[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT TOWNSEND: OK, we'll get started. Thank you for coming. Today is lecture 4, "Smart Contracts as a Solution to a Coordination Problem. It's still about distributed ledgers, but we're going to add on top of that not just information use, but also contracting.

The longer subtitle, commodity space with location dates and states. The policy objective is the same. The Pareto criteria, we will also have fragmented markets the way we did last time. The implementation is going to-- problem is going to have something to do with privately issued securities, which are circulating around as monies. And without the ability to coordinate with their smart contract, there will be market crises. And so we're featuring this multi-agent smart contract.

So a little more detail. Extensions of the theory that we started to cover last time, but this time, explicitly about time and risk. And then we'll talk about getting policy guidance from the data. As I said last time, I'm trying to put discipline on myself and the research, that we shouldn't just be exploring case studies or potential applications that potentially use these new technologies, we want to know what economic problem exists out there and whether or not we can help mitigate the problem.

So we're going to use this efficiency benchmark that we did last time, and I'll show you how to use it with actual data, although that's a bit of a summary. And then we'll, like last time, do some policy guidance that comes from theory. What's the logic according to the theory of what can go wrong when you have privately issued circulating debt? The idea is not to get rid of the debt, but to provide some coordination to avoid these financial crises.

And this comes up in at least four seemingly distinct venues, but they have in common this underlying theory. So again, I think you've seen the pattern here, which is, take the theory seriously. Even when it was designed and written without these kinds of considerations in mind, there are huge lessons to be learned from going back to it and thinking about the new technologies.

So again, it's about multi-agent smart contracts on a common ledger, and in particular, at the end, we'll talk about Ethereum and what that is relative to Bitcoin.

OK. So extensions of the theory to incorporate risk and intertemporal exchange. So we had apples and wheat. It's easier to think about borrowing and lending as in getting-- or giving goods up today and receiving goods tomorrow, so the time date is there. But in addition, we're going to have mother nature drawing rainfall and weather realizations. So we need to talk about the state of nature, and that's that little s t.

And actually, it could be the entire history from the genesis of states that have been realized over time. So we index commodities by time and by this history, including contemporary realizations. So little t superscript denotes the history, sub t denotes the contemporary value.

The utility of the agents is going to be-- I'll say it in words, discounted expected utility. So we're going to add up the utilities over the time horizon, and we're going to take expectations over these underlying states of nature. So it's discounted by beta, and the expectations are formed with these probabilities of states of nature.

And then we're going to, with this notation, do what we said we were going to do last time, which is turn this into a programming problem. So we're going to maximize this lambda-weighted sum of the utilities of the participants subject to resource constraints.

The words are easy. The notation is a little more demanding, but lambdas are the weights, here's the discounted expected utility summing over dates and states with that probability distribution. And this resource constraint says that consumption in a pure exchange economy, consumption should equal income.

So we're just summing up over all the, quote, "endowments" of the finite number-- capital I, it is-- of participants. That's the resources available for, say, distribution. So consumptions can add up to more than that.

Again, a very powerful tool, and I'm happy to elaborate on this. We set this max problem up as a Lagrangian, so we kind of repeat the objective function. And then we have these auxiliary terms which are shadow prices times constraint. So we have the resource constraint written out for each date and state. For every date in state, it's preceded by this shadow price, this Lagrange multiplier, theta, of s of t.

And the very nice thing about Lagrangians is as long as we have concavity of the objective functions and convexity of the underlying constraints, the first-order conditions of this max problem are necessary and sufficient for characterizing an optimum.

Here are the first-order conditions. And it simply states in words that the lambda-weighted discounted marginal utility of consumption in a certain history of states should equal the common Lagrange multiplier across all the households. For every date and state, we're equating-- we're reallocating consumption in a way that equates the weighted marginal utilities.

So this is-- and we repeat the resource constraint. Just a word about the math, and I'm trying not to do too much of that, there's a lot of these guys for every date and every history of states, and we're writing it out conditionally, but there are tons of them-- of these constraints, tons of Lagrange multipliers.

So this just says weighted marginal utilities for person 1 should be equated to any other person i, repeat the resource constraints, and we get, as a solution, something that looks deceptively simple, which is that the consumption allocation for an agent type i can be captured by a household i specific function g i, which has, as its argument, only the aggregate income.

Now the intuition for this is that we're operating kind of a mutual insurance society. And it's as if everybody were contributing their rice to a community fund, and then that is pooled together, here's the big pile of rice. And now, maybe under the auspices of the monks, that is handed back to households depending on whether their own harvest was high or low.

So the implication is, individual harvests are in there, but they're in the stockpile. And once you control for the stockpile in terms of its aggregate income y, the individual harvests don't enter. So you have this surprising result that consumption should depend on the aggregate and not on the individual shocks. And that's, as you'll see, a benchmark, which is taken to data.

There is another implication, because of the concavity, that you'll have co-movement in consumption. People's consumption should go up and down together. Not everyone should eat the same amount. Some will be eating more, some less, but if you track their consumption, it should never cross. You should never have any period where consumption of one household or region is going up while consumption in the other region is going down. And that's testable.

So again, we're working on policy guidance from the data, and the idea is, OK, that's the implication of the theory. How well are people doing? Maybe they're doing just fine. That would be wonderful. It sometimes happens. Or in other instances, things that look terrible, is there a remedy? That's the simple logic of it.

So we're going to take a peek at village India, village Thailand, and come back to this idea of targeting based on the theory more generally.

So this is a picture of villages in India that were sampled for 10 years by a crops institute where the economists and business folks were not telling households how to grow crops, but they were gathering the data. ICRISAT data. And here, you can see the year, 70-- 10 years, basically. The numbers here are 76 through 84. And here are household numbers.

And for every household, we're tracking the ups and downs of income. It's normalized around the overall average. So a negative number doesn't mean it's negative income, although incomes could go negative, it's just that you're below the average.

So you see this gradient here. This has to do with landless small, medium, and large landholders the way they were assigned numbers in the data set. And not too surprisingly, large landowners are eating more-- or have more income than these landless laborers.

When you look at the income, I sometimes refer to this as the Rocky Mountain picture because the peaks in front conceal what's going on behind the peak, which is arguably a valley. So incomes are not co-moving at all. They're very erratic.

But when you look at consumption, in this case, of grain, which is sorghum, which is the dominant item in the diet, it looks-- forgive me if you're from Kansas, it looks like Kansas, it's really flat. There's not much going on. And yes, I've shown these pictures before. I've had a student in class who said, "I'm from Kansas, and you're right, there's not much going on in Kansas."

The point here is how flat and a bit undulating waves of grain consumption are relative to income-- they are deliberately put on the same scale. I could blow this up if I change the scale, but the main message here is consumption is smooth. And then we can go to the data and see whether that more stringent test is consistent with the data, or whether we fail the statistical tests, and if we do, by how much.

So roughly-- hard to draw one number from this. If it were in dollars, not rupees, but for every dollar increase in income, there's a $0.07 increase in consumption. So roughly only 7% of the income fluctuations are making their way into consumption.

It's different for different groups. If you look at profits from trade and handicraft, then it's almost up to a quarter. Profits from animal husbandry in the two other villages, Shirapur and Kanzara, are, again, less than 25%, but close. So we're already into the, quote, "targeting aspect" of trying to identify sectors where the smoothing is not only-- overall it's not perfect, but the benchmark came surprisingly close.

This actually created a firestorm among policymakers. They couldn't believe it. Their presumption was that households in a poor village in India, replete with the traditional caste system with the Brahmins and the untouchables and so on, there's no way they could be doing anything close to this well, but it is a statement about what's going on within the village in terms of storage of grain, gifts and transfers, and so on. But it isn't. It gets a little further away, and these guys suffer from more risk than some of the other groups do.

OK, so I showed you the theory when there was pure exchange, as if endowments were-- income were given. It is possible to extend it to include production. So for the Thai villages, we have the output produced with capital inputs in sector I household J.

And there, the idea is a bit unrealistic. You can always undo your projects. You have some of the consumption, you set it aside as investment, you acquire machinery and equipment and so on, that's helping with production. At the time of production you have, say, the harvest or some other revenue, and you decide how much to eat and how much to save. So the idea here is you get to eat potentially everything that's available, which is the output from these projects, like agriculture over something over all households, and the capital stock.

But you can also reinvest, so it's not all consumed, but some residual is invested in the projects all over again. So sources and uses of wealth, you might say.

Now here, we're focusing on the production aspect. And then the key issue is how much risk does a producer bear of the fluctuations. And the aggregate would move everybody's incomes like before. And if you're restricting attention to what's going on within the village, there's not much you can do about that. Someone has to bear the aggregate risk.

But the idiosyncratic risk that hits one person's land and not another's, that's potentially smoothable in this mutual fund. So when we talk about a risk premium, it has to do with a financial object, the rate of return on these projects as if they were stock, equity. So the rate of return having to do with the net profits you get from the production activity divided by the capital stock.

If they're bearing some risk, they need to be compensated for that. The rate of return needed would be higher than what it is for a risk-free asset, or just holding inventory or money. And that can be tested. The first-order conditions are equating weighted marginal utilities.

That's an implication we already had from the previous slides. And then of Euler condition, which is telling us intertemporally how much to eat and how much to save and where to put the savings. So this is like an efficiency condition.

And it turns out, just looking at the investment and rate of return data, the percentage of diversified, idiosyncratic risk is, on average, in the 90s. So it's like some of it, maybe 10%, requires a higher rate of return, but 90% of the risk is gone. They've pooled it away in this mutual fund.

And another implication is what's aggregate has to be borne, but if you look at villages one at a time, the aggregate is not common. So it's as if each village had its idiosyncratic risk, and if you had a larger mutual fund, which they didn't have, could imagine a building a platform that would smooth risk across those villages. So a within-aspect and an across-aspect. Yes, Henry?

AUDIENCE: So for the diversified, idiosyncratic risk that's above 100%, for instance, for the third village, is the interpretation there that there's some consumption that goes to perhaps making up for some lost crops or something that isn't actually like personal consumption that's meant to improve utility, but some sort of investment? Is it like mismeasurement? Otherwise, how should you--

ROBERT TOWNSEND: I think it's mismeasurement of four. Buriram had experienced an enormous amount of transitions. It's sitting right on the Cambodia border. There were incursions, for one thing. And then they were growing rubber and they were making a lot of money, then there was a huge amount of Danish tourists that go there to retire. So things weren't stable. And so I haven't tried to actually come up with a particular reason for the higher numbers, I've just attributed it to data, but good to point it out.

OK. So just in case you're a little bit bewildered, back to the goal here. We have a criterion. I'm trying to illustrate going to the data and testing it, and then having guidance for policy. And there's a lot more to talk about today, but these slides come, for example, from other work.

I've mentioned, from time to time, 193 or the second half of 772, and there, we go over in much more detail the theory of risk-sharing, these two papers, and also what I began to allude to, the building of a digital risk-sharing platform.

But you don't have to wait to gain access to that material. I've already reviewed the lecture slides iterated with my administrative assistant, Deb, and those slides are ready to go. So we can post them if you like, and you can read ahead and-- well, read in parallel would be a much better way to put this.

All right, so the last word on the targeting is, it's not just Thai villages, although that here has to do with the networks within the village distinguishing the risk-sharing network from those that aren't in it, and the exposure to idiosyncratic shocks and how they channel through these networks.

But you could look at the whole country of Thailand-- and Angus Deaton, who's report here doing this for Cote d'Ivoire, you can see these cross-regional decompositions. It's not a developing country thing exclusively. You could go to the US. These workers beta-- betas, finance terminology, having to do with co-movements, distinguishing various groups using US data.

And this paper from the EU has to do with the countries in the monetary union are not in a fiscal union and their movements of consumption. Well, they don't co move much, so there is gains to thinking about having a better unemployment-- cross-country unemployment scheme for the eurozone.

OK, so now we get to policy guidance from theory. And we're going to focus on high-velocity private debt, or, for short, privately issued monies. A little more provocative to put it that way, although I will say, immediately, this is the monetary theory definition of a money as an object-- in this case, a security that appears frequently in exchange. Hence, privately issued securities are candidates to be money depending on how often-- what their velocity is.

And we'll look at these bills of exchange after we go through the theory that were being used in London. I have touched on that already-- in England and traded in London. We'll look at the first welfare theorem again as a review. View will attempt to implement the competitive equilibrium, but with fragmented markets and limited securities. And we're going to bump into this huge coordination problem.

So English inland bills of exchange-- so this is within England-- were used as a medium of exchange during the Industrial Revolution in the north of England in the 18th and first half of the 19th centuries. So if I say privately issued monies, and the comeback is, we only have fiat monies, that would be wrong. Historically, we have seen these privately issued notes circulating as media of exchange, as in these inland bills. And we'll come back to that.

So what is the welfare theorem? And this is a review. We have a candidate star allocation and prices p, and we have, say, a way of raising equilibrium. And then under certain conditions that, while raising an equilibrium, is Pareto-optimal. Despite the dates and the states and the risk and so on, it's the same theorem as it was last time.

So we are going to, as a continuation from last time, retain these fragmented markets. It's as if traders are meeting pairwise, or at least in subgroups. But we're going to extend this to the dynamics, which are now going to be explicit, as in maximizing discounted utility over time.

And the agents are no longer just code where we were trying to figure out how much information would be needed as an input into the code to achieve an objective here. The agents are self-interested, they're maximizing.

And it's always a bit uncomfortable, at least for economists, to think about agents that don't max something, so now we're in more familiar ground. But we're going to still need the computer science to talk about the code and how to correct the problem that will become clear when we get there.

So the focus-- the economic focus is on asset issue and ownership change of those assets, and I've already talked about money.

So this is a paper I wrote with Neal Wallace. Each trader-- here's the environment. Traders are endowed with a single good, it's an endowment economy, but there are multiple dates and states. Or at least multiple dates. And their endowments vary over time. They're going to meet in these trading venues, which are segmented markets or platforms.

And when they meet, they're only in one place in a given date, but different households are in different locations. And they will try to maximize their utility, which is derived from end of period of consumption.

Obviously just intuitively from the get-go, if your income is low, you might want to be a borrower or sell securities that you've already acquired to get your consumption to be higher than your income and vice versa if your income is high. They're all interested in smoothing the path.

So agents are going to maximize intertemporal utility at each date, taking into account the future, including the future prices of securities. Yeah, and we're going to get rid of the states of the world, so it's entirely about time variation.

So here's a-- not just an example, really the prime motivator for what we're going to do, is two locations. Again, agents meet pairwise. So agents 1 and 2 at date 1 are in location 1, agents 3 and 4 at date 1 or in location 2. In this case, trader 1 stays put in location 1, trader 4 stays put in location 2, but the other two guys numbered agents 2 and 3, they're flipping back and forth, back and forth.

So, again, for Ostroy-Starr, these would be the pairwise matchings. And it's deterministic, the sequence is known. So I'm going to focus on-- well, let me do it now. If you notice, agents 1 and 2 are together again at date 3 as they were at date 1. So depending on the endowment profiles, you can imagine that 1 or 2 are borrowing and lending with each other depending on who's high and who's low and when.

There's a lot of that going on. 1 and 3 are together, and 2, 4, and there's symmetry over here at location 2. So there's the potential for these debts, but in this case, they would be the more familiar debt where the borrower and lender agree to the loan, and the loan is to be paid back two dates later, not passed to third parties.

But, in addition, you have these chains going on where one issues a security, which is a promise to repay at date 4. And 2 could buy that and then trade it to 4 at day 2, 4 could trade it back to 3 at date 3, and 3 could cash it in.

Now let's just suppose, for the sake of argument, that there's no reneging. These debts have to be repaid when they're presented, otherwise, you go to debtor's prison which they had, severe penalties. So we abstract from that.

Again, without too much by way of apology, because the more we simplify and yet still encounter a problem that needs a solution, the more clear it is how pervasive of the problem could be. So no reneging.

But it's also true that 4 could issue this stuff to 3, 3 could issue it to 1, 1 to 2, and 2 redeems it to 4. So you're getting the idea here. We can look at these chains that connect the issuer at the issue date to redemption of the issuer at the redemption date with the change going through other parties. And those create the possibility of these circulating privately issued IOUs.

If we just looked at the first two dates and restrict ourselves as we are to these kinds of debt securities, then we're kind of in trouble. Why? Because the debt cannot be redeemed. You could imagine, one still issues this note due at date 4 to 2-- sorry. 1 borrows from 2, 2 carries it to 4, but that's the end of the world. If there's only two dates, so it cannot be redeemed.

Now this is interesting because it means the simple debt equilibrium gives us autarky, which is a horrible outcome. So the restriction to pairwise trading and to limiting ourselves to debt security trades can be quite damaging. If we had allowed the complete market solution with these four guys, you could potentially see some trade even though it's not bilateral. It's like 1 borrows from 2, and then 3 provides the flow going back in the other direction at date 2.

So it's a four-person multilateral agreement if we allowed, as the theory would have suggested, a decentralized competitive equilibrium in state and date contingent securities. But we're ruling that out for now. But we'll come back to it.

It is pretty much in the spirit of Ostroy-Starr, in the sense that not everybody's together all the time-- they're in this fragmented, spatially separated market, there's quid pro quo requirements as per the competitive equilibrium. They have a budget constraint, and then the issue is how much can they achieve with these kinds of sequential restrictions?

It's multilateral. Ostroy-Starr was not bilateral. It was about what the community can do as a whole even if the actual trading is restricted to pairwise. OK. That's what this says, basically. If you look here, and restricted it to two dates, you would see that the objects that agents care about, which depends on the location and date, is never common.

Agent 1 cares about location. Agent 1 cares about location 1 goods, 4 about location 2 goods, and the other guys keep flipping places. So there's no common date-state pairing if the economy ends with only two dates. Obviously with four, more is possible.

So these rules, traders can issue securities. They have to be redeemed. We allow short- and long-term debt. We allow circulating debt. Traders can buy previously issued assets and then sell them for value. This is like the quid pro quo or value for delivery swaps at any given date and location.

It's as if, though, that they carry these paper securities around, but goods can't move. Again, that's a restriction. We're not having endowments or carrying of the commodities around. That would change the outcome, they could do some smoothing that way. Here, they're restricted to use the securities for the smoothing.

OK. So if you think about that economy, we could talk about payment matrices and velocity. For payment matrices, I have shown you versions of this earlier. For simplicity, just imagine we have rows and columns in a big square matrix, and the objects consist of these consumption goods, but they're indexed by dates and locations. So you have eight of those guys because there are two locations and four dates.

There's also the potential of these non-circulating debts, and they can trade against consumption. If, in fact, each consumption good gets traded for one circulating security and one non-circulating security, then there's a non-zero element in each row corresponding to a consumption good or a non-circulating debt.

And for circulating debt, there are basically four checks because that debt could be traded in dates and locations and it's circulating around. So that's one way to think about the circulating stuff appearing more frequently in exchange. It fills in more boxes.

Velocity is related, but not identical. So velocity means the amount traded in a given date divided by the stock outstanding. And we take the average over all dates where the security is in existence.

So for that same table economy, a consumption good at location i at date t, the stock outstanding would be the endowment, and the amount traded would arguably be something less. This has to do with excess demand and excess supply. Rarely is it the case that if you're endowed with it, you give it all up to somebody else. So it's always a fraction of the total. So that's a number less than 1.

If you consider the non-circulating debt, it comes into existence when the borrowing-lending contract is signed. All of it is traded-- all of the stock is now held by somebody else. They hold it for one period, nothing happens, and then it's redeemed. All of it is redeemed. So that has a velocity of two-thirds because there's three dates and it's completely traded in two of them and not traded at all in the intermediate date.

And likewise, the circulating debts, of it is traded every date, so it has a velocity of 1. Either way you can think about these circulating debts as monies that appear frequently in exchange and have a high velocity relative to other securities and relative to consumption.

So it turns out to be useful to create more general notation that would cover additional economies, not just the one in the table. So let J be the number of locations, T the number of dates, the commodity space has dimension J times T. Same idea here. People get utility from commodities. They have positive endowments, but only in the locations where they're in residence.

If there are capital G people and they're living for these T periods, each person, little g, can be paired with another person or no one. Remember, in Ostroy-Starr, we had that a stand alone because trade was pairwise and there were five traders. So here again, we can allow these isolated pairings of singletons and isolated location.

So you're either assigned to be alone or you're assigned with someone, and that's the itineraries. There's one commodity still for every date. And location, goods are indexed that way, and there's no transportation, production, or storage of the goods, but of course these assets can be carried around.

So if we have person g in location i at date t, that person, if they are in that location, would have an endowment w g of it of the consumption. If you thought about a vector over all dates and states, they would have, say, positive endowments in the locations and dates that they visit and 0 elsewhere because they're just not there.

Likewise, for consumption, c g, c g i d being a typical element of non-negative consumption for person g in location i at date t. Well, now what I promise is coming back with a vengeance in the notation. We're indexing everything in sight by locations and dates, and then we have these people. i location, date t, and g for the person.

But they get utility from the consumption vectors, and now we have these securities, which would be the number of units of consumption promised by person f issued at date s coming due at date t. And we looked at this from the point of view of agent 1, for example, issuing a promise at date 1 coming due at date 4. And this is the general notation.

But we're only going to consider issuers f at date s coming due at t for paths or chains which allow this kind of interconnectedness. That hopefully was clear from the table. So when there is a chain, which is active on both the supply side and the demand side, we can determine a price. So we'll have this price of security f issued at s coming due at t, but priced at a different location i and some date u. Yeah, so sorry, we're up to five superscripts and subscripts and objects, p of st at i, u.

But where we don't need to price them everywhere, if traders h and g, say, are meeting at location i at date u, h may not want to acquire that security, there has to be some route from when agent h could acquire it, either directly presenting it to the issuer later or through other people who, in turn, are connected to other people who, in turn, are connected to the issuer at the redemption date. So that's on the, quote, "demand side."

Likewise, who's bringing this security in? They must have been either directly connected with the issuer at date s or connected with someone else who was connected to the issuer from whom they have acquired securities which they are now basically supplying.

So you can capture these rules for security markets and trades. The no redemption rule says that basically, the issuer is like a demander and a supplier, they're a supplier initially when they issue the security, but they also have to demand it back. That's the no redemption aspect. And they have to demand back everything that they issued.

And these are the notions that you have to acquire an asset before you sell it. Somehow, the English is a lot easier to say than the notation is to understand, but that's what this means, that the demand is non-negative when you're summing over previous histories. You have to demand it first before you can supply it.

And then this is key. This is the budget constraint of a person g at some location i and date t, and this is his or her endowment, and this is consumption. So if you subtracted consumption off from the left-hand side-- from both sides, you would see the difference between endowment minus consumption on the left. So that would be excess supply if it's positive. The amount that you own over and above what you eat, so you're a supplier of the good at that location i and date u.

But what are these other objects? These are the valuation at these prices p of the security trades. So you could take this excess value and basically buy securities. Or alternatively, when it's excess demand, how do you accommodate excess demand? How do you manage to eat more than you have? Basically you have to sell securities you previously acquired.

So when you're running surpluses, you're acquiring securities. When you're running deficits, you're selling securities. And that's what this budget balance equation is saying. But again, you're not-- we're not allowing direct commitments for location and date-specific goods, we are only allowing these non-circulating and circulating securities all captured by this d f of st notation.

So a securities equilibrium-- a debt equilibrium would specify consumption and debt demands for each person and prices such that consumption in the debt demands are maximizing utility by date and location subject to the debt trading rules, and we need market clearing.

And market clearing is just saying that the excess demands of consumption sum up to 0, but that also the security demand is equal to the security supply for every potential circulating or non-circulating security. This thing is summing up over g over both of them, over the total number of agents.

So again, apart from the struggle of absorbing all the notation, it's very similar to what we had before in a competitive equilibrium, except here, we are doing it with securities.

OK, so finally, we can look at our target, and the target is going to be as if we had the complete markets equilibrium all along. So remember, in Ostroy-Starr, we had the excess demands and prices that corresponded to a competitive equilibrium, allowing multiple goods.

Here, goods are indexed by locations and dates, so we could imagine a competitive equilibrium in date- and location-specific commodities where we price them all out initially at some initial date, and agents can trade them there, taking prices as given, maximizing utility.

And in an equilibrium, prices will be such that excess demand sum to 0. So those are the words of Ostroy-Starr and Walrasian equilibria in the context of this location, date. So these excess demands are consumption over endowments of person g at location i and date t. The sit are these unit of account prices-- you have to pick a numeraire. Maybe it's the ghost currency sit, the price of consumption at location i and date t.

And every person is going to have their budget constraint, which is that when you sum up over all the goods, sum up over locations and dates, the valuation of excess demands is 0 for each household, just as it was before. And also, the sum of the excess demands would be 0 when we sum up over g over the households. Yes?

AUDIENCE: The prices we had earlier, were we pricing goods or were we pricing the debt?

ROBERT TOWNSEND: Debt.

AUDIENCE: So those aren't the same?

ROBERT TOWNSEND: Not the same.

AUDIENCE: So you want to convert--

ROBERT TOWNSEND: Yes. Exactly. So we're going to get to exactly that on I think the next slide, but first, just a statement. What is the relationship between this complete market, date-location market initially and its equilibrium with the debt equilibrium?

So a debt equilibrium would consist of those debt prices and non-negative debt quantities such that the debt quantities and our target, complete market's excess demands, satisfy each person's debt budget constraints, which I wrote down on the previous slide, and the debt quantities and these excess demands are utility maximizing for each person given the debt prices. So that's what it means to achieve the date, location, complete markets equilibrium by way of these debts.

So here, I'm kind of reminding myself not to get too carried away on each one of these slides. I just want to tell you what's going on, and then if you like, you can look at the algebra later.

So this, Rafa, is the answer to your question. Let's guess what these debt prices ought to be. The debt prices are telling you how much consumption you get when you issue it relative to redemption, and redemption is one to one. You have to honor the number of consumption goods you promised in the d f of ft.

So you could think about the ratio of these complete markets' prices as how much consumption you're getting at redemption dates relative to the issue dates, and hopefully that's enough without getting too involved with it.

So the relationship between, say, the security issued by agent 1 coming due at date 4 traded at location 1 and date 1 is this Arrow-Debreu-- Arrow complete markets price in unit of account at date 4 relative to date 1. So that's the guess. Which means at the guess, we can fill in all the debt prices, and then we see if it works or not.

It's a very large system if you consider all these people, dates, and locations. Fortunately, we don't have to solve explicitly for all of the budget constraints. If we judiciously choose five of them, we'll be fine.

So there's two statements on this slide. One is important, which is that using our candidate for the prices and the target excess demands, and then we sum up over the budget constraints, but we're going to have those prices filled in here. We sum up over these sequential budget constraints with the debt prices filled in at the Arrow-Debreu unit of account, prices we can show that we will get out of this the complete markets single budget constraint.

So if we are able to come up with debt quantities that satisfy these constraints sequentially, we automatically have maximization because the complete market specification, they were maximizing utility there, and if we can find debt quantities that satisfy these sequential constraints, then they can get that utility level.

And then the next part is, we don't need to solve this for every person, every date, and every location. There are, in fact, 32 excess demands, but we have constraints for the way excess demands sum up for the location and date, market clearing conditions, and for places where things are not visited. So we end up with more-- we end up with five things we have to fill out, not 32.

And again, forgive me, but I'm at risk of falling into the trap going through that algebra too much. Let's just say we're counting equations and unknowns, and there are five unknowns. We just have to choose them very carefully. This is a matrix notation with five rows, and the columns correspond to the potential debts, which are being traded. And then we end up with two systems.

We actually did a little linear algebra and took one row multiplied by a constant and then add it to another row. That's perfectly permitted matrix algebra. The advantage of that is the first four rows will pin down the non-circulating debts that are part of achieving the target allocation, and the fifth row is this equation.

So this-- if this equation is satisfied, then we can achieve the target. And this equation has the valuation of excess demands on the left of what would have been the target complete market's equilibrium and a configuration of debts on the right.

So for example, if this valuation of excess demands were positive, all we need to do is issue this quantity of the debt issued by person one coming due at date 4 and ignore all the other debts.

Or, instead of doing it with person 1 at location 1, we could do it with person 4 at location 2. There's two negative signs here, so this is a-- so a positive number here with all the other debts at 0 will also work.

And now we're inching up on the problem. How can agents know in the decentralized world, where they're born into a certain location, what is going on in the other location? And again, it's very akin to Ostroy-Starr where you don't know what's happening, unless it's all recorded on the ledger, what happened in other pairings.

So there's an infinite number of equilibria. The underlying allocations are unique and identical, but the way of achieving them through the debt, there's an infinite number of ways of doing it. They could get it right, but they're going to need to coordinate. To get it right, they're going to need some information about what's happening in the other location.

So this compares to Ostroy-Starr where we said money might potentially work if there's enough of it but they would have to know which object is the money, and the budget balance information had to do with information being recorded on the distributed ledger. So the point this is akin to Ostroy-Starr. It's not different, it's a different context and a bit of a different model. Yes?

AUDIENCE: Sorry. The question went a little too fast. Can we go back to last slide? So this is something that has to hold at every debt equilibrium that implements the Walrasian allocation of this? And how can we see from there that-- what the chains are and how the debt is circulating, and what the possibilities are? So what does this mean?

ROBERT TOWNSEND: Well, this is the stable economy.

AUDIENCE: OK.

ROBERT TOWNSEND: And this-- we don't have a statement, we have an example of where you can actually reach the target.

AUDIENCE: And it satisfies that equation?

ROBERT TOWNSEND: And it satisfies this equation.

AUDIENCE: Yeah. And where are the degrees of freedom? What could you vary now as you--

ROBERT TOWNSEND: This is the given as an Ostroy-Starr. The excess demands are given and the prices are given. And the statement-- that it corresponds to the complete market's equilibrium.

AUDIENCE: So once you pick--

ROBERT TOWNSEND: So that's data, right?

AUDIENCE: Right, that's data. And then on the right-hand side--

ROBERT TOWNSEND: How do we get to this number? Choose.

AUDIENCE: Any one?

ROBERT TOWNSEND: Any one that satisfies the equation.

AUDIENCE: You don't need to specify all the different d's that are possible? You only need these four d's or are there--

ROBERT TOWNSEND: That's true, you don't even need all of them if you-- you can choose one at a time depending on this number.

AUDIENCE: Yeah.

ROBERT TOWNSEND: What I featured was this is, say, positive, then just choose this one. But another is to set it to 0 and choose this one. Right. And I didn't go through the other cases. If this is negative, something like this also happens.

AUDIENCE: And the one that would say minimize-- velocity is always 1, we said, in all of them?

ROBERT TOWNSEND: Yeah.

AUDIENCE: So that's not going to change? OK.

ROBERT TOWNSEND: Yeah. Yeah, all the work that I didn't really share in the algebra is going through all those debt budget constraints and making sure they're satisfied that these gas prices-- so that we are achieving the excess demand of-- which is just popping up here at the end.

So yeah, so just like Ostroy-Starr said need some kind of coordination, either information directly or agreement about what the money is in the case that that works, for us here, you need some information about the security issues.

So the question is, what happens if somehow they get it wrong? What happens in the model if location 1 believes they're the only issuer of the circulating stuff, location 2 believes they are the only issuer? And then this is an outline. We'll come back to Bagehot's London Money Markets, low and middle-income countries, and fragmentation of liquidity in decentralized finance.

So the answer is, if they get it wrong and they're both issuing this circulating money in such a way as to violate the equation, the people who have gotten a hold of the circulating debt in, say, location 1 thinking that they will be able to unload it are going to discover that its price has plummeted when they get there. There's enough information for everyone to realize the mistake at date 2, but it's too late to undo it.

They will start changing their borrowing and lending that's done at day 2 and coming due at date 4, but it's not enough, it's not enough of an adjustment. It's really kind of a slow-moving crisis. Even though the prices change abruptly, it's partially mitigated, as I said. And in the third date, the people who are carrying around this circulating stuff, they just take a big hit in terms of consumption.

And we work out-- Spector and I work out an example with log utilities, constant aggregate endowments that vary, and who has what fraction of it over time, and we're able to compute the solution. This table is a bit confusing because there's so much going on. There's the two locations. Scenario 1 is where they get the debt issues correct. Scenario 2 is where they guess wrong. And here's where you can see that consumption is getting buffeted around, often severely.

Eventually in day 3, you can see, when you have time to study this, that the prices are collapsing because otherwise, in the complete market, the prices are all equal to 1. And here are the utility consequences. Namely persons 2 and 3 who acquire and carry the circulating stuff have a loss in utility. Summed with the units, this log utility is negative, but you want something to be less negative, that's a good thing. OK.

So oddly, also, agents 1 and 4 actually benefit from this, and we haven't thought that through. The impact is heterogeneous given all the heterogeneity in the economy.

So, again, things could go really bad in the model, and the fix is to coordinate on the issues. You could imagine respecting the itineraries-- so people are only allowed to trade consumption when they're in a date and location and paired with someone, but the big accounting matrix is keeping track-- and now restricting especially the security issues. So I'll get-- and now I'm going to get to applications.

But everyone is talking about digital assets, we're going to focus a lot on that in the next class, and tokenization and so on. This risk has not been articulated so far as I know. It's a somewhat nuanced message. It's not like the singleness of money.

It's not the idea that you should ban private securities or the circulating stuff, but rather, coordinate in the issue of that stuff if you know enough about the underlying economy. There's a lot of literature, It's quite the rage now on tokenization and what you can do with dynamic ledgers, but so far as I know, no policymaker has taken on board and talked about this kind of coordination problem.

Back to London. We had those inland bills of exchange. They were circulating around as the money. And the brokers were making money off of this. So you could think about a small-medium enterprise selling stuff. The purchaser has issued an IOU. The small business is suffering from low liquidity, so they go to a bank or, in this case, a bill broker, and for a discount, are able to get liquidity. They're selling the security to the bank.

And the bank is a broker-dealer, so they trade to other people. So this was the London money market, so it's a familiar institution, these kinds of IOUs, and it was a big part of what was going on in England, except that there was a crisis. And, yeah, the model actually predicts that can happen if you don't have coordination over these IOUs. It's entirely plausible that that was the source of the crisis.

In any event, then there was an issue about what to do, and this book of Bagehot has to do with the role of central banks. The way to summarize this literature is real bills versus the quantity theory, and this debate is alive and well today, what causes inflation?

And one view is, you take any money or money-like object, that's the quantity, then you have to restrict it in order to keep prices stable. The intermediation point of view is, no, intermediation is fine, that's what the financial system should be doing. These banks were basically lending backed by the eventual receipt of value from businesses, hence real bills linked to Main Street. And Wallace and Sargent have a paper about this.

The same issue comes up with low- and middle-income countries, although, again, they have not yet made the connection. This has to do with the double use of both electronic accounts and currency, and I actually referenced this in a different way last time.

All of this literature is on the syllabus. So if any of it interests you, that's a candidate, either a two-page review of it to share with your fellow students or adopt it as your bigger research project on this. And here's something about fragmentation is killing DeFi. So again, you could make a connection to the lecture today.

The last thing I want to end with-- I haven't really allowed myself enough time, is to review what we mean by smart contracts on a blockchain. And these slides are from Sam. Thank you. So basically, in about three minutes, I can tell you what the content is.

This is a review slide. A blockchain is a distributed ledger. That was lecture 2. It keeps track of transactions, which are recorded in blocks. And it's a common listing of these transactions. There's a consensus algorithm to come to an agreement about which blocks are valid. And the transactions in Bitcoin involve movement of cryptocurrency from the owner to someone of a different address.

And there is this consensus going on. The blocks are being proposed. There's a validation algorithm. And we've talked about this before, so this is really a review, and we'll come back to it in other lectures featuring Bitcoin as the state machine. And finally, we come to smart contracts in Ethereum.

So the main slide here is we're generalizing the notion of the state of the blockchain. And with Ethereum, we have a complete language for coding which allows arbitrary functionality.

I will say that, according to the literature, almost immediately after the advent of Bitcoin, people were trying to do more with the coin other than just transfer balances. Trying to put some kind of conditionality. Maybe limiting the issuance, or directing it to certain participants, or you can imagine umpteen other possibilities. So Ethereum was born from Bitcoin.

And this is really the main slide here, which is the two types of accounts. There's the Bitcoin type account for balances and transfers and there are contract accounts. And the contract accounts contain the code and the data storage. The contract accounts receive, quote, "messages--" say, transactions, and the contract accounts update the state and send transactions out as part of the outcome of the contract being executed.

So what's the point here? There's a common language. And I mentioned this at the beginning. State in Bitcoin or money means balances and transfers of balances. Here, state is generalized, so not as to preclude balances and transfers, but to accommodate contracts.

And currency, currency, and then you run into these contract accounts, which are nodes, they're valid nodes, but they have different functionality relative to the other agents who are doing transactions.

So as a database, you can talk about the transaction structure, which distinguishes the contracts from the currency. And I'm not able to do justice to all of this. So you can take a look at these slides, and Sam's here also as a go-to for you to absorb more of these details, including Ethereum's virtual machine.

It seemed like today was the day to talk about smart contracts because we talked about how to use them and couple them with the distributed ledger, and on Ethereum, this is what those contracts are building on.