the moment, then the individual faces a terminal condition. So in finite time, it means they cannot borrow. You die with that. And they also don't want to leave any money at the end of their life. And there is an equivalent of that in infinite time.

And then there is otherwise, an asset and accumulation equation, which we already saw when we discussed education, which is that your assets is what you had plus what you earned minus what you consumed, multiplied by the growth interest rate. To solve this problem, we use the standard dynamic programming tool. So we write the value function of the problem, which is the expected utility of an individual who makes the optimal choice of savings today, given a certain level of assets he has. So that's the Bellman equation.

So the value of your assets is the maximum over the best-- the maximum value over all any possible value of omega t you're saving, of the instantaneous utility of consuming a certain amount plus the expected value at t plus 1, of the assets you're going to be left with if that's your choice of saving and borrowing. All right? So at this point, we are making no assumption about whether-- I'm talking about savings but it could be negative saving which would be borrowing. So that's the Bellman equation.

Once you have the Bellman equation, all we need to do is to derive it in order to get the optimal decisions. We take the first order condition and we get that the marginal utility of consumption must be equal to this term, the expected value at time t of the derivative of the value function at t plus 1 multiplied by the interest rate. Now that's not yet very helpful because we don't know what this guy is.

So we need to take one more step. We need to express this guy as a function of a quantity we know. And for this, when we derive the Bellman equation with respect to assets A at the optimal saving rate omega. And the result is that the derivative of the value function with respect to assets at the optimum is simply the derivative of the utility function of consumption.

Why that's true? Why is that true? What-- why don't we have that thing to mess this up when we derive with respect to assets? Yeah, Bishon, You look like you--

AUDIENCE: Is it the envelope theorem?

ESTHER DUFLO: That's the envelope theorem which is, at the optimum, then this is the derivative of the-- of this with respect to-- is the envelope theorem. So we are great. We are happy.

It becomes very simple. We can combine equation two and three, the first order condition and the derivative of the Bellman equation with respect to assets, to get the lambda of C or the marginal utility of consumption-- is the expected value of the marginal utility of consumption in the next period times the interest rate. So that's the Euler equation, which is a pretty central value of economics.

So it says that the marginal utility of consumption today must be the marginal utility of consumption tomorrow, up to a discount factor. Now if the-- we have some special cases that are well known. If we have exponential utility and a constant interest rate, the Euler equation simplifies very nicely. And what's the result there if I put it on the board? This you probably saw even in undergraduate. Macro-- if you've taken undergraduate macro?

Or you can recalculate it in your head. That's how it simplifies. And if in addition it's actually quadratic, so it's even simpler. Then it's fur-- and if r equals to delta, which it should be in equilibrium, then the equation further simplifies to, the consumption is equal to the expected value of consumption tomorrow. So consumption is a martingale. That's the whole result. So it is-- make sense?

Now if consumption has-- of course this martingale result will not hold for measures that-- for different utility functions. In particular, if the utility function is not quadratic, the derivative of it will be a little more complicated. So it will not simplify to this very simple formula.

And there is a case that's particularly interesting, which is one where the utility function, of course-- the derivative is positive. That is better. The second derivative is negative. That is your risk averse.

And the third derivative is positive, which is called precaution. So what does this-- that becomes hard to think about, third derivative of utility function. So if you were to express it, the third derivative being positive, in words, what might that reflect?

AUDIENCE: That you get less risk-averse as you get wealthier?

ESTHER DUFLO: You get less risk-averse as you get wealthier, exactly. So that's why it's precaution. So the poor particularly dislike changes in their income. And therefore they are going to want-- the poorer you are, the more they are going to want to be in the buffer stocks of savings, or ought to be. They ought to be able to borrow in order to compensate a loss in income if they can.

I do think there is something on the fourth derivative being negative, which is called temperance. But I forget exactly how it interprets in world. But I do think there is one more that people have been thinking about. So exactly like I said, positive-- third derivative means that at low value of consumption, the variations are more painful than if you are starting at high value of consumption.

That means that an increase-- the consequence of the third derivative. Why are we interested in the third derivative? Well, it's because our Euler equation is expressed in terms of second derivative. So if we want to say something about it using the-- using our famous, looking at [INAUDIBLE] things, we need to say something about what it's-- the second derivative of the Euler equation.

So to use the Euler inequality-- so the consequence of the third derivative being positive is that an increase in the variability of consumption will increase the height and side of the Euler equation. That means that consumption must decline and savings must increase if people aren't able to borrow to compensate in period t-- to compensate for an increase in the variability of consumption. That's just restating what you told us. So that's kind of a testable prediction of this model, which is, there should be a precaution-- what's called precautionary savings. So if you think that your income tomorrow is going to be very variable, you should start putting some aside. All right?

So there is a test-- very classic paper by Christina Paxson is a test of this hypothesis. The Euler equation implies that the change in permanent and transitory income should have different consequences on consumption. So if you take a change in permanent income, what should you do? What should you do with it?

They just got the Nobel Prize. There is a transitory component. He's going to get some big check. And there is a permanent component, which, he's going to be invited to give talks that-- where they're going to pay him a lot of money to do it if he wants to. And then he can also bargain with MIT to get some increase in his wages.

So that's so he has both-- he just experienced on Monday a transitory income shock and a permanent income shock. So what should he do with the permanent part of the-- what should we see on consumption? What should he do with the permanent part of his wealth?

AUDIENCE: Increase consumption.

ESTHER DUFLO: Yeah. He should go ahead and increase consumption to that extent. How about the transitory part?

AUDIENCE: [INAUDIBLE]?

ESTHER DUFLO: Yeah, that should be saved. That should be saved. And maybe it should be saved a little less because this-- his permanent income increased, so therefore he becomes less risk-averse. So he needs to save a little less for it.

So that's an empirical test, not just of precautionary saving by the way, but any-- the Euler equation itself, pretty much with most forms of the utility function. In the simplest case, a permanent change in income should be entirely consumed and a transitory change in income should be entirely saved because that's going to go away. It's going to-- it needs to be smoothed over infinite number of periods so it's really nothing.

So does this correspond to what we see? So Paxson's test is going to test this using data from rice farmers in Thailand. And her key idea, which was then used in some form of another in many other papers, is that if you use rainfall variation, the year to year variation in rainfall changes transitory variation in income.

A permanent change, for example acquiring irrigation, would make a farmer permanently richer. That should be consumed away. But the fact that the rain happens to be good or bad this year shouldn't really affect our consumption at all, because it should be smoothed away either by savings or by borrowing. So the test becomes, whether that translates into larger variation in savings than a permanent source of valuation income, such that individual ownership of land or whether people have irrigation, or whether it's good land, bad land, and so on and so forth.

So if you write your estimating equation as, saving is a function of permanent income and transitory income, you would like to see a 1 here and a 0 here, and a 2 being-- oh sorry-- a 0 here and a 1 here because this is saving with the-- and then potentially-- and then the variance of income. People who have income that are more variable, irrespective of the realization this year, might want to save more. Then we might in principle interact valuation with permanent income because the poor would want to save more.

Traditionally what people were thinking about is, oh, let's treat the transitory income as a nuisance because it's just noise in the income process. Let's try to predict the part of the income that is permanent and see whether people respond to permanent changes in to things that affect permanent income. So people run this type of regression of savings on-- or income-- sorry. That should be-- y. So the y could be savings or consumption on a set of household characteristics. That would predict more money-- so whether people had a job or something like that.

Now it has a bunch of problems, of course, because people who have occupation, education, for example, which was the thing that people put here on the right-hand side to predict permanent income-- they also predict a million other things. So it's pretty bad. And so the idea is that, well, that's actually easier to make the case, that we might have a random or quasi-random form of transitory income than a random or quasi-random form of permanent income.

So maybe the identification of this coefficient is much easier because we can just forget about the permanent income or not pay much attention to the coefficient we estimate. But focus on valuation and income that can be credibly considered to be random. If they are random, they might disappear in the next year, but then that's good since this is precisely what we're looking for here. So that was the insight.

And this was before any of these natural experiments were really developed. So it was a particularly clever insight by Christina Paxson. So instead of trying to find instruments for permanent income that are uncorrelated with [INAUDIBLE] and consumption, or patience, which is kind of a hopeless enterprise, you can find variations in transitory income that can be much easier to consider than random.

So what does she do? This is a very small table but we only need to look at here. She's saying, well, let's look at-- not rainfall in a place, but the deviation in a particular year from the trend. That could be reasonably assumed to be random with respect to each individual farmer. And that can be reasonably assumed to be transitory.

And in fact, when you look, for example, at variation lagged one period-- so this is the variation this year. You can see that if the rain was good this year, people make more money. This is a t-statistic, so that's a t-statistic of 2.

Then she'd like to see that, well, if it rained last year, the income should not be larger. That's actually not the case. The t-statistic is also high. But for the rainfall last year and the year before, it doesn't affect income anymore.

So maybe the income process is one that responds to this year's rainfall and last year's rainfall. She also makes the argument that this year's rainfall has no predictive property of next year's rainfall, otherwise the realization of this year would give you some indication of what's going to happen next year. And therefore you need to take that into account in making your anticipation about what your income is going to be. That's going to affect your income.

So she shows that in fact, income is not serially correlated to [INAUDIBLE]. And therefore what we can do now is to see whether it appears that saving is strongly correlated with this year's rainfall variable. And she has various measures of savings. That's a bit of a problem in any empirical data, is to measure saving-- it's pretty bad. This is no exception. This is pretty bad.

But she's trying three measures. We don't need to get into the details because frankly, it's more about the method than about the results given the measure of savings that she has to play with. But she's finding it to be quite correlated with savings at sort of reduced form.

Then once you have that, you can run an IV and she was saying, well, what's the effect of transitory income on savings? So this one is two steps, which is just like literally constructing the IV by hand. This is maximum likelihood. That's closer to a 2SLS estimation. She also a instrument for permanent income, but we can forget that for now. I mean, we can look at it, but we know that it's pretty hard to find instruments for permanent income.

But focusing on these guys, you can see now using the rainfall in period one and two as an instrument for income, which is seeing pretty strong coefficient of this transitory component on savings. And then once you've done that, you can try to see whether you're rejecting one, and you don't hardly ever reject one. So she cannot-- you cannot-- on the other hand, you can reject that these two are the same. So that would be a test that people-- of this particular implication of the standard utility maximization model, that people save most of their transitory income and consume most of their permanent income. Yeah.

AUDIENCE: I had a question on how to measure savings.

ESTHER DUFLO: Sorry?

AUDIENCE: Why is it hard to measure savings?

ESTHER DUFLO: Because people are saving all sorts of ways. So you can't easily-- so these are people, farmers in Thailand. They don't really have bank accounts at that time. They were there saving their money. Today they probably are doing it more.

But in the time where she collects the data, or even today if you went to an Indian household, they are saving is like all over the place. They have some in the bank account. They have some in the form of grains that have not been eaten this year. They have some in their mattress. They have some that are saved in the form of having lent them to a cousin-- you know, second cousin twice removed-- that might be giving it back.

And so accumulating-- like summing it up over all the possible way that you actually can save is very difficult especially when you try to think about, what's the value of that saving? Is the second cousin twice removed going to give it back or not? So that's one-- so the direct measurement of saving this stuff.

So what people then try to do is to say, well, let's try to do it by measuring income minus consumption. That's saving. And the difficulty with that is that, first of all, income is very poorly measured because people have all sorts of sources of income, many of them informal. And then they don't know how much money they made. Like if they're farmers, like you need to value everything minus the input. It's like super painful.

And consumption is better measured but not great either. And then you have durable expenditures. That is a form of consumption but is also a form of savings. So you should really value the services from that consumption in order to have the full measure of consumption. OK. And you just got me started. It's pretty hard to measure savings. Yeah.

AUDIENCE: Related to the last point about durables, how should we think about illiquid assets, or [INAUDIBLE] liquid forms of saving [INAUDIBLE]. Like suppose I-- I don't know, I buy a cow or a--

ESTHER DUFLO: A bullock, yeah. So it's such a super interesting thing. In fact, there are interesting-- there-- there is a very nice paper by Fosten and Rosenzweig-- Rosenzweig and Wolpin-- old paper on bullocks, showing that people invest in bullocks. And the issue is, bullock is not only illiquid, but it's also a productive asset.

And so in bad times, you have inefficient liquidation of bullocks because people-- all their saving is in the bullock and they need to sell to consume. And then they don't have the bullock to actually kill. So that creates this-- so this is about sort of, the-- we'll go in a minute, talk about credit constraints, savings constraints. It might be difficult.

One of the reasons why people don't save in to bank accounts is they don't have access to bank accounts. And in particular, they don't have access to bank accounts that are cheap and easy to manage. And so they save in all sorts of ways, including potentially in bullocks, which has this feature that it's both productive investment and savings.

In particular, bullocks go by pair. So if you have to liquid one bullock, then you have your one bullock that still needs to eat. But you don't have a pair of bullocks so you can't really use them to harvest. It's a partic-- that creates one more. It's a nice paper, although there' isn't-- that's a great question.

So this is where we're going to constraint. Let's start by adding-- we already started to talk about savings constraint and talked to-- we'll go back to that a little bit later. But there are also-- what the macro literature certainly has paid attention to is the possibility that you might not be able to borrow. So then that's really adding an additional constraint, that the assets have to be positive at any given time, or the consumption cannot exceed the sum of the asset and the income you're making or cash in hand.

So now you have two regimes. In the unconstrained regime, you have your Euler equation which is satisfied as usual. Your marginal utility of consumption is the expected utility of future marginal utility of consumption up to a discount factor. But once you constrain, you just consume everything you can. And therefore the marginal utility of consumption today is too high compared to what-- you make it as close to the desired utility as you want, but it's too high. That means you're not consuming enough. All right?

So if you put the two cases together we have that the marginal utility of consumption is the maximum of the marginal utility of consuming everything you have today, and the product of the standard Euler equation-- the expected utility of tomorrow's consumption. So Deaton has a pair of classic paper on a book, showing the solution to this problem can be characterized by a consumption function where your consumption is a consumption of what your cash on hand in total. With a C-A-R-A production function, that's-- utility function, sorry-- it can be shown that it has the following form. If you have less than some thing, you eat everything you have. If you have more than that threshold, you save a fraction of the difference between the cash you have in hand and that threshold.

So if your income is less than-- so you can write it-- so, basically, what you do is, when you can, you save a little bit to constitute your buffer stock. And then when you have enough, then you can just consume normally.

So the propensity to consume out of your assets will be smaller when your assets are higher than when they are low. And when they are low, eventually you're going to-- when the assets are very low, you're going to deplete them completely if the next valuation of income is low. And then you're going to end up with zero assets.

So what you can do is to simulate and to see how much you can achieve visually. And that's something he has in his book, *Understanding Consumption*-- which is actually a very nice book. And he has the simulation-- is to say, well, suppose that your income is IID with mean 100 and standard deviation 10. And you construct a rule of thumb for consumption that actually looks pretty similar to what would be the optimal rule, which is, consume everything plus 70% of what's above the threshold.

And the threshold is 100, so for a mean of 100. So what it then shows is that that rule of thumb is very, very similar to the optimal path. So we can simulate that easily. So here's a consumption of someone with this variable-- the income of someone, mean 100, standard deviation-- what did I say? 10. Boom, boom, boom.

So you can see that there are a couple of good realization of income. So the person has been building up some assets. And he's like, she or he is doing-- has lot of assets at some point. And then there are a few bad years. And then he's trying to keep the consumption more or less stable, so the assets get depleted.

At some point they are so low that the next one is another low income. In fact, that one's very low. They're eating everything and their consumption takes a big fall. And then they can start again building the assets.

So that is someone that's-- a conservative saving policy. That's a less conservative saving policy, where they are saving a lower fraction of the excess. And you can see that the consumption is higher on average than in this one. But the asset level is lower but sometimes the consumption takes these huge dips.

So depending on someone's level of risk aversion, they would choose between one of these two processes. They can choose more or less stable with the occasional dip, or they can choose much less stable with more dips, but then they don't have to hold all of this cash that they're actually not consuming. And you won't get any value from the assets. It doesn't make you happy to have all these things in the bank.

So that's the neoclassical model for you. That means that if someone has a very variable income, which is the case of the poor in many developing countries, and if they are risk-averse, then they should really save a whole lot. In fact, even here where they are saving a whole lot, they can't fully eliminate bad times. Sometimes the bad times arrive, and then when they save less, then they have a lot of lean periods.

So risk-averse households should be saving a ton if they don't have access to good credit, which-- you can grant me for a moment. We are going to do credit for a long moment because you're going to see credit later with spend. But basically, it's a fair approximation to say that it's pretty hard for households to rely the credit market, on the well-functioning credit market, to finance consumption when they need it. And, therefore, they should be saving a ton.

Even if they don't have great savings options, even if they don't have a savings account, but particularly if they do-- for example, if they could invest in their business and get some return for it then they should. And in fact, I'll show you in a minute two papers-- one by Pascaline Dupas and Jon Robinson, and one with Simone Schaner, which shows that they are a huge benefit to the business. For people who have small businesses, they are huge benefits to the business of saving a bit.

So they basically have access to-- households, in particular self-employed households, have access to saving opportunities. And even if they don't, they should just save under their mattress or in bullocks or in bad savings opportunity just to avoid these huge dips that are going to happen from time to time if you cannot borrow. So that-- so far, there is no puzzle. That's just a prediction. And where it becomes a puzzle is that they don't.

So after this-- the people that I'm going to show you, where they found that when people in Kenya have access to a savings account, that has great impact on their businesses and on their consumption and everything is great. They replicated this experiment along with Dean Karlan and Diego Ubfal in Uganda, Malawi, and Chile. And they found that basically very few people had any interest in this account.

After all the fees were waived, and people were given this basically free savings account-- so you would think if the problem was saving constraint, then people should get hold of it. And in fact, basically people run this experiment on savings, open savings account to people. And then people say thank you very much and they don't use them at all. So in Uganda, 17%, Malawi, 10%, Chile, 3% use the account.

So these models are not going to-- I'll now elaborate a little bit on this point, but these models-- the basic model do not rationalize that savings behavior because people should be willing to save at negative, even if the returns are negative, just to avoid these big dips in consumption. So before we think about other models, let me just elaborate a bit on these two papers.

So what Pascaline Dupas and Jon Robinson did, is that they opened a free savings account for small business owners at the local bank. So we were discussing about saving in bullocks or illiquid asset. So why would people do that? Well, it's because it's pretty hard for the poor to open savings accounts.

And do you know why it's hard for the poor to open savings-- why is it so expensive? Like $7 to open a savings account for people who make $2 a day. That's a lot of money proportionally.

So why is it so expensive to open a savings account? Why isn't just-- why isn't it just free? After all, they are putting their money in the bank, which the people can then use to invest. Right? Yep?

AUDIENCE: Are banks charging for security?

ESTHER DUFLO: Yes, so they are-- in particular, what's the issue with savings as opposed to credit?

AUDIENCE: Your money's somewhere.

ESTHER DUFLO: Well, yeah. They put your money somewhere. But presumably they're relending it, so it's not security in the form of giving you access to a safe because they are banks. They are just relending it and all that. There is probably some security, but I don't think the issue [INAUDIBLE] because the money is put to use.

In fact, you would think there is not so much capital all moving around. People should likely like to have liquidity. So why is it that they need to charge so much to people to hold the money-- for the privilege of holding the money?

Well, think of your money that you have in your bank account. What feature does it have for being in a bank account or in savings account? What happens if your bank fails? Do you lose the money?

AUDIENCE: It's insured.

ESTHER DUFLO: Yes. No, you don't lose the money we spend. Good news! You do not! It's FDIC insured up to a limit that I think is relevant for-- I hope it's relevanr most of you.

[LAUGHTER]

So it's called FDIC insured, which is-- it's basically-- you cannot lose your savings. The government insures it. Same thing in Kenya, same thing in Uganda. The money that is in savings account and deposit account is insured.

You cannot, as a bank or the government, run the risk that poor people who have put their savings into your bank would lose it. So that's why you were right with security, although that's not like physical theft. But it's the kind of loss-- bankruptcy of the bank or this, that. So the money is safe. It's FDIC insured.

But what this means is that as a bank, you need to do all sorts of paperwork to make sure that you know who has lent to whom. And so there is a whole, know your customer, set of forms that need to be filled regardless of how much money is in the account. So there are big fixed costs of operating a savings account.

So they don't really want your money, in fact, because it's too little money. Compared to the cost of actually creating the account, it is not worth it at all for banks to hold savings of the poor. And because it's not worth it for them, then they don't particularly do it. And then it's not really a product that's offered by the market in the absence of specific intervention by governments, NGOs, or others.

So in this case, what they did is that they just paid the fee-- $7 opening fee. Once you have your money and your basic income-- by the way, the account has no interest and they have a withdrawal fee of $0.50 for transfer below $8, $0.80 between $8 and $15, and 1.5 above. So it's actually pretty expensive to have your money in a savings account there because taking it away is expensive. So it's not necessarily the most attractive account, but very representative of the type of savings account that would be accessible to people.

So we were talking about saving constraints. The first one is just this very basic one. It's just not that easily available. And in fact, at baseline they're working mainly with bicycle taxis called boda boda, and food sellers, market sellers, et cetera. Only 2% of them had an account.

So what they did is that they took a baseline. They opened the account to half of them and then they asked them to fill diaries, that they were periodically checking with them once a week. They were checking their daily logbooks for about three months. And that goes a little bit in trying to understand better their profit and business, you know, going back to the discussion we had about measuring savings.

So here is what they found. So what's the main finding here in a sense? Yep.

AUDIENCE: Do you use it almost immediately? Put whatever you're going to put into it and then--

ESTHER DUFLO: So yeah, that's the number of transactions in the first six months. So this thing is not that you use it immediately. That big massive zero is telling us something else than what you're saying. Telling us that the 55% of people have invested exactly-- have done it zero times-- have never done anything with the account.

So that's-- and then these are the people who used it. And most of them use it very little. You're right, maybe they're putting everything they had in the first six months that applies-- oops. And then some people use it more. So that gives you a bank account usage of about-- first stage of about 40% of people that actually have used it at least once.

So on average, including the zeros, they have more money in the bank. And then they look at whether they have more money in other forms. So one is animals, the famous bullocks or other things. And the other is ROSCA contribution to this Rotating Savings and Credit Association where people save together.

So they don't find a substitution from the ROSCAs. If anything, people are more likely to have animals, perhaps because they accumulated enough money to buy animals. And they have more in the bank. You notice that they don't try to dodge savings. They don't even attempt to measure total savings. So this goes back to your question of, it's hard to measure savings.

And then you can look at business outcomes. And they found that they are able to invest more in the business. Business investment is lumpy, so they try with and without trimming a little bit. But you get the same kind of results.

This is mostly the women as opposed to the men who are cycling driver, maybe because they don't need much of an investment. They just have their cycles. And then when you look at their expenditure, you find that they are able to invest a little more in their daily expenditure.

So this is coming back to this idea that saving allows them to invest in their business, which actually has a positive return. And then they have some more detail down on consumption smoothing, and they show that when people have a shock, for example when they get malaria, they consume much more in periods where they have malaria and they can't work, because they have some money in the savings-- either in the bank or from the fact that they are making a bit more. So it seems to be that their business is actually a profitable saving investment opportunity.

But the argument they are making here is that it's somewhat lumpy. So if you're selling charcoal for example, you can buy one bag of charcoal or two. And then in order to get to the second one, you need to put your money somewhere until you can get the second one. So they are this-- the return of being able to save from time to time in the business itself. And this is on top of the consumption smoothing benefits that we discussed before. So it's suggesting that they have profitable savings opportunities, and they should therefore use them. Yeah?

AUDIENCE: Just a quick clarification, I'm still a little bit confused though because if the accounts are paying no interest and it costs them money to withdraw them, why is anyone finding this-- is the primary value of the account just being able to put your money somewhere rather than being tempted to use it on something? Or, and so that--

ESTHER DUFLO: Exactly, so to that safety plus own internal saving difficulties, which I'll get to in a minute. So in fact, the very feature of this money, that it's hard to withdraw, might be an attractive feature. If you go back to the Callen Zinman-- Callen, Ashraf, Yen that's-- in a sense that's-- this sells a little bit of these features which is, it's going to be a bit costly to withdraw the money. So you are not going to do it until you really need it. It's going to put some constraint on you.

AUDIENCE: [INAUDIBLE]

ESTHER DUFLO: Yep.

AUDIENCE: Does [INAUDIBLE] businesses instead of [INAUDIBLE] change how we should think about the model we went through or anything?

ESTHER DUFLO: Yeah, yeah, yeah. So here what I'm trying to make the point is, that for households that have a small business, they are investing-- saving has potentially higher returns. Now for a family that doesn't have a business, then you might not have that. But in effect, the fact that savings has potentially higher return, compounds the puzzle that was in the first slide, which is the take-up of savings account and any form of savings that we can measure is pretty low.

So that's the point. The point I'm trying to substantiate there is that a lot of these people, a lot of the poor, have some form of small business activities even if they actually are otherwise working. And then if they do, it seems to be from the evidence we had, that it's actually beneficial for them to save in this form, in the form of that.

Then there is another paper by Simone Schaner which-- I'm not going to show you the table in detail-- but it deepens in a sense that [INAUDIBLE] because she gives saving incentives in the form of a high interest rate for a few months. And then she followed them immediately and then 2.5 years later. And what she finds is that there are large benefits to the small business 2.5 years later.

So it seems to be that people who got the initial saving incentives started to save, and then kind of continued to save somehow and then eventually reach benefits for their businesses. It's a bit of a puzzle why. It could be because of some lumpy investments the business is capable of doing. It could be because it creates a saving habit that they are continuing with. We don't quite know.

The reason why I don't want to spend too much time on that is that I want to go back to our core puzzle which is that households should save. They don't save that much. In fact, not only they don't save that much, but they borrow at very high rates-- and when they can borrow. So borrowing is difficult. But they can borrow-- they borrow at very high rates. And that for households who borrow at very high rates, borrowing less is a form of savings.

If we go back to the Euler equation, it tells us that the ratio of marginal utility should be at least beta r, with r is now the growth interest rate because it's beta r when they can and it's greater when they are constrained. Just putting some numbers, it suggests that the ratio with CRRA preferences, it simplifies to this. That means that the consumption should raise really fast for the poorest, because they have interest in savings and it is beneficial for them to save because they are so unhappy being so poor.

So the return to saving in terms of their utility is very, very high because when you are poor, your marginal utility of consumption is super high. So people-- they should save and their consumption should increase fast. And of course, we don't see that

So what could be going on? One experiment is kind of more deepening the puzzle in terms of what could be going on. It's an experiment that was set up originally by Dean Karlan and Sendhil Mullainathan and stayed idle for a long time. And finally, Ben Roth came and helped them take it to the finish line. So it's Karlan, Mullainathan, and Roth.

Very nice paper in the ARInsight. And they provided evidence of that Euler equation puzzle, which is, they look at vegetable vendors. They have a very simple production function. You purchase fruit in the morning, you sell them through the day, and then the next morning you do the same thing again.

So you constantly need working capital for one day because you buy it in the morning and you-- so how do they finance it? Do you know? How do they purchase the vegetables in the morning? Yeah?

AUDIENCE: They take out a pretty high interest daily loan.

ESTHER DUFLO: They take out a daily loan. So what they do is that they purchase it from the vendor who gives from the wholesale grocer who gives them a loan. And these loans are 5% a day. So basically people pay 5% a day. So the question is, why are they borrowing at this rate? Why are they not saving it?

In other words, they have access to high return investment, which is saving a little bit. 5% a day is really a lot. You can write it in an Excel sheet and convince yourself that after 30 days at 5% a day, you'd be quite richer. So imagine for example, that you drink one less cup of tea every day. So you are saving some small, small amount of money and you reinvest back in the business such that you have to borrow a little less.

In 30 days you would have doubled your income by not having to borrow. And then you can-- you know, every time, then you can take a little bit once you are fully paid off, when you don't need any money to pay off what you take normally, you can take a bit more. So why don't they? So one possibility is that the true cost of the loan is different.

This is how the moneylender slash wholesaler is working with them. Maybe you're not allowed to borrow a little less. You have to do everything on credit or everything without credit. So there is some fixed thing.

Maybe you can't cut your consumption. Maybe you're extraordinarily impatient. Maybe you don't understand compounding. Maybe you don't have access to a good savings opportunity. Or maybe the preferences are not what we thought and there is something that makes them more impatient today than in the future, either in the household family dynamics or self-control issues.

So we now know to learn a bit about that. There are two treatments. One, where they just paid the entire loan, so that eliminates a lot of the issues about the fact that the loan would have to be fixed. And so they just paid for 3,000 rupees worth of vegetables.

And then the second, in order to make sure that this is not about misunderstanding, is that they give a crash course in financial education. So they teach these vegetable sellers-- have the class where they worked, how much they spend in total, show them the benefit of cutting down, discussed compound interest rate, discussed what they could have done with the money and brainstorm on whether they could cut down consumption. So the training is to pick up the effect of potentially not knowing.

If you can't cut back consumption because of the [INAUDIBLE] utility function, then the one-time buyout should put you on a path to save. And even if you can't save, you should be able to stay where you are unless you get hit by a shock, in which case, you restart from scratch. If they are impatient or if they have self-control issues, then you should pull back fast. So that's how we can interpret these two treatments.

They did this in India and the Philippines, as I said. In Philippines, they have the survey every two weeks, six weeks, 10 weeks. And in India, three months, six months, 12 months. And just to show you results for India, there is no impact of financial education. So I didn't even put it on the slides because it's a big table.

For the treatment of buyout-- and so the treatment of buyout only and buyout plus financial education is the same and financial education has no effect. So these are a few-- the first few numbers with buyout only, where you can see that at 2-4 months, they are less likely to have money than they are debt. At five to eight months, that's gone, and it's then [INAUDIBLE] gone. So originally, of course, none of them has debt. Two to four months later, they still don't have debt but then it goes away.

So the fallout is relatively quick. It's not instant, as in, they immediately blow it in a big party-- the 3,000 rupees they were given for the day. But it's frittered away fairly quickly. So this is the main impact. And that means that originally-- yeah. So that's the first interesting fact that they seem to-- in a sense, they are sitting under a money tree just by-- which is, stop borrowing, and they don't take advantage of the fruit from that money tree.

There are many other examples-- a bunch of nice papers that try to-- that makes the same effect. For example, John Robinson and Laurie Beaman have a paper where they simply observe that often people in shops don't have a cell phone account. And a cell phone account is what people buy to get prepaid money on your phone.

It doesn't spoil. I don't think it can even be stolen because if they were stolen, you could call the office and say, the card got stolen. It doesn't take any space. And people keep coming to you and buying them in small increments since they don't have any savings. They want to-- so it's-- and on each of them you make a small amount of money.

So you think you would always want to have this cell phone card. And yet very frequently, they just don't have them-- the little shops, so they miss out on the sale and someone else is making the sale. So that's another kind of missed opportunity that comes from just not having the liquidity and not having the ability to save in money cards. Saving in money cards seems like a perfect investment-- a perfect way to save for a shop owner.

Another example which we'll come back to at the very end of the lecture is-- Michael Kremer, Jon Robinson, and myself have a series of paper on fertilizer, and asking why people aren't investing in fertilizer-- in Kenya again. And despite the fact that-- our first initial hypothesis that maybe it's not profitable. But we conducted a series of experiments showing that it appears to be quite profitable.

So despite the fact that fertilizer is profitable, people don't seem to be saving in the form of fertilizer-- which is basically buying fertilizer when it's time to plant it. Then it's in the ground for a while, but then you have much more fertilizer. And people aren't able to-- so what they say is that they don't have the money when it comes time to planting-- which basically what it means, they haven't accumulated enough to save.

AUDIENCE: I was going to ask why people are not doing it either.

ESTHER DUFLO: Yeah, that's what they say is-- and we finally were frustrated-- that doesn't mean anything! How can you not have the money? It's divisible. So if you have $0.03, just buy $0.03 worth of fertilizer. But it's like, oh, we have to go to the place, to the shop, and we have to buy the fertilizer. It's not worth it if you're not going to do a little bit more. So I don't know whether I agree with that but it's-- basically they say that because at the time of planting, they don't have the money. And--

AUDIENCE: [INAUDIBLE] options that are available [INAUDIBLE]?

ESTHER DUFLO: Exactly, so let me go back to that. That's kind of what we introduced. But that seems that basically at harvest, they have lots of maize, which is basically lots of money because the maize gets sold for money and they buy and sell maize. So at harvest they have lots of money, but it's not the time to put the fertilizer in.

At that time they don't go to buy the fertilizer. Then a few, three months later, it's the time to put the fertilizer in the ground but they don't have money anymore. They feel poor. They are not sure they are going to be able to stretch their means to the end of the season. So they don't want to put money into fertilizer.

So that's a puzzle, that's another puzzle. So let's hold that puzzle in mind for now but that's another example of a missed opportunity. People are not able to save enough to buy a bag of fertilizer.

So how can we explain that? Well, one way that we have already circled around and hinted in the whole class, is the idea that the discount rate is always really low. Or in fact, people have individual inconsistent time preferences. So the discount rate might be different for the present and for the future.

So you can-- so the standard way that this is modeled, and I think Frank introduced it briefly when he did his lecture, is the idea of hyperbolic time discounting where the entire future is discounted more viewed from today. Then tomorrow is discounted, then the day after tomorrow is discounted with respect to tomorrow. So that's the hyperbolic discounting model.

More generally, again as you saw in Frank's lecture, is [INAUDIBLE] good money. People seem to think differently about the future than about today. So for example, there is an experiment where people are given an opportunity to choose a movie video from a set of 24 titles, including, for example, *Four Weddings and a Funeral.*

That dates the study. I don't know-- probably you weren't born when *Four Weddings and a Funeral* was shown. But it's fun and it's lowbrow. And you also had the *Schindler's List,* which is not lowbrow.

And the point is that when choosing for today, 56% choose lowbrow-- so *Four Weddings and a Funeral* versus the *Schindler's List.* But for choosing for next Monday, 37% choose the lowbrow. And for choosing for the Monday even in two weeks, 29% choose the lowbrow. So that shows that-- so it could be put in the inconsistent-- it's certainly inconsistent time preferences in terms of good. You seem to have your different preferences for today, then for next Monday, or for the Monday after.

It could be couched in hyperbolic discounting. But more generally we could write it and say, well, there are only two types of things that we-- two type of goods. There are the goods that we really enjoy now and we enjoy in the future, and there are the goods that we only enjoy-- we enjoy if we consume them today but we don't look forward to consuming them in the future. So with watching *Four Weddings and a Funeral*-- so let me first take the example of a big, sweet, double-glazed Krispy Kreme doughnut.

[CHUCKLING]

If it's in front of you in the shop, you said, oh, I would really like-- I really want it now. And you'd feel quite happy eating it. But if you think about yourself tomorrow and say, would I-- with everything taken into account, would I really like my future self to eat that thing? It's like, not really. So the doughnut is an example of a temptation good, which is, would enjoy it if you eat now but you don't derive any joy from thinking about yourself eating it in the future. In fact, you might see the opposite.

You might become-- similarly spending time on Facebook. There is a nice paper by Matt Gentzkow and Hunt Allcott called "Digital Addiction", showing that basically he paid people to give up on the internet. And showing that people when they think about the Facebook, Twitter and the like-- Instagram-- they would like to consume less of it in the future. So they do not enjoy thinking about themselves doing it in the future.

But on the other hand, today they don't want to stop. You would need to pay them a lot to stop today. So that's another example of a temptation good. And similarly of course, with the example of *Four Weddings and a Funeral* or *Schindler's List.* You would really-- it would be very fun to watch it now, but in the future, you want your future self to watch the *Schindler's List* because that's more educational. So that's the point.

So within-- so there are two periods that-- and within each period there are potentially infinite number of periods. But even two is sufficient to start creating this-- to create that inconsistency. And within each period, you maximize your utility over the non-temptation good and the temptation good. You could separate the goods into two, and that's kind of a reduced form for all of the good in that temptation and all the good that are not temptation.

And then the intertemporal preferences maximize today's, plus they discounted some of the non-temptation good subject to your usual asset accumulation constraint. So as soon as you have two periods, you have a commitment problem, which if-- it's actually not the case in hyperbolic discounting. You need [INAUDIBLE] to make one appear. Here you have it initially because today's self wants-- even with only two periods, today's self wants X and Z but for tomorrow they only want X.

This next hyperbolic model-- if the consumption is CRRA with the same discount rate, you can see that you can rewrite exactly-- you can rewrite this model in a way that will have a higher discount rate for today and for the future. But it's actually much easier to work with because the key result in the paper is that basically it's like the Euler equation.

When you maximize this problem, you can maximize-- you can choose how much of x and z is going to be consumed in period one. And what you're going to get for the consumption of x is like a Euler equation, except there is that term, which is a consumption tax, which is in the future, in the temptation tax or in temptation tax. In the future, if you save some money to be consumed in the future, in the future some of it is going to be devoted to the z good, which you don't like.

So this is now z of-- and that gives you a z of c. So this z is a new function that you haven't seen yet. But what they show is that basically it is just an invert of the x of c of the u of x. This is an inverse of the consumption. It's an inverse of the marginality of consumption with respect to x and z.

So you can easily rewrite the problem. Actually, the proof text-- one paragraph and I have it in the next slide. You can just have a look at the paper. But it's very intuitive. But the result is very intuitive.

It's saying that basically the person who is subject to temptation is going to maximize their utility. That's going to give you pretty much the same result as the standard Euler equation except there is going to be this temptation text. Already that with respect to the z function is increasing. So its z prime is positive. So 1 minus z prime is less than 1. Yeah?

AUDIENCE: Does that mean people in this model-- people are sophisticated about [INAUDIBLE]?

ESTHER DUFLO: Yes, people are completely sophisticated in this model. And so what does that mean? That means they understand-- just to-- sophisticated means they understand exactly what model they are talking about. So for example from this model, you're going to get the demand for commitment that you're finding to some extent in Callen, Ashraf, and Yen paper, which is, they know what's going on. They want to protect their consumption against their future self.

So that's the first result you're getting from here, is a demand for commitment. You get even more interesting results in the case where the temptations are declining. That is, the z prime function is decreasing in c. So the temptation tax is decreasing in c. What does that-- intuitively, what's the intuition for decreasing of our temptation tax that declines with your consumption levels? What does this reflect? Evisha.

AUDIENCE: Is it concave?

ESTHER DUFLO: Yes.

AUDIENCE: So I guess like, the first chocolate bar tastes really good. But then the second, third-- [INAUDIBLE]

ESTHER DUFLO: Exactly. It's saying that the shape of temptation is exactly-- it can be mapped to the shape of temptation. The utility function from the temptation good is concave. That means that temptations are larger as a fraction of your consumption, when you're poor than when you're rich.

So for example-- so if temptations are mostly visceral goods like chocolate, Krispy Kreme, et cetera, that makes a lot of sense because in proportion of your income, there is only so many Krispy Kreme you might want to eat. So even if the number of Krispy Kreme increases, it increases less fast than the amount of money you have to give to that. So the share of the consumption good as a function of your budget is going to decrease.

So that's what it reflects. It doesn't have to be true. For example, you could say, well no. For a rich person, the temptation good is going to be a yacht and a yacht is super expensive. So you graduate from Krispy Kreme to a home TV to a yacht. And finally this is a non-decreasing or even an increasing shape.

Note, for example, that the hyperbolic discounting function is with CRRA, with the same coefficient. In that case, it's non-decreasing because the temptation always remains a fraction of your income. So the hyperbolic discounting is actually a case-- can be mapped into this model with a constant with non-decreasing temptations. Yeah.

AUDIENCE: In [INAUDIBLE] this models predict that poor people or people with less consumption have more temptation for [INAUDIBLE]-- say you took--

ESTHER DUFLO: Yep.

AUDIENCE: --it's [INAUDIBLE] because as model predicts something, there is this temptation. But also there's the other way around, which is that there are different people, say, more patient and less patient people. Less patient people would end up poorer.

ESTHER DUFLO: Yeah, yeah, of course. So you're right, which is-- in the data, the model would predict that. But that would also be predicted by heterogeneity. So what you propose as a prediction is not a testable implication. I mean, it is an implication but it's not going to make the difference between this model or another model based on basic heterogeneity. Absolutely.

I'll give you another example where-- so we should think about what this model predicts and then we should think about other tests for whether or not we seem to have temptation in the world. And that's related to some of this implication. Could be related also-- we could devise consumption-- we could devise tests based on the shape of consumption function for some good. Before I go to the implication, think of another example.

So one form of temptation goods that would satisfy this declining temptation tax is these visceral things. Another is being taxed by your relatives because you could think of this-- map that into a household model where part of your temptation is, someone will come and ask you for something. So your child wants a toy or your cousin wants some money because they are starving.

So the declining temptation would be where, if it's easier to say no to your cousin who wants $1,000 to start a business when you're quite rich, than to say no to your cousin who wants a bowl of rice when they are quite poor and you're quite poor. So this is-- so I don't know whether that's the case or not, but that would be that.

That would be that-- proportionally, you have to give a proportion of your income to your relatives-- even though you have to do it at the time, but you don't enjoy it. That would be a form of temptation good. So maybe that's easier for the rich to not give in proportion of their income. Or maybe not, but that's the example.

So assuming that you do have this declining temptation, then you have extra implications that are interesting. The first one is, the precautionary saving argument might not take place because for someone who is very poor, if they are thinking they are also going to be very poor in the future, and therefore most of their savings is going to be eaten by themselves in the form of things that they don't want. So the consumption tax is going to be high.

So someone-- and then in fact, when there is more variability of income and when they are particularly poor as well, when their future self is going to extract all of this consumption in the form of z, so they might say, I might as well eat it now. If I'm going-- in other words, if I'm going to eat some Krispy Kreme tomorrow, which I don't like the thought of, I might as well eat it now where I'm going to enjoy it. And therefore you might not have the precautionary saving motive, and therefore you might have the poor saving less than the rich.

And you also have-- and what is interesting is the perspective that you have for the future. So in a precautionary saving model, someone who anticipates getting richer can consume more today because they will be richer in the future. So they don't need that saving. Someone who anticipates that times might become bad should save a lot. That's the third derivative of the utility function being positive.

Now if you think that in the future, you're going to be richer, therefore the temptation tax-- in a fraction of what's available, the temptation tax is going to be lower. And therefore you can afford to save more today because you think the time will be better. So you have the opposite of the precautionary motive, where the hope for the future might actually helped people getting started in savings. So that's the first implication of that model with decreasing temptation. That doesn't work with non-decreasing temptation.

Another one is that you have a poverty trap even if there is no-- you could have a poverty trap even if there is no other source of discontinuity. Why? Because the poor people won't save because they know that they are again poor tomorrow. And any savings is going to be spent in the z because it's a large factor. z represents a large fraction of their income.

If you're sufficiently rich, it is worth saving because you know that you are not extracting such a large fraction of those savings in the form of z that you don't enjoy. And so you're going to save more. And so the richer people are becoming richer and richer because they are saving, saving, saving more. And the poor people are becoming poorer because they are saving, saving, saving less. So the poverty itself is the source of what makes you poorer, which creates the potential for a poverty trap in the absence of any other discontinuity.

It also says that people might prefer to have-- might not find it-- going back to my fertilizer thing when I was telling them, but you could buy just so little. Why do you-- even if you have a tiny bit of money, buy a tiny bit of fertilizer. Well, a tiny project is not making you rich enough that it's already changing your temptation tax. So most of that extra money you would make from the fertilizer on a small scale, is going to be frittered away on things that you don't care for today, so what's the point?

So that might explain why, even if there is no fixed cost in going to the shop, there is no reason to do fertilizer unless going to do enough because the extra money that is going to be generated from the fertilizer won't work. And you might choose a microfinance loan at a high interest rate in order to put money in-- rather than a credit card loan because you want a large project.

You might like this kind of funny account where you have to pay for withdrawing mone because since the amount you're paying is smaller when the amount to withdraw is smaller, you won't be able to access it. It's not going to be worth it accessing until there is enough. And once there is enough in one go, then you won't want to fritter the whole thing. That-- it won't make-- if you were saving and you were able to withdraw $2 to pay your Krispy Kreme, you might. But if it's expensive to do so, you might wait till you have $50, which is where the withdrawal costs are smaller, and then you can do something more meaningful with the money that you value from the point of view of x.

Very quickly, you read the paper and it's pretty self-explanatory. But it's nice in the context of this model. This is the first paper to set up an experiment for demand for commitment. It's a very imperfect paper, as many of you have pointed out, starting with the fact that the results are pretty mushy because the saving marketing is, in fact, almost as large the impact as the demand for commitment thing.

And the other thing is, they are trying to test whether people who have hyperbolic dis-- who look at hyperbolic discounting in fact are more likely to take up. But in their graphs, there are a lot of people who seem to be confused. So you would test hyperbolic discounting by saying, if you are more patient in the present than in the future. But a lot of people are not on the diagonal on both ends, and that's not predicted by any model.

So that's confusion. But then if that's confusion, then maybe that's also confusion. Since then, since that paper, on the ability to measure-- with the work of Charlie Springer in particular, we are much better at measuring demand for commitment. But this is how science progresses. This paper is a wonderfully creative thing of going ahead and creating the thing.

And then you've seen Frank's paper building on it with his alcohol paper. If you're interested in what I think is like the state of the art in terms of a commitment savings product, it's the paper by-- an Afghanistan savings account by Michael Callen. It is kind of the 2.0 of the Ashraf, Callen, Yen using the work of a big firm in Afghanistan. It's excellent.

In fact, it's Callen and I think it has the son of Ashraf Ghani on it as well. So-- who is actually a economics professor in business school somewhere in the US. So it's interesting also for that purpose. And I suggest that you read it. There is also a lot more on savings now, which is inspired by-- on savings paper, more recent empirical literature inspired by both the combination of the type of models that we discussed today and this original paper-- which maybe we'll have a chance to see in recitation. Great. Thank you so much.