[SQUEAKING][RUSTLING][CLICKING]

ROBERT TOWNSEND: OK. So let me update where we are in terms of the calendar. I guess a notable thing is that we have reached lecture 10 today, Thursday, in terms of the reading list. We are-- having done all the risk-sharing. Yeah, we're at lecture 10.

So there's a couple of things that are starred here, and they obviously both come up in the lecture today. One is the early version of this paper with Alvarez and Pawasutipaisit sit cash, which has a star on it. It's a little bit rough because it's kind of, even then, in the process of revision. And I have added new material in the slides that haven't even made their way into the working paper yet.

And then there's Samphantharak and Townsend, which is describing the tools we're going to be using a lot today-- how to organize data, and think about flows and stocks in terms of financial accounts and so on. So, again, the two starred readings.

Other things you might enjoy, time permitting, there's something on Kenya, M-Pesa. And I have one slide describing-- actually two slides describing that today; something on lifecycle planning. But none of those things have stars on them. OK. So that's the reading list.

The last thing, turning to the study guide or the lecture from last time, which was "Risk and Return in Village Thailand." So explain intuitively in words why multiple technologies can be chosen to be operated simultaneously, and why the rate of return on some might be higher than the rate of return for the others that are also operating. And for that matter, why some projects are not chosen at all. If you can, try to use the language of risk premium, but don't use it as jargon. Try to imagine you're telling someone in your household or a friend kind of what we're studying in class.

So I've been calling on people. Let me just ask for volunteers today. Of course, I can call on people. Any volunteers? Carrie, you want to take a crack at this?

AUDIENCE: Is it because some technologies have a lower rate of return because they're riskier? And the probability of a high rate of return is lower for them?

ROBERT TOWNSEND: That's true. At first blush, the rate of return is simply all the possible returns multiplied by the respective probability. So that's kind of inherent in the technology itself. But if that were the only thing going on, then why would they not choose the one with the highest expected return?

AUDIENCE: Can I take a crack at it?

ROBERT TOWNSEND: Yep.

AUDIENCE: Well, I think they've got-- yeah, I'm a bit behind in the readings. But what it looks like here is that you have some technologies that have a higher expected rate of return but more risk, and some technologies that have the low rate of return and lower risk.

And basically-- and some of them-- so some of them are-- some people-- it's like some households might be more like-- it's like, you can have like-- and if these aren't perfectly correlated, then you might be able to get a better-- sort of a less risky rate for the same-- less risk for the same rate of return by diversification.

And the reason that you would have lower-- you would have-- you would have lower rate of return technologies in the same-- used at the same time as higher rate of return technologies, as the lower rate of return points, have a-- the lower rate of return points have a lower risk, and the higher rate of return points, you have to basically pay a risk-- there's risk premium there, which is why it's less desirable or something like that.

ROBERT TOWNSEND: Yep. OK. That's pretty good. Thank you. Yeah. What you're describing is correct. There's like a mean variability trade-off here, a downward-sloping frontier, if you want to think about it that way. As you move across different projects they may have higher and higher means, but they come with higher and higher variance. So there's a trade-off. And since households are risk-averse, they don't like the variance. So they pay attention to that.

And you said this also. You could have two technologies with the same mean return-- one with a higher variance than the other one, and they would never do the higher variability one because then there's no trade-off. Since they get the same mean but higher risk, they wouldn't choose that.

There is a diversification aspect that we featured quite a bit in terms of land plots and how they want to diversify, or in village India over crops and income sources and so on. And that's kind of captured here in a new way, but it is in here. Namely, a guess of what technologies to operate overall, you can compute the variability. And then you look on the margin about whether to add or subtract one more. And you look at its variability relative to this village average.

So if it's going to add to diversification, it's going to have a return stream which is low when-- well, say high when the rest of the village is low. So that's kind of the way the diversification aspect is coming out in this mean variance framework. But great. That's very helpful.

All right. So here's another one. Explain how it can be that while the magnitude of idiosyncratic risk is higher than the magnitude of aggregate risk, nevertheless the idiosyncratic risk premium is smaller than the aggregate risk premium. So, again, I'll take volunteers.

AUDIENCE: Yeah. So is it-- so I think it's because when it comes with idiosyncratic risk, it's possible for the village to distribute the returns amongst each other. It's like ex-post. Whereas, for aggregate risk, it's like is risk affecting the whole village. And so there's not-- and so redistribution can't really benefit anyone when it comes to things that affect the village as a whole.

ROBERT TOWNSEND: Yep. Perfect. Yeah, the last part, which I didn't read out, explain in words what gifts and loans have to do with it, but you kind of did already. That's the way they pool away most, if not all, of the idiosyncratic risk. Whereas, in the aggregate, it's common. It's hitting everyone. So essentially, you can't pool it away.

And this next part is related to that. Is there full insurance in these Thai village economies or not? And to give it away a bit, not. And, in particular, what is it in the data that is inconsistent with full risk-sharing? Volunteers?

AUDIENCE: We find that in the data household consumption is still-- is positively correlated with the idiosyncratic shocks. So that shows that it's not full risk-sharing. But there's still some risk-sharing because the coefficient is really, really small compared to other coefficients.

ROBERT TOWNSEND: Right. Compared to the coefficients on gifts and lending, which absorbs the bulk of the risk. So that's a perfect answer. I am so happy with this class. I enjoy teaching you tremendously. It's very gratifying. OK. Thank you. There's one more-- see if we can get it in.

Describe one example from Thai villages of propagation through networks that would be inconsistent with full risk-sharing and investment efficiency. And I'm abbreviating. So this happened at the end of the lecture, just as a hint. Can someone described what we did at the end of the lecture?

Well, it's kind of given away here in this bullet point, but I'll read it out anyway. How can an individual specific illness shock turn into an aggregate village-level downturn? Volunteers?

AUDIENCE: I think at the end of the class, we were talking about how some of the networks-- how different kinds of networks are different from each other. I'm not sure which specific networks were being compared, but yeah.

ROBERT TOWNSEND: That's true. So what happens, do you recall, when a household gets hit with an illness shock that increases its expenditure?

AUDIENCE: Right.

ROBERT TOWNSEND: Did nothing happen? Something? You remember, or anyone else remember?

AUDIENCE: I remember that the households who are directly connected to the household who got illness shocks would [INAUDIBLE] down their working opportunity much more than households who are unconnected to those illness households. So through the network, households who are connected, who are related to each other much more maybe shock-- there was an actual shock between them. So sudden illness shocks to these households may transfer to the aggregate village through the [INAUDIBLE].

ROBERT TOWNSEND: OK. That's a pretty good. And it's a bit tricky, this question. Because there are different networks in the village. There's kind of the risk-sharing insurance network. There's the labor network. There's the supply chain production network. And they don't coincide.

So if you got hit with an illness shock, but nevertheless, you're in this risk-sharing insurance network, then part of the-- or almost all of the shock gets covered, in come the gifts, just like we were describing with the production shock.

However, when that illness shock is not covered by gifts, then the household has to respond on its own in some way. And it does this by cutting the number of laborers that it hires, and by shrinking-- saving costs by shrinking its production operations. And that is what spills over to other households along the production network.

So networks could be good or bad. Networks are good for sharing risk, but they also can be bad because they propagate contagion. So that was--

AUDIENCE: So if there was a full risk-sharing, then the-- then there were no effect on the employment for the household who did not get any shock? Network-- I mean, network-- that was a factor. Those houses who did not get any shock because they are more related to the households with illness. But if there was full risk-sharing, those households will not have any effect?

ROBERT TOWNSEND: That's right. So the emphasis here is on a very specific household idiosyncratic shock. And this-- I drew an analogy with COVID-19, but it's a bit unfair. Because COVID, as we know, has a contagion effect literally in the health sense. And if one person gets it, people around that person get it.

So here, though, it's a different kind of illness. But it-- and it could be just an elderly person in the household. But still, there are big expenses-- double the normal expenses. So they don't have the budget to manage, and they have to start either getting a gift from someone or cutting.

So when they start cutting expenses, that's when it spills over to other people that they hire or people. So this is very reminiscent, as I said in class, of what's going on in the supply chain networks in the US, for example, and many other countries. Although, that's more of a different kind of health shock.

OK. Any questions about-- any other questions about the lecture? OK. So then we come to the lecture for today-- ledgers and management. In fact, there's a lot in the lecture today. I would say I'm going to give you a tool to think about, namely, financial accounts-- income and balance sheet accounts, in particular. And then we'll go to the Thai monthly data that we were just talking about, and use the tool to think about planning-- about wealth planning.

Then we'll take a step back and add another account. I'll motivate that account by thinking about Bitcoin, as I promised in lecture one, and link the ledgers of Bitcoin to the statement of cash flow. And hence, back to the Thai villages. In particular, I'll talk about distributed ledgers, which is a big part of Bitcoin-- Bitcoin, blockchain, and distributed ledgers.

And then having introduced that concept, I'll focus on the use almost exclusively of paper currency in Thailand, the use of almost exclusively electronic transfers in Sweden, and something in between in Kenya, which has two different kinds of monies, essentially-- privately issued e-money from a cell phone provider, and a public Fiat money, Kenyan shillings.

And then I'll end-- and so these three things for Thailand, Sweden, and Kenya are about cash-- various concepts of cash and cash management. And then we'll come back to the ledgers again and talk about trusted third parties, some problems and management that may or may not show up at the individual level, but definitely show up at the wholesale level at the level of financial infrastructure. OK.

So here's a statement of income. And forgive me if this is a review for some of you. These things aren't hard, but they're a bit tedious. So we talk about sources of income and uses of income.

And the sources of income, as in the Thai data, come from cultivation, livestock, fish and shrimp. This is a wage-earning household primarily, a little bit of livestock. So the line entry here is almost entirely from wage, labor, income. Those are the sources.

And then the uses, so to speak, or you could say revenues versus costs-- you have expenses from each of the activities-- livestock. If they're in debt, they have interest expenses. There's that depreciation that we talked about last time.

And then this is charges against revenue. So in addition to having expenses for production, they might pay taxes-- these guys don't. But they can either consume the rest or save.

Advantage of the financial accounts is that everything balances. So in this case, revenue is totally balanced by charges against revenue. And if you're-- as we did-- taking the Thai data and thinking of all the line items then creating the accounts, and if these left- and right-hand sides don't balance, we know we something wrong. And we have to go back and iterate and figure out-- excluded a variable, or put it in the inappropriate spot. So it's a great discipline device.

The balance sheet is stocks of stuff at a point in time, namely, liabilities and assets. So the obvious liability would be debt, which this household has. This is really somebody's numbers. This is a real person, I want to remind you. More on that momentarily.

And the debt's about 37,000 baht. And the assets that the household has could be real or financial. The financial assets are a big chunk of it in paper currency and some deposits at the bank. And real assets, in this case, household fixed assets.

So we went through some examples last time when we talked about rates of return, because we had in the numerator profits, which is on the previous account revenue minus expenses. And here, profits over capital. And here we have the real and also financial capital.

Now, again, these things balance. But here, there's a big item that I deliberately didn't mention at first to get your attention, which is if you take assets minus liabilities, hopefully net worth is positive. If liabilities were greater than assets, this guy's bankrupt, at least in the accounting sense.

So for almost all these households, assets are greater than liabilities. And the difference is something we call equity or net worth. And that enters oddly on the right-hand side with liabilities, precisely so that the two sides of the financial account, or the balance sheet account, are equal to each other. OK.

So now with those concepts, the statement and balance sheet, would take it to an application. So this is a wealth planner, which we implemented to that household and all of the other ones. What is the household we imagine to be doing? Namely, maximize discounted expected utility.

So you've kind of seen versions of this before. You saw it first with respect to storage of seed and so on in the medieval village economy. We've got a finite horizon here that happens, but more on that momentarily. So it's discounted utility. This is constant relative risk aversion, with a risk averse coefficient of kappa. Delta is the discount rate.

And the end of life, which is T, they may bequest assets to their kids. It looks like a similar function, but there's an extra term here, which is beta, almost for bequest. And again, the notation changes from one lecture to the next. But anyway, so that's another parameter. We have delta, kappa, and beta. We're going to have to figure out what they are because they will impact the solution.

Subject to the big constraint-- now, again, you've seen a version of this in the medieval village economy because we had resources available today, which they could eat or store in two different ways, which would determine resources available next period.

Here, they have wealth today, and they have resources available today, which is the beginning of period stocks, their wealth, plus income-- more on that momentarily. And they can either eat it-- and the difference between what they have and what they eat, they invest. So some fraction, Z, goes into risky assets, and the residual 1 minus Z fraction goes into safe assets. So this is a lifecycle portfolio allocation problem.

A little bit more detail about these numbers-- so the income term is the second and third term. If gamma is equal to 1, then the mu, which is the mean income, cancels out. And this household is entirely on its own with income fluctuations around the mean. When gamma is equal to 0, this whole term is gone, and the household has its mean income. So gamma here, between 0 and 1, is capturing the risk-sharing that we know is prevalent in the Thai villages.

So that's the way all that risk-sharing stuff entered into the individual program. It looks like self insurance, but it isn't. Because the income stream can be attenuated due to the gifts. But it can also vary over households. We have to elicit from the households how big or small gamma is.

This is a premium if the-- sorry-- an indemnity if the household head dies, all these village villagers get together and contribute something to help the bereaved household with expenses and so on. c is consumption. Big C is committed consumption. So these are things the household tells us they're going to do no matter what. And so we enter as a particular expense. All right.

So here is a bit of what we show the household. We explain to them-- we go and interview them. And that we have been getting data from them all along, but this is extra.

So we tell them we want to know how much they're planning to give to their kids, if anything. Try to elicit their current income profile and future movements and income. Get their goals, and get their current wealth.

And then we're going to tell them how much they should be saving, and how to save it in risky versus safe things. And particular, whether they might want to be borrowing. I don't know. Some people don't like these little cartoon characters. It's an attempt to communicate clearly.

So we don't say, tell us everything. We say, let us demonstrate to you how this works. We'll show you the results of another villager. We asked that household whether or not they want to leave money to their kids. And this household said, no. My kids are going to be much richer than I am. They're going to work in Bangkok. I don't think I need to leave anything to my children after my death. So this household said zero baht.

I made the mistake once of having this interview with the mother and the daughter present in the room at the same time. Obviously, that was awkward. [LAUGHS] OK.

Collect inputs-- so we need to calibrate this income process. And we give, for example, Chen. Now you're going to see some Thai. Sorry. Anyway, wage earners earn a regular income of 10,000 baht. Expects to work to age 65, and then income drops to something very small, which is a pension. So that's this picture. 10,000 baht, age 65, the household's currently young, actually. Drops to zero.

Or this guy-- different example-- has a little shop as a business-- very volatile income. Expects to keep its business running until age 60-- currently age 25. And then receive some small income, again, from the kids. Oh, yeah. Kids give money to the parents, you guys. You hear that? [LAUGHS]

OK. Depends on where you're from whether you can relate to this.

Another example-- this household has regular increasing income for 30 years. But, again, income starts to drop. However, they're going to go into a second occupation, and income only slowly.

So I stay-- we used to stay in the village with the headman. But eventually, Thailand has motels. And you'll often see an elderly couple running the motel. That's now their new business, relative to what they used to do, et cetera. So these are just three examples.

So what about goals? So do you expect to send your kids to college? Is that a planned expense? Or in the future open a business or buy a new TV? But what we want are expenses they're going to be no matter what. So yes or no, depending.

In this particular example household, we give them car purchase. Was that obligatory? No. Send the kids to college? Yes. How many years away? 15 years-- kids are young. And it's going to cost $10,000 for three years-- 30,000 baht. That's about-- it's a lot less than US expenses. It's not an MIT tuition. All right.

So now we have wealth, so w. This is tricky. Money available to save or spend or invest. So if you have a car that's worth 10,000 baht. You think-- it's used, but you could sell it easily. Then we would count that as a household asset or business asset-- could be sold.

If they own land worth 100,000 baht, they may never plan to sell it, and may not even be productive land. So we don't count it unless they said it was. If it's productive, we want to avoid double counting.

So here's the thing. You would count the asset, or you would count the return stream. You cannot count both. The value of land should be the discounted capitalized value of net profits from farming. That's how much the land is worth. You don't want to count both the land as an asset, and also ask the households about the income stream.

So generally, we don't count assets that we are using in calculating the income stream unless they say, I'm going to stop farming at a certain date. And how much would the land be worth then? If we count income up to that point, then we count the land.

So let's run the planner. I'll give you examples. We created an Excel spreadsheet with the answers to all the household's questions, and then enter it into the computer and run it. I'm not showing you all of them. The second household-- we told this household to cut its expenses on consumption quite a bit. They weren't saving enough, and to put almost all of it in safe assets.

More interesting is the life cycle picture. So here's what the planner recommends that the household should do. Namely, let's start down here. We've got income in red. That's the elicited profile over time that I just told you about. The consumption is endogenous variable, obviously. It's in black.

So here, we told the household to cut its consumption expenditure a lot and much lower than its current income in order to build up the assets. So here's the total level of assets. And then when income drops to little or nothing, they eat off the assets. So this is just saving for retirement. Most of it's in safe assets, not so much in risky, and there's no borrowing.

Another household basically had committed consumption that is going to get it in big trouble, even though they said they were going to do it no matter what. When we ran the planner at the time that they have this high expense-- I think it was for starting a business-- they really didn't have the resources to cover it.

Now, they try to increase their savings with income over consumption. But in this median simulation scenario, it's just not enough. So when that big expense comes, on average, they use up all of their assets. Consumption isn't dropping to zero. You remember, it's constant relative risk aversion. So marginal utility of consumption at zero is infinite.

Although, you can't see this, but we gave them some subsistence bundle. So they don't really go bankrupt. They survive. Well, they don't die when they go bankrupt. They continue to survive because of that subsistence bundle.

What else is going on here? They kind of know in advance-- the planner knows-- state by state, year by year that they're in trouble. So actually, they start borrowing.

Of course, given our subsistence consumption rule, if the median scenario occurs, they're not going to pay back the loan. They get to default on the loan, and they know that. Call this gambling for redemption. They're starting to do risky things too, and borrowing, leveraging up, with the idea that there's a good chance they're not going to make it.

Here, the planner is not the bank. Whether the bank would make this loan is another issue. Anyway, they survive this episode and continue to live another day with their fluctuating income. So the rest of this looks like the other diagram to some extent. OK.

So we show the households these pictures, and try to explain to them what's going on. And there are all kinds of households in the data. And some of them, in fact, are very conservative, and they have fluctuations in income and consumption. Consumption is never more than the lowest possible income draw, and mostly, a lot lower than their income draw.

So they're saving like crazy, and their net worth is growing at 20% per year. Sounds good to me, by the way. But they're probably over-saving. But we're very conservative about what we tell them. We don't tell them that that's bad. That's fine if they want to do that. It's the ones who are eating too much and not saving that we're more worried about. OK.

So this looks like Thai villages, but the life cycle model is universal. And just to shift gears, this is Larry Kotlikoff's website. He's both a professor and an entrepreneur.

So he is selling this product, this ESB planner. But he has a lot of great economic research on life cycle planning. And no one wants to splurge today and starve tomorrow, nor the opposite. Lifecycle consumption-smoothing balances our goals.

Now, you would think, could everyone do that? The answer is no. There's only two companies I know of in the US that do this. And it's not Vanguard. It's not Fidelity.

They're adopting kind of a rule of thumb. They ask you what's your current expenses? And then say 70% of that is target for retirement. And then how much do you have to save to make that goal? And so consumption could be low because you're saving a lot. And that can end up in a very imbalanced profile in which you're starving today in order to splurge tomorrow.

Whereas, we're using this strictly concave utility, and so is Larry, to say we want to smooth over time. Just like we want to smooth over bad shocks, we want to smooth over time, as in the medieval village seed and so on. OK.

So let me make a transition then to the other account, which is cash flow. And I'm going to do this by explaining a little bit about Bitcoin, and many of you probably know about this already. So every transection in Bitcoin is an encoded message sent out with a-- generated with a private/public key scheme, which has the advantage that it's known which node sent the message.

And there's full commitment to the message. You can't go back and deny that you sent it. And it prevents fraud. Nobody can transact on your ID because nobody else has the private key, unless you get careless.

So a transection in Bitcoin is the moving of the coin from one person to another. The message sent might be, I want to give the coin to person X. Arguably, because person X is a merchant and you're buying something or whatever. The various transactions are accumulated into a block, and the block is subject to a validation routine, which is mining and it's quite costly. I'll come back to that later.

The point here is that Bitcoin is combining two different things. One is the statement of cash flow as the coin moves from out or in to the person in question. And the other is the balance sheet, which is at a moment in time how many coins do you own? So the accounting of the cash flow works for Bitcoin. It also works for paper currency in Thailand. So let me show you a statement of cash flow as an accounting statement.

Remember, before, I showed you income and balance sheet. I haven't shown you the cash flow statement until now. It's pretty small print. Let me just give you the big picture here.

We have cash flow from production, which you can imagine, if you sold something in got paper currency as a consequence, you'd have a plus. If you're buying fertilizer to grow the rice, you have a minus sign. So you have inflows and outflows associated with production. You can spend money on consumption. That's typically an outflow. You can spend money on investment, also typically an outflow.

So we have cash flow from consumption and investment, and then we have cash flow from financing, which is, you can put cash in the bank as a deposit, or you could borrow cash. Hopefully, you get it as an inflow, through gifts and so on. So cash flow from production, consumption, and financing are the three categories.

And this stuff up here is complicated. When I said I'd shown you the income statement, but not the cash flow statement, you may have been thinking, well, aren't they the same? And the answer is, well, they could be, but not necessarily.

So if you're a business, and you sell something and the guy doesn't pay you, then he's going to pay later, that's an account receivable, but there's no cash. So the sale would count as revenue in the income statement using the accrued notion of accounting, but it would not count as a cash inflow. So oddly enough, this top stuff here is just starting with the income statement and making adjustments to it according to whether the movement of cash you have to adjust one way or the other.

So, I mean, if you have an account receivable, and you didn't get the cash, then you have to subtract off from accrued income to get the cash flow because you're counting something that is not yet associated with the cash inflow. All right. Mainly, I'm just trying to give you the big picture here. This is the statement of cash flow.

And it is true, and this is automatically true in Bitcoin, that you can check whether the change in cash holdings that come at the bottom line of the statement of cash flow are exactly equal or not to the change in cash that we have on the balance sheet. So you have cash in a moment in time, and then you have these inflows and outflows. You have cash in a future period.

In each of those periods now and in the future, you have two balance sheets. Take the difference in the balance sheet, that's the change in the cash flow on the balance sheet, and it should be exactly equal to what's showing up here in the cash flow from the statement of cash flow.

So we run a check on the Thai data. And, again, if we're making a mistake these are not equal, and that's a big guide to us. I should save this for later, but I will tell you that no top US survey does this. Not the Survey of Consumer Finances, not the PSID. They did not gather variables and check what they gathered against these comprehensive statements of integrated accounts.

And Bitcoin is automatic. Because you either have it in your account and you spend it, or you don't spend it and it's still in the account. So that's kind of the unspoken advantage of Bitcoin.

There is something that people talk about a lot on Bitcoin and that is consensus. So I keep giving this example in Bitcoin that Bob wants to transfer a coin to Alice. We have both Alice's accounts and Bob's account. And what Bob gives up is what Alice gets.

So the accounts have to be consistent with each other across the households. And that is true in Bitcoin. But it is not-- well, and again, this is a bit technical for now, but the whole thing about the mining and the validation of the transactions in Bitcoin is to precisely be assured that Bob is not double-spending the coin. You can't spend it twice. And they kind of figured out a way to do that-- a very ingenious way to do that.

But anyway, the concept of reconciliation and validation could be taken to accounting in the Thai villages, in the sense that we really did try to measure all the transactions. And when a household says that it bought something with cash, if the person they bought it from is in the village, then that person should get the cash, while the person in question is giving it up.

So it should appear twice, both in the accounts of the person buying and the accounts of the person selling. So with one of your former classmates, who was a student in this class about two years ago, Chad, we set out-- he wanted to help on these projects. He was an RA. And I'm not pushing that hard, but if any of you are interested, let me know. There's a lot of interesting stuff to do.

So Chad was very interested, so here's what he did. He looked literally to see whether what shows up in one person's account is also showing up in the other one or not. And these are some of the details of going through and doing that as an algorithm. He created these accounts. He checks to see if it's reported. If not, et cetera, et cetera.

Now, right in the middle of this slide it jumps from Thailand and the thing that Chad was doing to a very different application, which is the stock market in Australia. So the stock market in Australia is about to run on distributed ledgers the way Bitcoin does. Not using Bitcoin, but using the ledgers.

And so the idea is when someone owns the stock, and there's a buyer and a seller, the traditional legacy system is that they enter into a deal. And then only after a day, or maybe three, the process of reconciliation kicks in and then they settle. So the trade of the stocks is separate from the settlement of the stocks. And believe it or not, there's often huge discrepancies.

For one thing, prices have moved. So a household could claim, I didn't buy it. And in the New York financial markets, there was an attempt to provide this in the repo market. Because every day in the US repo market with treasuries and so on, they spend two hours at the end of each day after the end of trading trying to reconcile accounts.

State Street Bank here in Boston is a big third-party broker for people that hold accounts, and it does this kind of reconciliation. It's quite costly and quite time-consuming. But if you're entering these trades onto the ledgers, there's instantaneous reconciliation. And so they could settle right away if they wanted to. The decision to not pay for it is a credit decision. They're not forced to settle immediately.

But in principle, trade and settlement can happen at exactly the same time all through the distributed ledgers. And the one in Australia is able to process 115 million daily trades. It's 6,300 trades per second continuously over five hours. That well exceeds the volume in Bitcoin. Bitcoin is not capable of doing this because of the mining. But these other ledger technologies are a lot faster and a lot cheaper.

So let's come back to the village and think about how an individual household is managing its currency. Let me back up for a second. That seemed like an abrupt transition. I featured in lecture one that we will say a lot about Bitcoin and cryptocurrencies, and it's spread out throughout the lectures.

So Bitcoin, blockchain, and distributed ledgers-- it's the first time it really came up. And I started with a ledger. So when we're talking about financial accounts, that's why I started to bring in Bitcoin. And we will revisit Bitcoin in other lectures when we talk about monetary economics and so on. OK.

So now, back to paper currency management. I say paper currency just to constantly remind myself that cash is kind of ambiguous. When you talk about cash in the US you mean, currency, yeah. But also, demand deposits and other variables, like maybe even debit cards. But in Thailand, it's paper currency. They love it. That's what they use. We're on that momentarily.

So I'm going to, in notation, describe a Miller/Orr model-- the management of cash for firms, for business firms, which do not have regular nor predictable flows. And the result solving the model is rather intuitive. There's going to be bounds.

So cash holdings can drift around. But when they go above an upper bound, it's time to make an adjustment and deposit the paper in the bank, or lend it out. And likewise, when it goes below a lower bound, they're going to have to withdraw money or call in a loan.

So sometimes that lower bound is zero, obviously. You can't run out of stock of something. It can't go negative. So that's going to be what the solution looks like.

So let me introduce a little bit of notation, e being expenditure outflows of paper currency, which looks like c. But it's basically more than consumption. It's all expenses out, and then y being stuff coming in. OK. So expenditure is outflows minus inflows, which could be positive or negative.

So we postulate a stochastic process for the cash flows, expenditure flows, and estimate in the Thai data. We postulate that mostly things are pretty smooth and predictable. Delta refers to the length of the time interval. And it's easier conceptually to think of delta as equal to 1.

So with roughly probability 1/2, they're going to have an expenditure of c, but it could be a higher order of magnitude, sigma. Or with residual probability roughly 1/2, it's c minus sigma. So these are the ups and downs of expenditures-- probabilistic basis. How big or small depends on how big sigma is.

And I keep saying roughly probability 1/2, because there are these zapper things that are very large inflows and outflows. So expenditure is an outflow. So positive zapper, Zp, is a very large expenditure that happens with probability kappa. Again, say delta equal 1.

So this could be like 2% or 3% of the time, there is an outflow which is more than two standard deviations away from the mean positive outflow. Or it could be an inflow. So these are rare revenues coming in with different kappas and different Zs, which will be estimated in the data.

And these shocks-- you take the probability of the positive outflow, the probability of the negative outflow add them together, delta, say, is still equal to 1. So with probability 1, something happens. With probability 1 minus the sum of those kappas, you're in this middle branch. And that's the 1/2 adjusted up or down. As always, I'm assuming you will interrupt me if things are not clear. OK.

So let's think about how much cash they have in house that's paper. So they have a certain amount of cash in a moment in time. They have an expenditure outflow, which could be positive or negative. And they're at t, and we're going to try to do the accounting for how much money they would have at t plus delta, where again, delta could be equal to 1.

Money at t less outflows, and then they could adjust. They could either withdraw money or deposit it. So these two things the w and d is endogenous. And the model is going to have to solve for w and d, when to do it, as I was saying, and how much to do. And that will result in this endogenously determined statement of cash flow.

You can think about this like life cycle planning, except now it's not over 50 years. It may be a matter of a couple of weeks or 30 days. In the data, it's 30 days because that's when we get the data.

They're trying to minimize costs. Now, there's two costs. One is kind of obvious, really, for economists, which is if you're holding cash and it doesn't earn interest, you've lost the interest you could have earned if you had deposited it in the bank.

So that's a real cost to the household. That's the opportunity cost to the household. We're going to let the interest rate in the bank be capital R. If you hold too much cash, you're giving up a lot of foregone interest. That's a cost.

There's also a physical adjustment cost to get to the bank, which you can think of as the hassle, or in these villages you've got to get in your truck and drive to the village, to the town, et cetera. So that should have had the notation here, but it's a little b. You'll see a little b momentarily.

So the next slide-- don't panic. I'm going to walk you through the value function. It looks pretty intense at first. They're sitting here with a certain amount of money, m, and then deciding whether they want to stay put and carry m into the second period or adjust.

Now, if m is already negative, they're in trouble. Can't go negative, so they have to adjust. And you get this cost, b, which is gone. This cost function is not utility. It's not net profits. It's a cost, and they're trying to minimize the cost. So this is enters positively because it's a cost.

They've gone to all the costs, to getting to the bank, they get to make an adjustment, which means they can choose a new m prime going out. And so they'll try to minimize the value from that point on by that choice of m prime-- minimize the cost from that point on.

Here, m is positive, but they still get to choose whether to adjust or not. And if they don't adjust, they keep it. They suffer the opportunity cost, the lost interest in the bank. And then they'll have their value function tomorrow, which is discounted by little r. That's like a discount rate. And things next period are stochastic.

So you take expectations over the future e values following that stochastic process. So this is v of m minus e is all expectation, is all written down here in terms of these branches, which is positive zappers, negative zappers, ups and downs, roughly probability 1/2, et cetera. So this is the expected cost in the future if you don't make a monetary adjustment today.

So this is the second or third time you've seen these value functions. I introduced them at the end of the medieval village dynamic programming problem. We then used that last class to talk about implications of risk sharing for production, with the planner's value function. This is the third time. This is the value function for a household minimizing its costs.

More fun-- let me show you some pictures. So these are real live people. This is the cash flow of this house-- the cash holdings of this household. Up and up and up and up she goes.

This one also had a trend, and then boom, down, and then to zero, and then back up again. You can't see it very well. I'm sorry. The pictures are too small. But what's plotted over here are net cash inflows or outflows, along with, say, a deposit when you have a net inflow, or a withdrawal when you have a net outflow. So there's three lines here, and sometimes things correspond the way I was just implying. You're short, so you withdraw, et cetera.

But when-- you can't see much here, but when you integrate it out and look at the consequence for the levels using this flow equation, then you can see what their stocks are. It looks like a lot of money. It is going to be unbelievable-- unbelievable in the sense of three years research running to try to understand this pattern.

So how much money does a typical household have-- currency, paper, squirreled away in the house? In the rural villages that we were looking at, roughly 32 months worth of consumption. What's in your wallet? You're going to say credit card.

These guys-- it's not in a wallet. It's in the house or hidden in the backyard. Almost three years worth of expenditures in cash-- cash equivalent consumption. We measure the consumption for each household, and we see how much cash they have, which we measured in the data. 32 months-- it's unbelievable.

And we also get deposits and withdrawals measured in the data. That's a lot of detail to absorb. Let me show you what's going on in the model. We estimate the cash flow. We recover all the parameters that we need, take a stand on the interest rate in the bank, and then look at how small or big the cost would be. b is varying down here. And the consequences for currency, especially but also withdrawals and deposits.

So taking the reality to be the calibrated expenditure flows we see in the data, along with the few other reasonable assumptions, I assure you we can look at how much cash they should be holding relative to consumption. That's this blue line. And we can barely get it up to three.

The number in the data is 32. I mean, you could imagine increasing those costs and decreasing the interest. And this is like asymptoting out to something flat. So you end up with something ridiculous. The model isn't working. You don't go to the bank and forego two years worth of consumption while you're doing your banking transection. It's just not a reasonable value.

The sign makes sense. The higher the cost, b, of going to the bank to either withdraw or deposit, the less you're going to do it. So when you do it, the transactions are bigger. And meanwhile, you're holding more cash. If b goes to infinity, then it's all in cash. Forget the bank. And when you go rarely, then you have to make up a big difference. And this is the frequency, which again, as b goes up, the frequency of going to the bank is going down.

Question-- so because we did that life cycle planner, we wanted to evaluate as a randomized controlled trial how well or poorly it was working, and in the process, we gathered monthly data for the urban household. So we're not in villages anymore. We're in towns or even city neighborhoods.

In each of the four provinces, we basically replicated all the existing surveys. So the same questions are being asked. And this is kind of buried here, but if you look at the normalized cash in terms of consumption for the urban sample, and look over here, we got a number of 3.87-- much closer to 3. Nothing like the 37. Well, actually, I should be over here at the median. So it's 1.94. But when we do some adjustments in the data, I'm happy to tell you about them.

That's within reach of the model. We calibrate the model again using the cash statistics-- the zapping, big shock, small shock, how often they happen, et cetera et cetera, and get a model prediction. And the implied value, m, depending on whether you use the mean or the median, is something like 2, which is close to what we're getting in the data.

So now we're in kind of business, in the sense that a reasonable economic model is actually predicting pretty well how much cash they're holding. Whereas, in the rural data it was a disaster. We didn't get close.

So now we switch into how can we understand the error that the households are making with their cash management? And the guess is-- several guesses. Some turned out to be correct and some not.

So why would cash holdings be so much lower in the urban area relative to the rural area? Is it because households in the urban survey deposit more often in larger amounts? Or is it because they withdraw in the urban areas less and withdraw smaller amounts?

So for deposits, it turns out households in the rural area deposit more and not less, and they deposit larger amounts. So this is exactly the opposite of what we would have predicted because the cash is going into the bank or into the local money market.

However, for withdrawals, it works the way we would have conjectured, that when urban households put money in the bank they don't withdraw it very often, and they withdraw a lot less. It almost looks like they're putting their proxy for their salary into the bank. And any other money which they have, from business and so on, goes into the bank and stays there. They just withdraw what they need for consumption, taking a, quote unquote, salary equivalent out and leaving the rest in. So this is why their money value's a lot lower.

There's still something to fix, which just like in life cycle planner, we could create from the research cash flow app and put it on the phone and tell households how to manage it. And, in fact, the question is, why are they having cash at all? Why don't they just put it on their cell phone?

And that brings us to-- oh. And the costs of mismanagement of currency are very high-- something like 9 and 1/2% of monthly consumption. That's huge in economics. Say, exaggerate a bit-- 10, or call it 5 if you want. 5% of monthly consumption could be saved or used for consumption if they just managed their currency better. So one way to put this-- enormous cost of mismanagement of petty cash.

Businesses get cash advisory services. They pay for it. Why should these poor households not have access to better management tools? OK. Sorry. That's my view. OK.

Go to Sweden. How much paper currency are they using? Close to nothing. They're gone-- it's gone. This is the amount of currency relative to GDP ever-declining. This is the absolute level of cash, going down in an absolute not only relative basis.

So I got this and this slide from the governor of the Riksbank in Sweden. So I don't know whether you will appreciate the humor here. But this is a picture of a stone. So economists and others have written books about the first coins, the first use of cash, and there were claims on stones that were buried at the bottom of the lake.

So now in Sweden, they're transferring the stone around on their cell phone. Well, not really. It's Swish, basically. They have a bank account, and for no money at all households can transfer cash to each other, buy things, and so on. And it's replaced the paper currency.

It's actually a bit of a concern for them, which is why they're thinking if there's no paper currency other than reserves, there's no role for the central bank. So they feel they should be providing a payment service as a backup. So they're going to create, almost surely, a central bank digital currency. So households will have digital accounts in the central bank, which is what South Africa is doing. China has a different version of it and so on.

This happened in the Bahamas. They created something called the Sand Dollar. The virus is accelerating this problem because people need government aid, and they can't get the aid-- truck the paper currency over there. When this disastrous Hurricane Dorian destroyed an island in the Bahamas, they couldn't get emergency assistance. So these are things that are happening all over the world right now.

This last slide is Kenya. And I mentioned they have two kinds of currency. One is Kenyan shillings-- used to be a British colony. And the other is e-credits at Safaricom, which is a cell phone provider.

So what you do is you take your shillings, you go to this Safaricom agent. You give him your Fiat money, and you get credited cell accounts called M-Pesa. And then you can use that to give to your mom back in the village, who goes to the agent there and transfers it back out into shillings, or you can buy stuff.

So here you see the use of the paper currency, but also this e-currency, which is created by Safaricom. And it has reached almost everybody. Some economies try to do this. We kind of tried in Thailand. I have not yet succeeded.

In Kenya, even the lowest income quartile-- we need to update this graph. Tavneet Suri did this. Roughly, 85% of the low-income population, and a similar, higher number for those with no bank accounts, are using Safaricom.

With no bank accounts-- in other words, you don't need the traditional legacy infrastructure. And this is, again, the allure, the so-called disintermediation that the fintechs are bringing to the world. Of course, the commercial bankers don't like it much.

We talked about risk-sharing. And Tavneet and her co-author, Billy Jack, looked at how much consumption is being smooth in Kenya, comparing households who have not yet gotten early on the M-Pesa expansion versus those that did. And there's a lot more smoother consumption, which you see these common themes.

We did risk-sharing. We did gift giving in Thai villages. Now people are giving gifts and transfers through M-Pesa. And so it has an impact.

Well, what's the big deal? It's a big deal for Kenya. Sweden already has this infrastructure. It's not a big deal. But it makes a big deal for low-income countries.

There are some qualifications, though. First of all, it's not on a distributed ledger. So you have to trust Safaricom to keep the accounts properly, and they have the cash. So they're supposed to put it in a commercial bank. You trust them to do that. So you have to rely on a third-party vendor in Kenya.

Not only that, the money's in the bank, and Kenya is subject to runs. So how safe is that? You may think you're going to get the money back, but will you really? And in other economies, they don't have that kind of trust, so it doesn't take off.

The other thing is those agents, these guys, they run out of stuff. So you go and take your shillings and say I want M-Pesa. And they can't create it. There's no net creation. It has to be a one-to-one transfer. They have to have it in their accounts already. And if don't have it in the account, they've got to say, well, wait, or come back tomorrow.

Likewise, and perhaps more obviously for the currency withdrawal, say here's my cell account. I want to cash it out. Oh, I don't have that cash in my drawer. What Tavneet discovered is these agents are kind of linked to each other in an informal money market, giving gifts to each other, actually, and borrowing and lending in order to smooth the fluctuation, from smoothing at the household level, to smoothing at the wholesale level, and similar models would apply.

So this brings us to the theme of building financial infrastructure. How can we not just get currency on ledgers and transfer it around in a cheap way? How can we build in the trust? And how can we deal with this shortage problem among over-the-counter vendors and agents and so on? And next lecture, starting next time, we're going to do mechanism design and private information, and I'll tell you how to design-- at some point-- how to design this system.

All right. Went over by an extra minute today. Questions? All right. So otherwise, I'll see you next Thursday or in my office momentarily.