[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: Greetings, everyone. So before I start sharing my screen, let me ask if there are any questions from the last lecture. OK. So, where are we in terms of class?

We have a review session today. The exam is next Thursday, and it covers material from lectures 8 through 15. And I'm going to take questions on those, or scroll through them, or review, or something. So that's the main thing that's happening.

So the overview of that is we did the theory, and then we started in risk-sharing, with two applications, including production. We went to financial modeling. Then we went to contracts to get at obstacles, like information and moral hazard problems. We did a theory lecture of sorts, links to Bitcoin, and ledgers, and then another application lecture. We went to international trade in the context of all raising equilibria, and we did two trade lectures, including the one we did last time. So that's the overview.

Some words about the exam, again, even though in principle the thing is comprehensive, that's because it's building on material that had previously been covered. But the focus will be on the material from lectures 8 through 15. And again, because it's limited to 80 minutes, there's no way I can ask a question about everything you might be anticipating. But hopefully, the questions themselves are reasonable, and fair, and not terribly surprising if you've been following the lectures, doing the p-sets, and keeping up to date.

And you have all the review session material, as well. That's the best guide to the lectures, or there was review questions we do at the beginning of each class. Which brings me to the next thing. So here's the study guide. And let's start with what we did last time, focusing on the US. I talked about China shocks and accounting identities, which led up to this question using the lens of the accounting identities. What is the preliminary evidence for the impact of China shocks on the variables of the current account?

Somehow, I'm in the mood to take volunteers, rather than call on you. Any volunteers for that? Quiet group today. Well, I will-- how about Charles?

AUDIENCE: Oh, let's see. I think I haven't got around to watching your last lecture recording yet. So I'm not exactly sure. But oh, yeah. Sorry. Is it this one?

ROBERT TOWNSEND: Yeah it's in and out, synchronous and asynchronous. How about is someone ready to volunteer? China shock financial accounts and what did we find?

AUDIENCE: I can try. I'm also very unsure. But is this about looking at the current account, and how the trade balance would be negative. And then that's kind of giving us evidence about how-- the China shock was about how since low-skilled workers, we have less of that commodity and more of high-skilled workers. So then we're like low-skilled workers are going to be losing in trade.

ROBERT TOWNSEND: That's also true. So I was in part looking for this thing about the current account deficit. And so, yes, it's interstate trade and states that are hit the worst by China shock, meaning that their local manufacturers had to compete with these Chinese imports. Those states export less to other states, and also to some extent, import more.

Does anyone remember what else is on the left-hand side of the current account? Current account consists of the trade balance plus something.

AUDIENCE: Net income?

ROBERT TOWNSEND: Net income, yeah, so meaning gifts incoming or remittances for factor payments earned in other states. And does anyone remember? Did it go up or down?

[INTERPOSING VOICES]

AUDIENCE: I'm sorry. Go ahead. I was going to say that I think transfers were positive.

ROBERT TOWNSEND: Yeah, good. OK. So this linked up to risk-sharing. And I'm not sure if I emphasized it last time. But given that it's quote, a "shock," a China shock, or at least an adverse trend, the question is, what arrangements would cover these shocks?

Public transfers in principle could have done that through fiscal taxes and transfers. But what we found in the data was that it was the private transfers that seemed over a 10-year period to go up. So it definitely relates to all the risk-sharing stuff we did in the last few lectures and that we are about to review today.

And what about the other side of the current account, financial stuff, anyone remember?

Maybe these financial accounts are tricky. Well, we'll review it today. But we didn't find very much adjusting. On the right-hand side, we should see a given state, losing-- losing its claims on foreigners in this other states, in the sense of running a deficit. So they should be running down their savings accounts, and so on. But so far, we haven't found that in the data.

But it's still good to know what was the financial adjustment on the right-hand side, because the accounts have to balance. Let's do one more. What in practice in the US is the difficulty in reconciling income flows with the change in the balance sheet? That is to say explain what accounting identity ought to hold, and what in contrast we actually have in integrated macro accounts and household surveys.

AUDIENCE: I think I could take a shot at what identity might-- I think I know what the identity is. I'm not sure exactly how it might feel though.

ROBERT TOWNSEND: Go ahead.

AUDIENCE: Isn't like, the identity should probably be something like the expenditures by households which should be matched up by the incomes to whatever firm is selling the goods or something like that. Those two sides of the balance sheets should match up if it's within a country within the country. Right? So it's something like that, is it?

ROBERT TOWNSEND: Well, it has to do more with what is the relationship between flows and the income statement, and assets in the balance sheet?

AUDIENCE: All right. Yeah, I'm sorry.

ROBERT TOWNSEND: Anyone else?

AUDIENCE: Sure, but basically I think it all should be able to match in a sense that we should be able to account for these dollars on both sides, in both the flows and in the assets. But because of the way that the data is structured, there's no basically central accounting place where all of these are accounted for on a regional level in a way that just like has all of the data. And so basically nothing matches up. But we had that table with errors in the-- I forget what it was. Like 20 plus or percent or something like that.

ROBERT TOWNSEND: Yes, that's correct. Great answer. So the logic of it is you have flows in the income statement. You have revenues less expenses equal profits, or for households net income. And what they don't eat or pay out as dividends, they should be saving. So the changes in the balance sheet, some balance sheet items ought to go up if households or firms are saving. And in practice, since as was just said, the balance sheet and income statements are not collected at the same time, we see very big discrepancies between the two.

So the rest of the class last time, almost up to the very end was about what this discrepancy means for attempts to measure inequality, or the impact of COVID, and so on. So I'll just say what I was trying to do in the class is not to get you to memorize financial account identities, but rather to think creatively about how to use financial account data, and so on.

Which brought me to the very end, where I talked about trying to fill in those gaps, and gathering detailed level information from households in West Virginia, and so on. But I won't ask about that. Well, that may come up again today, actually.

So let's go to the review from lecture 8 on, so I will try to highlight. We did risk sharing in India, in the villages, and also in medieval England in terms of dividing up the land. I talked at the beginning of the lecture about the high-risk environment and the way households could respond to it, either by diversifying or by helping each other out with financial gifts and transfers, and then ways to test.

So in India, we decided to test by looking at incomes and consumptions, and not looking at the mechanism. And in contrast, with medieval England, we focus more on whether or not the mechanism of dividing up land could achieve an optimal allocation. So then we went through characteristics of village economies. And the main thing here was the risk.

In particular, in village India, we looked at the risk of harvest yields on various crops, of crops by soil type, of income by type. And we uniformly saw a very high risk on all these categories, as well as some diversification possibilities. So for income, not everyone has the same income sources, and there's high risk of each. But the covariances are sometimes relatively low.

So the focus here was on the fact that households are not completely diversified. If all the households did the same thing if, they planted the same crops in the same soil, and have the same income sources then their initial outcomes would be essentially the same. And if that's true, then there's not much left to share. But we see the opposite. We see they're not diversified over all these various crops, and soils, and activities. So there's a lot of idiosyncratic risk. And that was what was clear from this picture of what we call the Rocky Mountains, where the ups and downs are not together, not going up together across households in a given year, and so on.

Whereas, they do achieve a lot of smoothing. So something is happening in the villages that whatever mechanism it is, is allowing them to smooth. Now one question that's come up in Piazza are these profiles. I mean so this is a bit of a review, but kind of interesting. We see this income profile from low to high. Households with a lot of land are consistently getting higher income, despite the fluctuations.

And there's a big compression here. So something is transferring income from the landed households to the landless. But a little bit of this profile does remain. So it's not equal treatment. It's not as though everyone is getting the same amount of everything. There is heterogeneity in the treatment. What that mechanism is, is not so clear. It's some kind of, at least over these 10 years, redistribution of income from the wealthy to the poor. But it doesn't work perfectly. It doesn't completely equalize things.

So this lecture was featuring something we had done in the previous one, which is to determine the pareto optimal allocation of risk, we maximize lambda weighted sums of expected utilities, subject to the resource constraint. And this way of dealing with problems generalizes considerably and in the subsequent lectures.

So when we solved this problem, we looked at the first order conditions. We could see that weighted marginal utilities should be equalized to a common Lagrange multiplier. And you can do the algebra, and actually solve for the consumption that's supposed to result, and get these kinds of closed form solutions having to do with consumption moving with the aggregate, and not moving at all with idiosyncratic shocks, if the risk sharing is optimal. So here is another version of it, individual consumption co-moves with aggregate consumption for all the households, but should not move with things like individual incomes once we control for the aggregate.

So we tested this model to see whether these coefficients were 0 or not, and largely they are. I mean there is some residual risk that is being borne by the households, depending on the village and the source of income. But the benchmark was doing quite well.

AUDIENCE: [INAUDIBLE].

ROBERT TOWNSEND: I would be delighted.

AUDIENCE: Yeah. I mean here, we assume-- I see the sample. There was about 40 households in each villagers. But in this model, we assume that the households should be large amount. If not, then the idiosyncratic risk might not be shared, right? And for example, if there was just two households and two situations, then although they face the idiosyncratic risks, but there was like 0.25 probability both of them were facing negative shocks and/or both of them face the positive shock. So they cannot share the risks.

So in this model, we assume the sample is large enough. That is why we think that the idiosyncratic shocks will not affect people's consumption? Yeah, each household consumption.

ROBERT TOWNSEND: OK. So it's a great question. What can we test if we only have a subset of the households, and are we being misled? Well, actually you could think about it as testing for subsets. And the theory holds. Of course, what is aggregate consumption for two households would be different from what it would be if you had data on all of them. But the theory still works for two, or three, or just a handful. If they are bearing risk optimally, then they should be bearing risk optimally in any subset.

Let me put that a different way. The aggregate consumption within the group may be moving around. Individual consumptions in the group should move with that. The group consumption actually could look like an idiosyncratic shock if we considered the whole village. So yet more restrictions are implied when we and if we had data from the larger sample. But you can still test the theory with a small number of households. It's just that it's not as powerful. You could reject, however. Does that help?

AUDIENCE: Yeah.

ROBERT TOWNSEND: OK. It was a great question. So then the next thing we did was to take the same theory to different applications, back to medieval villages, and the landholding patterns which are spread out all over the different fields. And the idea is, does that make sense or not, and would dividing up the land in that way exactly achieve an optimal allocation of risk bearing or not.

So we took the different types of land to have a different vector of yields. We talked about the endowment, so to speak, the yields of land type k across states epsilon. And then we have capital K types of land and capital S types of shocks. We can look at the covariance and the risk, just the way we did in India. So again, we get this picture, very, very risky with some covariances pretty low.

And then I took you through a series of examples to see how close we could get to an optimal allocation of risk bearing. The first was household 1 is all of land type 1. Household 2 has all of land type 2. But these returns do not co-move, hence this would not constitute an optimal allocation of risk bearing. Household 1 would go 10, 6. Household 2 would go 2, 4. So although the aggregate is going down, I should say, from 12 to 10, household 2's consumption if they ate all their own crops would be going up. So that's a restriction.

Then we considered a different way of doing it. Give those household shares. Household j say having a share alpha j of the two different crops. So now, this thing goes in the aggregate 12, 10, 7. Obviously whatever household 1's share is, if it's like one third, it's 1/3 of that aggregate. Household 2's shares would be 2/3 of that aggregate. So now they're each bearing a constant share of the aggregate and their consumptions are moving with the aggregate. So that might actually work. That's this.

And then we looked at special utility functions, and started fussing about intercepts, and risk aversion, and so on. But then we come to the general idea that we have as many land types with independent return vectors as we have states of the world.

And then we have a target optimal consumption allocation. So we're trying to achieve this allocation that came from maximizing a lambda weighted sums of utilities by dividing the different land types into shares. Except now we have even more power, because we do not require that household j have the same shares and all of the land types k. We have this alpha jk for household j in the land type k. And it's not perfect. But in general, since this vector of endowments is non-singular, you can solve for the consumptions. And one hopes that the shares will be positive in the end.

So if there are as many types of land as there are states of the world, they have a shot at achieving the optimum by entirely dividing up the land, and not having to give gifts or transfers to one another after the fact. So then we come to-- we put in particular utility functions, and we assumed that we're going to solve this problem, maximizing lambda weighted sums of utilities, by the shares alpha jk. But we're going to assume that there are fewer land types than there are states of the world, namely only two types of land, and a larger number as states of the world.

So now, we know we in general cannot achieve an optimum. But we were trying to describe this constrained optimal arrangement. I guess what I want to emphasize is despite the restrictions, we're still solving this programming problem. We're maximizing lambda weighted sums of utilities subject to constraints, namely the shares have to add up. So the method of achieving characterizing pareto optimal allocations is the same.

And then we looked to see with those particular utility functions how far off we would be from achieving the optimum. And here you can actually see, as the aggregate endowment moves the adjustment. In other words, even if they did as much as they could by dividing up the two land types, it would not be perfect. And there should be incoming or outgoing transfers to try to make up the difference to get to the full risk sharing optimum. And that was the end of this lecture.

Are there questions about what we did or what we were trying to accomplish in lecture 8?

So then the second of the two application lectures included production in lecture 9. And the main difference between what we did just now and this lecture, is we actually have physical production, rather than treating output as an endowment. But we still solve this programming problem. So we're maximizing lambda weighted sums of expected utilities, just as before, subject to resource constraints, just as before. The choice variables now include the transfers, just as before, but we also are choosing which projects to fund and by how much capital investment.

So the reason that transfers are the same as before, is because we have transfers or consumption. You can move consumption around by moving transfers. So it's arbitrary choice whether to solve for c or tau. The more interesting part, because it's new to this lecture, are these capital stocks, which has to do with deciding what to do with the current wealth, namely eat it or invested in projects for tomorrow.

Here is the starting point, w, which is the trees and the fruit from the trees. And this was the ending point, what to do, either eat or save. Oh then, back to the accounting identities, you either eat something or save something, as in the flows.

So we got two first order conditions out of this problem, one very familiar, lambda weighted marginal utilities are equated to each other. We also got a new equation having to do with the marginal utility of consumption today, and the marginal utility of consumption next period. And those typically will balance each other off for any project they're running. There was an issue that this thing could have been negative, in which case that project would not be implemented, because they would actually not want to put stuff in there. They'd want to take stuff out. But they hit zero.

So the rest of it was to take expressions for these first order conditions, and start manipulating them. I'm not going to review that today. It's algebra. It's straightforward. We can even make some assumptions about linear technologies, subject to risk, and quadratic utility functions. And when we did that, we got two expressions directly from the theory that we could take to the data.

One, we can measure from the data how household j returns are moving with the village aggregate return, this beta being a measure of the covariation of household j's return with the village average. Then having used the panel data to get these beta j's, we look in the cross section to see how the betas are related to each other, and in particular to the overall average or expected return. So the main empirical finding here was that returns have to be higher when the betas are higher, because those particular activities from household j co-move move a lot with the aggregate. So in effect, it's very risky, not so many diversification possibilities come from including household j in the village portfolio.

And so if they're going to be running those projects of household j, there has to be a compensating differential, and that's the higher expected return. So then we started adopting this language of risk premia, which is the difference between the expected return and the risk-free rate. The risk premia has to compensate for the risk.

So then we went to the Thai data, where we have everything we need. We have the income and the balance sheet. Oh yeah, so here are these accounts. This, I think, was the first time the accounts started to come up explicitly. We used the Thai monthly data over 13 years, and I featured net income to be profits, to be revenue minus costs. And then we looked at the assets to be assets of various kinds, typically physical tangible assets.

So here you have the income statement and the balance sheet. I just didn't call it that at the time, but that's what we're using. So then we tested the model, and found that the betas, the expected returns, are indeed higher when the betas are higher, just like the model predicted. And here, we are back to idiosyncratic and aggregate risk, which was a theme I mentioned in the previous lecture. We can decompose the aggregate variation of household j, and explain the unexplained part. The explained part has to do with the covariation with the village average. And the unexplained part is just there were some idiosyncratic risk.

And we decompose the risk premium as well into the part that's coming from the village as a whole, through the beta j, as well as the sigma square sigma j on that epsilon. And then got another decomposition.

So it turned out that, as a quick review, idiosyncratic risk was the largest contributor to total risk, and aggregate village level risk was a small part. So there were a lot of idiosyncratic shocks in these villages, just like we imagined that there were in medieval England, and also that we documented there were in village India.

But when you looked at the risk premia, it went the other way, namely most of the risk has to do with the aggregate risk. So that's consistent with them finding a way to diversify their idiosyncratic shocks. And there was evidence that I presented that they're doing that by gifts, and borrowing, and lending. And that's the mechanism.

We started out talking about how we weren't going to test for mechanisms. Then we looked at ex-anti-division of land. Now we're looking at ex-post gifts and borrowing as a way of getting from their, quote, "production endowments" to their ultimate consumptions. And most of the shocks are covered, not all of them. So they're not completely following the full risk-sharing program, but they definitely diversified mostly 90% of it.

So then we got to particular health shocks, and specific Illnesses, and I showed you that again those illness shocks, not production shocks, per se, also covered by gifts largely, although not entirely. And when they were not covered, they spill over into the labor network with less hiring of employed labor, and also less purchasing of intermediate goods from other households.

So this was documenting that. So this is about kinship groups and risk sharing in village villages in Thailand. It also was about an illness as a shock. It was not about COVID, although it was about illness. The illness was idiosyncratic and not aggregate. The closest we could get in the Thai data to an aggregate was watching those shrimp ponds suffer catastrophe in one of the four provinces. And because it becomes an aggregate shock rather than idiosyncratic shock, it's the gifts didn't come in as much, and it's very harder to diversify.

So that's a theme I keep mentioning now for the last 10 or 15 minutes. Idiosyncratic stuff, in principle, you can diversify. Aggregate stuff you cannot. But the other thing I want to draw your attention to is this network stuff. So the issue is not just in Thailand or India, for that matter, but do these networks exist in the US. And what's going on with COVID?

So that inter-state risk-sharing stuff suggests that networks are not necessarily local. They could be global. But your relatives or other family members are engaged in risk sharing when you experience a shock back home. And then this project in West Virginia, well Michael, I hope you don't mind mentioning it. But Michael, thinking about networks, wondered about networks in the US and the advantage of collecting our own data, indeed would be to better measure the networks among relatives, or friends, and so on, in towns in rural areas in West Virginia.

So everything that we've been doing featuring Thailand, because we have all that data, could be done in other contexts like the US, if we can add that data. And we've definitely decided to try to measure these networks as a suggestion that Michael had from last time. All right. So any questions on this lecture 9?

AUDIENCE: I have a question. In the slides 19 data, when we're decomposing the risk, yeah, here. I remember before we have R [INAUDIBLE] that was represented the risk-free return. But why here, we just have beta and the idiosyncratic? That is because we both subtract the alpha from left-hand side and the right-hand side, or because the results we get before that alpha is not significant, so we don't care about it?

ROBERT TOWNSEND: Yeah, another great question. So there are two ways to do it. We should have been thinking about the rate of return minus the risk-free rate. But in fact, in the Thai data, we took the position that the risk-free rate was 0. So that's why we kind of dropped it.

AUDIENCE: Thank you.

ROBERT TOWNSEND: It's like having rice. You can put rice in storage. And you're going to get back more or less what you put in, et cetera. So that was the rationale for it. Other questions? OK.

So that was all the stuff on risk sharing. Then in lecture 10, we started to introduce these ledgers, and financial accounts, and other measurement issues. Actually, there's a lot going on in this lecture, as I'll review with you. We did the wealth planner stuff, looking at the balance sheet and the income statement. Then I went to the other statement, the statement of cash flow. And we thought about Bitcoin as an example, having the balance sheet, and the cash flow in one place. And with Bitcoin, that motivated a review of distributed ledgers, and how it works with currency, physical currency, as opposed to the ledgers. And then we touched on Sweden and Kenya.

So this lecture introduced Bitcoin. It also introduced distributed ledgers as financial accounts, and also value transfer. So we talked about this statement of income. By now, you may or not remember, it doesn't really matter. This household A appears twice subsequently. This was the first time we introduced it as an example of the income statement, and also household A's balance sheet.

And then we started to talk about the wealth planner. So we implemented this in Thailand, maximize discount expected utility over the lifetime, even allowing for bequests to their kids, subject to the stocks and flows constraint, investing in risky and safe assets, and allowing for the risk sharing which you don't usually see in wealth planners. But we know we needed it for Thailand, because you've already seen how much risk sharing there is. So we want to smooth out the income profile consistent with the ins and outs of those gifts. And the residual is the household's exposure to their own income shocks.

So we explained to those Thai households how the planner works. And again, to motivate this, think about going to the doorstep of a household living in West Virginia, having collected some of their data. Or better yet, what data do we want to collect in West Virginia consistent with running this wealth planner, so we can give them advice? The setting would be different, but the idea is the same.

So with the wealth planner, we find out how much they're planning to give to their kids, if anything. We calibrated their income processes with trends and fluctuations. We looked at their committed consumption goals, and we're very careful about their initial wealth. Again, you come into the period with a certain amount of wealth, and you could consume it, or you could save it. You can save it and risky or safe assets. In principle, you could also borrow, and then that is going to determine wealth for next time.

So I probably didn't emphasize it as much. But this is where the stocks and flows are coming together again early on, through this budget constraint. Initial wealth and consumption is subtracted off, and the remainder goes into risky and safe assets, which will determine wealth next time.

And then we looked at actual flows in the data, the second household, the recommendation of the planner in terms of life cycle which is saving up enough when you're relatively young to be able to survive through older ages.

Some of these households get in trouble. So we try to help this guy, giving advice and actually talking about the realism of the goals. And I made the point that this planner exists for the US. Although I don't know that it's particularly well designed for households in West Virginia, it does take into account details of the US tax code, and Social Security, and so on, but acting as if there's no networks, no risk sharing. Who knows? We'll eventually find out.

Then I made this transition to Bitcoin as an account. So Bitcoin accounts have the ownership of the coin, and also a way to buy something else and transfer the coin to a third party, by sending the proposed transection to be validated. So the Bitcoin ledgers are a combination of the balance sheet and the statement of cash flow. And with the introduction of the idea of cash flow, I went back to the Thai financial accounts to look at that for the first time.

We'd already looked at the income and balance sheet statement. Now we're looking at the cash flow statement. Cash flow from production, consumption, financing, and so on. So then the issue is, what do the ledgers do?

Well, in Thailand, they're largely using paper currency. So they do not keep accounts. And there are no common ledgers. Bitcoin is all on an account, and there are common ledgers. That's why it's called a distributed common ledger. So we actually went to create ledgers that would be common in the Thai data, and discovered a lot of data discrepancies because we hadn't collected the data that way in the first place.

This is also a reminder, if we're collecting data in West Virginia and we're getting the identity of the transacting partner when we interview that partner, those transactions should also show up. So ideally one would do that in real time almost, when memory is fresh, or for that matter their understanding of a transection with a commercial bank which is opposed to what commercial bank is doing. At least you think that, of course, the commercial banks are honest. Just remember Wells Fargo that actually created accounts for households they didn't know about, and then charged them money for it.

So we then talked about the advantage of Bitcoin like distributed ledgers, namely data reconciliation. But in practice, of course, they're using paper currency. So with that opportunity, we looked at another problem of management, which is managing the currency, managing the cash flow, which we modeled and fit to the data, as having trends and stochastic fluctuations. We have the evolution of cash flow, adjusted taking into account the ups and downs with shocks of revenues and expenses, and adjusted by withdrawals and deposits.

The two costs were going to the bank to make the physical adjustment, and the opportunity cost, foregone interest, of holding the currency in the house. Which led to another optimization problem. Now minimizing the expected costs, let me just say that our vision for West Virginia is to combine the high frequency cash flow with the longer term life cycle. Because in practice they are intimately related. So we'll have a synthesis of those frameworks in mind.

And in Thailand, they didn't manage the currency very well. In the rural areas, they're holding far too much of it. And in the urban areas, they actually did better. Of course, one advantage of getting them to switch from paper currency into electronic entries that is then in principle, it's easier to manage portfolios or to send reminders about cash management on apps on phones, if all the transactions are electronic. And I give you some examples of Yodlee and other apps that do this, that collect electronic data.

So anyway, that led finally to, well, do you have to use paper currency? The answer being no. Sweden doesn't much anymore, although they're worried about it in the sense of worried that the private sector takes over all of the payments, and the Riksbank isn't needed anymore, the central bank. And the stuff in between, which was Kenya, so Kenya eMoney, digital money, is created by the cell phone company, allowing people to cash in and cash out, and electronically transfer the credits on the cell account, which turned out to help with guess what? Risk sharing.

So you see a tie-in here back to those risk-sharing lectures. In this case, it's in the context of Kenya, and the fact that your kin back in the village have had a terrible harvest can be mitigated by having the migrant worker in Nairobi send more money. And evidently, that is exactly what happens, as opposed to having to give the Kenyan shillings, paper currency, to a friend. Put it in an envelope, hope for the best, or in a very costly way go back home and hand over the cash. So risk sharing on the ledgers actually helped.

But these are private accounts. So they don't have what we typically refer to as distributed ledgers, where there's third-party validation. All of these transactions, cell phone companies accounts, are trusted to be accurate. Households can check them, when they happen to see if they get credited the way they're supposed to. But other than that, Safaricom could steal the money if they wanted to. They don't, evidently. They put it in the bank, but still, and the banks are subject to runs.

So there's a lot of risk remaining, even though this digital currency is an innovation in the context of Kenya. It certainly would not be in the context of Sweden. They have their own interbank method. So are there questions about this lecture 10?

Again, it's kind of a mouthful. This lecture does both the financial accounts, income statement, balance sheet, and cash flow, and is also looking at two problems namely cash management and lifecycle management, and finally merging into a discussion of what distributed ledgers, Bitcoin and distributed ledgers, do or don't do, and potential substitutes. How good or how bad is currency and Safaricom digital credits, relative to validated systems? Although we haven't really talked about validation. That happens I think in the next lecture. So that was lecture 10.

Lecture 11 was about contracts and mechanism design. And I'll just give an overview. There were two parts to this lecture. The first part is figuring out what to do when there's private information. And the main conclusion was there is a way to write down additional constraints, called truth telling or incentive constraints, that capture the problem of private information. And we will be able to append those on top of the usual programming problems that maximize the lambda weighted sums of utilities.

So we've got a tool, and the tool is getting better and better, because we can do more and more with it. So that was one of the main points of this. And the other was to think about mechanism design, private contracts, having to do with messages, and then back to the ledgers which have to do with the cryptography part of ledgers, and keeping messages secret, and whether or not they have to be validated.

So the motivation for the first part was the fact that some of these villages lie far away from the central monastery. Some are very near. Some households work as serfs, slaves. Others far away pay rent. So we wanted to explain this kind of variation in the organizations that we see. So we introduced a simple two-agent model, where only one of the two, agent 1, had some privately observed endowment. Agent 2 was the villa or the monastery. And we wanted to know what would be an optimal resource allocation mechanism, granted that agent 1 has private information.

So we did some derivation, had agent 1 sending messages, looked at what would be a maximizing message, derived a series of inequalities, and then changed the problems to sending messages about the underlying endowment, and ended up with something we call the truth telling constraint, which captures all of the problems associated with private information in this context. So then we're able to solve for optimal resource allocations just by adding on these kinds of incentive constraints, in addition to maximizing the lambda weighted sums of utilities.

We did this in levels. Then we went to lotteries to make it more general. We had a good discussion at the beginning of one of the classes of the advantages of these lotteries, which we had talked about in the first one third of the class as well. Oh, and at some point, there was a starred article on the reading list, which was how to actually code these things up. It's more related to the next lecture. But anyway, trying not to be abstract about it.

So then we went from one period to two periods. We solved the two-period problem, and had incentive constraints, truth telling constraints for both periods. In this case, what I was emphasizing both the dynamic program and how to work back from the second period to the first period sequentially, but also the fact that moving forward from the first to the second, the message is about underlying unobserved states would be key to implementing the optimal information constrained solution. And by emphasizing messages, I was already anticipating contracts, smart contracts, written on ledgers where messages replace transactions.

And the optimal solution to that two-period problem was neither full insurance which wouldn't be feasible, nor limited borrowing and lending which was not optimal. And again, in class, we had a very good discussion about this hybrid borrowing and lending scheme.

We did costly state verification. I won't dwell on that, except that over a range of shocks, you don't need to be verifying anything at all. But it's still effective, because there's the threat of an audit for low claims. And we went back and argued that costly state verification can help explain what we saw in the initial part of the lecture.

So here is where the transition happened to distributed ledgers again. And there was a lot of material here. I kind of argued that encryption was one third key element of distributed ledgers. We've seen it historically. We see it in contemporary applications. It really breaks down into encryption, of incoming and outgoing messages, of hashing of data and others so that it's uninterpretable unless one wants to prove that it was the original, which is possible to do in crypto graphic puzzles, which are hard, but not impossible to solve.

Bitcoin is using various of those things, including a way of organizing the data through mining as the validation. And I emphasized that mining was not the only way to validate transactions on ledgers, if you don't trust Safaricom or some third party. And I also emphasized the encryption part, which Bitcoin doesn't use, homomorphic and that's not quite fair. The encryption part for the messages is related to homomorphic encryption. But Bitcoin doesn't use multiparty computation.

So then we came back to the essence of the lecture, which was, what is the difference between validation on distributed ledgers and incentives and mechanism design, and the different philosophies of economics versus computer science? Computer science assumes that most nodes are honest and reliable, and designs algorithms on the assumption that the honest nodes will follow the algorithms, and just have to have enough total nodes relative to the faulty nodes in the population for those protocols to work.

But if you take the mechanism design point of view, the protocols may not be incentive compatible. Economists assume everyone's very selfish and maximizing. So even the miners in Bitcoin have to have incentives to follow the prescribed protocol. And that point is true more generally. So that kind of led to the idea that why not separate the two things, write down the contract or mechanism design problem in code as in a smart contract, let the parties agree to it and make sure that it functions the way they think that it does, and then send messages to implement it? Those messages would be encrypted through HE and MPC. And that's a contribution.

The mechanism design literature assumes we have a planner. I hate that language. And the planner is trusted to implement the mechanism. What encryption does nowadays is to get rid of trusted third parties and planners. Agents can keep secrets from each other and implement the optimal contract.

The ledgers still play a role. I mean as in borrowing and lending, you can sequester the collateral and have it automatically transferred to the lender if the borrower were to default. So we're not claiming that we only need the cryptography. We also need the ledgers.

On the other hand, you don't need to have all the line items of the code in the ledger being done, as with Ethereum, on a distributed ledger, because that's really, really costly. So again, there was a lot in this lecture. Questions about lecture 11?

Lecture 12 was the application of contract theory. And we took it to the Thai villages, and looked at occupation choice. There's a couple of things going on here. One, we now know how to incorporate private information, and incentives, and collateral by just appending extra constraints onto the problem. So let me start skipping a little bit.

Will the entrepreneur take appropriate action? Not necessarily. If it's not seen, then you have to induce it. Will the entrepreneur pay back the loan? If not, you have to have collateral, which is another kind of constraint.

So here we have the household's utility in consumption, and a disutility of effort, and potentially risk neutral, potentially risk averse. And households vary by talent, wealth, and schooling. And we have they're running a firm with stochastic output as a function of effort and capital. And they can borrow. This is the slide I was kind of racing to get to. So here is the problem, which we posed as maximizing the utility of the borrower or entrepreneur, subject to the bank making zero profits.

So it's worth emphasizing that in the previous lecture, as we've been maximizing lambda weighted sums of utilities subject to resource constraints, and incentive constraints, and so on. Well, here we made a shift from maximizing the utility of one party, subject to a fixed utility of the other. Or here, the utility of the entrepreneur subject to a fixed and zero profit level of the bank.

One could have solved this problem the other way around, to maximize the profits of the bank subject to a fixed reservation utility on the part of the borrower. You should get the same answers. Of course, there are many, many solutions-- if you allow profits to be positive, if you allow some of the surplus to go to the bank, et cetera. In fact, the borrower could be pushed down to the point that they're almost indifferent to participating if the banks are getting the bulk of the surplus.

But what I want to emphasize is it's all the same in the end. Maximizing lambda weighted sum of utilities is the same as maximizing the utility of one type, subject to fixed utility of the other, and then varying those.

So the rest of it was just implementing the solution, limited liability, moral hazard, with separate equations. And then we got to this diagram where I said, if there were no constraints at all, that would be like the full information, full commitment solution, hit the bull's eye, it could be entirely moral hazard and not limited liability. It could be entirely limited liability and not moral hazard.

Although in this case, if you had both constraints, you have two equations and two unknowns, the capital and the effort level. So when those two constraints are satisfied, if they're both in the program, you get a solution that has nothing to do with utility maximization, because there's only one point in the input space where both constraints are satisfied. But in general, if you have only one or the other, you have both max utility and then subject to constraints.

The rest of this lecture was about implementation, namely solving linear programming problems and trying to pick parameters in such a way that maximize the probability that the model is true, given what we see in the data. Maximize the probability that the data we see in reality could have come from the model by choosing the parameters of the model. That was the application. And then we had the findings that the constraints varied by region. Are the questions about this lecture?

So lecture 13 is also loaded. Although it's more recent. I'm somewhat comforted, just like the first review session, by the time we get to 13 and then 14, 15, you've already seen 15 tonight. Today, we just reviewed it. So we're actually getting close to the end here.

This lecture did two important things. One, the theory which was to define the notion of Walrasian equilibrium within and without transfers. And the second was to apply it. So we get general equilibrium solution in a small open economy with production. All of this slide had to do with, remember what an economy means. This is partly review of the first one third of the class-- numbers of households, numbers of firms, what consumers do, what firms do, aggregate resources.

We went over private ownership. And again, we had a very good discussion at the beginning of the following class to review this slide in terms of individual endowments and ownership shares. And the key definition of Walrasian equilibrium with three items in it-- maximizing profits for firms, maximizing utility for households, and the feasibility of the allocations that result from those two optimization problems. The key being that prices are the things that allow the reconciliation.

You can do partial equilibrium. We did. You can solve for profit maximization taking prices as given. You can solve for utility maximization, taking prices as given. Here in general equilibrium, we are only allowing certain prices where those decentralized optimization problems are consistent with each other in the sense that supply is equal to demand.

Oh, and then this Edgeworth box, which is like economists' favorite environment, a sad statement, putting the two types of households on the axis and talking about what a competitive equilibrium would look like. And the price equilibrium with transfers and the welfare theorems, which so far I have only stated, although the last one third of the class next week will be focusing in part on these welfare theorems starting next week.

So then we did production in the small open economy. And we again went through this algorithm pretty carefully. So I don't think I want to say too much about it, too, too much. But we simplify the problem down into firms hiring inputs and trying to find the factor prices that clear the input markets, taking out external prices of goods as given.

So we had the firms maximization problem. We wrote down that the sum of the input demands equal the endowments, and reminded ourselves of Shepard's lemma. So we had a different statement of the input demands, which we could then solve, especially in the 2 by 2 model. Input demands are entirely a function of input price ratios. And with constant returns to scale, it's even easier. We can talk about cost minimization to achieve a unit of output, and then to scale up and down when we get to determining output.

We talked about the intensity of goods, the factor intensity of the different goods with a definition. Here was the cost curves and their properties. Concave, we solve for the equilibrium factor prices, where both firms face common prices. Hence, there's only one point where that's true on this diagram. And then we went from that to production. From production, we were able to solve for the input quantities and the outputs. And then we got the external goods prices by considering the free trade equilibrium with two countries determining a price ratio moving from Autarky to international trade.

So there were three theorems. We had this Heckscher-Ohlin theorem, which someone mentioned at the beginning of class today, that the returns to factors depend on whether the country is abundant in one factor versus another one. The Stopler-Samuelson, having to do with how factor prices are moving when we vary the price of outputs. And that's all.

OK. Are there questions about this lecture?

AUDIENCE: Here the HO theorem that it says that in the Autarky, the price of the capital intensive goods in the capital abundant country will be buy down, so here buy down means the relative price wars in this capital-abundant country was lower than another country, right? Not the absolute prices?

ROBERT TOWNSEND: The good that is being produced, which is capital intensive, if produced in the capital abundant country, should be relatively cheap, because the other country is abundant in labor and scarce in capital. So--

AUDIENCE: Yeah my question is that these are cheap means that the relative price was cheaper. I mean in this capital-abundant country, the relative price of these capital intensive goods was lower than the same relative price of the capital intensive goods in another country, like labor-abundant country, right?

ROBERT TOWNSEND: Yeah. When you look at this-- another great question, thank you. So when you look at this diagram, it's really only the relative price of the goods that matters. It's like p1 over p2. So there's a lot of things that are relative here that look like they're in Heckscher-Ohlin stated as absolute prices. But it's really only the relative price.

AUDIENCE: I get it.

ROBERT TOWNSEND: OK. All right, so that was lecture 13. And the other two lectures, 14, we took that trade model to Thailand and did some more with the financial accounts to get not just individual income and balance sheet statements, but to aggregate up those statements to get to village asset accounts. Village level savings accounts, flows across the villages, and so on.

But my time is up. So I can't review it. Of course, the other thing we're doing is taking that 2 by 2 model to actual data with the exception that the model here actually did not have constant returns to scale. It had capital mobility, whereas the 2 by 2 did not. Factor prices would be equalized, even without factor flows. The thing that's mucking up the Thai data is that there's a borrowing constraint. So we have both real and financial flows, subject to borrowing constraints.

And lecture 15 was what I reviewed at the beginning of class. Any questions about 14? OK, well I hope you find these reviews useful. I must say, you should all feel really good, because you're doing well in the class, and we are covering a lot of material. The constant themes here are theory, data, both, lots of tools, lots of concepts. By now, you're functioning like actual real economists, and the conversations on Piazza reflect that. So I'm really happy with how the class is going. And I hope you found this review to be useful, and to be a reminder of how much you have learned.

All right. Thank you very much.