

[SQUEAKING]

[RUSTLING]

[CLICKING]

**ROBERT
TOWNSEND:**

Thank you all for coming. Today is the third lecture, the second after the introduction to the class, focusing on distributed ledgers as a solution to an information problem. The longer title is "Fragmented Markets, Policy Objectives, Regulatory Solutions, and Distributed Ledgers as a Technology Solution." And largely, this lecture involves statics or kind of pseudo dynamics, but not real time.

So here's the outline. Efficiency as the objective of public policy and private sector innovations. I want to dwell on that for a second, because so often you will see central banks coming up with a use case and trying to work it out, spending resources and so on. I have no idea sometimes why they pick that use case. They never give the motivation for it. It's more like something they think they can do with the technology, rather than something that's motivated by the underlying economic need. So I'm going to take the time to go through efficiency the way economists think about it.

Then problems from fragmented markets. So we're going to take the standard setup and chop it up, create some frictions, and then see what happens and potentially how to overcome those frictions. When you have frictions like fragmented markets, it's typical that regulators get involved. Being well intended, they will try to remedy the problem.

And you'll see that in the US, this is the case. But even imposing common prices, we're still going to be left with this information problem and a natural role for distributed ledgers. This, by the way, quite a few years ago was my first entry point into thinking about distributed ledgers, because I made this connection. OK, so that's a little long.

Efficiency is an objective of public policy, private sector innovations as well. So these slides are about Pareto efficient allocations through the planner's problem and competitive markets. Now, you saw some of this material on Friday, for those of you in the recitation. Repetition arguably is a good thing, so I'm not going to skip over it too, too fast.

Let me just say, especially for those of you who weren't in the recitation section, those slides are available, and also [? Ralph ?] and I are available to walk you through. We certainly recognize that the economic background is something that you may need some help with, and we're here to provide that background.

That said, this treatment, Pareto efficient allocations, competitive equilibria, it's all self-contained. It's not like you need another reference to understand the definitions and the concepts. And we will be working with it a lot today. So you'll see the basic definitions and then how we use that in analysis.

So long story short, this is a pure exchange economy with two goods and two households. We're going to have a representation of an efficient allocation that has the property that you can't make one party better off without making the other worse off. And that's going to be this dotted orange line. Note that these indifference curves representing preferences are tangent, a little bit of jargon. The marginal rates of substitution of the households are equated.

And if you pick one good as a numeraire, the prices at which they're trading are equated. So there's a whole bunch of these efficient allocations, and a lot of allocations which are not efficient. When we're off the Pareto front, the Pareto set, we might want to fix something.

Another way to find these efficient allocations is to convert it into a programming problem, which is to maximize weighted sums of the utility of the agents subject to resource constraints. So here the weighted sum is represented by this red tangent line. And given the weights λ of the agents, u^* would be the solution to the programming problem. And u^* would be some particular allocation in this Pareto set.

If you change the weights, you'll rotate along that frontier, and you would be tracing out all of these efficient allocations. So the programming problem makes it easy to find stuff. In fact, we can convert them to linear programs. And finally, a particular economy may be characterized by private ownership where we have endowments. Their starting point, we have prices and a trading line running through the endowment, the slopes being the prices and the outcome, the competitive equilibrium outcome is a particular Pareto optimal allocation.

So hopefully these three pictures are worth thousands of words and a little bit of notation, but I will go through some of the notation in more generality. Commodities could be apples and oranges. General equilibrium theory has a bit of a stigma that somehow it kind fell out of fashion in favor of partial equilibrium and game theory and so on.

I think part of that is just not realizing the applicability. When you start to index goods by time or by realized states of the world, we can deal with intertemporal allocations and uncertainty, and we don't have to change anything at all about the underlying way of proceeding. So it's really quite a robust thing.

So the economy consists of more than two potentially goods. L is the number of commodities. A finite number of households, say more than 2, capital I . We're going to add production and firms J , finite number. And then consumers have sets of feasible allocations, which could generate utility. They have preferences over those bundles. Firms have technologies characterized by production sets. And there is potentially, in addition, some aggregate resources available. So L is the dimension of the commodity space. I is number of households. J is the number of firms.

And then we have this utility possibility set for any economy. So economies are now being summarized by consumption sets and utilities for each household, production sets for each firm, and this aggregate endowment. And we're going to experimentally think about varying the consumption and the production allocation to be feasible for households and consumers and look at the requirement that you can't consume more than what's available. What's available in this case being from the aggregate endowment or from production. And for any feasible allocation, we're going to generate this utility possibility set. So this example was two goods, two households, feasible allocations in the box, generating this pink area of utility possibilities.

The utility possibility frontier is the outer locus of feasible points. Here there's nothing northeast of any of these points. If you had a point on the frontier with something northeast of it, you could make one household better off without hurting the other one, and that would violate the definition of Pareto efficiency. There is nothing northeast if you're on the frontier.

And likewise, the formal definition of Pareto optimality is an allocation for consumers and firms that is feasible, and such that there is no other feasible allocation with primes here that would make no household worse off and one or potential subset of households better off. So Pareto optimality is defined by what it is not. You cannot find such a bundle that Pareto dominates, that would make some better off without hurting the others.

And then as a final step, we talk about competitive markets as a way, hopefully, to implement an efficient allocation. Let there be markets in everything in sight. And then we have to specify ownership. So not just the aggregate endowment, but the vector for household I of its endowments of goods.

Some of the components of that vector could be 0, but some could be positive. And these θ_{IJ} 's are the equity market. So θ_{IJ} is the share that household I has in the profits of firm J . In other words, net profits will be distributed according to those thetas.

So of all raising equilibrium, competitive equilibrium is a specification of prices, a particular price p^* , and these allocations for consumers and firms also with stars, x^* , y^* , such that each firm is maximizing profits, scanning over all feasible bundles in its production set. Each household is maximizing utility subject to two requirements that the bundles be feasible for household I and its consumption set. And in this case, that the value of expenditures on goods not exceed the value of its available resources, which is the valuation of endowment plus the distributed profits from firms. And allocation should be feasible. You can't eat more in the aggregate than what's available.

Now, let me just convert something in words. Excess demand. So excess demand for household I would be the positive difference between its consumption allocation and its endowment. It could have little or no endowment. Who wants to eat more in a competitive equilibrium? So we could relabel this. We could get rid of production and just say the valuation at price p^* of excess demand is less than or equal to 0. Likewise, this resource constraint getting rid of production just says the sum of the excess demands are equal to 0.

I'm mentioning this because shortly when we get to this information problem of decentralized exchange, it's written down in terms of excess demands. But you can go back and forth between the underlying demands, keeping track of the endowment versus just talking about excess demands. Excess demands are negative when you're supplying more than you eat.

And these great theorems without here all the qualifications of all Walrasian equilibrium is Pareto optimal. So it has these efficiency properties. And in fact, any Pareto optimal allocation can be supported as a competitive equilibrium if you're able to redistribute wealth lump sum. So that's a 10 minute summary of these theorems.

And now we'll launch into these two problems that we get from fragmented markets. The first is the intention that everyone should be trading at a common price. And the second has to do with this impossibility theorem of Ostroy-Starr.

So a reminder, again, the Edgeworth box. We start with an endowment. We have a competitive price. Trader one is moving, giving up good one and getting good two. Trader two is moving in the opposite direction. The valuation of the excess demands are equal to 0 for both parties. And in particular, if you thought about the good on the x-axis, so to speak, as a numeraire, they're trading at a common price. The price of good two relative to good one is the same for both.

The US has lots of financial markets. In particular, it has multiple equity markets. Not just the New York Stock Exchange, but also local exchanges in the Midwest and so on for other securities. But it's regulated and there's something called the US national marketing system that seeks to remedy a potential failure that could happen when you have multiple exchanges quoting different prices. Trades could be executed on one exchange at a price different from what would happen on the other one.

And the Security Exchange Commission framework, again, called National Marketing System, NMS, from 2005 is order protection to ensure that retail orders are executed at the exchange with the most favorable pricing. Or another way to put this, an exchange as an institution has to match the best quote from any other exchange.

So exchanges are competing for orders. Broker dealers are placing those orders. There are other things that happen, like giving discounts to brokers who bring in more volume, trying to attract traders, and other obvious things that happen with exchanges. But the outcome is checked against this standard. Chester Spatt at Carnegie Mellon and his colleague have just written a rather lengthy paper called "Regulating Market Microstructure."

Oh, so now that reminds me to say something, which I forgot to say last time. When I scroll through these slides, I am mentioning, briefly, other papers. And those papers are on the syllabus often with stars. So this is like a reminder. Maybe you're interested in getting into the nitty gritty of what the National Marketing System is about, as summarized by Chester and his colleague. Maybe not. There's lots of other options, but I just wanted to highlight these when I go through it.

So they've amended this. The Securities and Exchange Commission changed the National Marketing System regulations. The key idea would be minimum pricing increments, making those increments smaller, dealing with access fees that are charged by the Securities exchanges, and making orders more transparent.

I will spare you the details, although they're written down here. But I wanted to compare this with something else the SEC was doing, which has to do with an antitrust suit against Amazon, which I presume has now been dropped. But when I wrote up these slides, it was still being argued in court.

Amazon was charged with antitrust. They filed a motion to dismiss the charges, arguing that, well, all they're doing is innovations that lead to consumer benefits, and they're pushing their competitors to follow suit. In particular, Amazon matches the discounts provided by competitors. Also, if the items are high priced or overpriced, it will feature deals. And it also features efficient delivery options for Prime subscribers.

Well, at this point, hopefully you're feeling a little bit puzzled, because the spirit of the antitrust suit goes against-- in a commodity, Amazon as a merchandise platform go against the standards of the National Marketing System. So this is not a uniform standard. I don't know whether there's something about financial markets that makes regulation of those markets different or more difficult than regulation of other markets. So something provocative to think about.

The second leg of this. And let me again just say regulation is a tough thing to do. If you had technology available that you could implement, it sometimes mitigates the need for the regulation. So when you're advising central banks in other countries about distributed ledger technologies, this kind of discussion comes up. Better to use the new technologies if possible rather than be stuck with a regulatory problem.

So Ostroy-Starr I wrote this paper, granted, in 1974, which will hopefully through the slides, you'll see, make it even more remarkable, called something about the information problem of decentralized exchange or the heading here is the information problem of fragmented markets. So we're going to again partition agents exogenously into pairs. They can trade with each other at common prices in these pairs. In fact, to make it "easy," quote, we just impose common prices. So the pricing is not the issue, but the fragmentation still creates frictions.

So in this model, first in English, there's a finite number of goods like the apples and oranges. But as I said, they could be assets, even digital assets. The same framework works. They own stuff and they want to get other stuff and they're willing to trade. So they're endowed with goods or securities, depending on the interpretation. They have utility functions over end of period holdings of these assets.

And the goal is to achieve the Walrasian equilibrium allocation. That's the target. It's a good target. It's Pareto efficient to achieve that, even though they're doomed to have to trade with one another in the sequence of markets. What you can't do here is have everyone trade with everybody else initially in a completely centralized market. We're going to respect the fragmentation and nevertheless try to achieve the objective.

So traders are matched pair wise. The pairings are known in advance. Equilibrium prices are used to value these swaps or trades as if they were digital assets. Atomic swaps of digital assets. Each bilateral exchange satisfies a quid pro quo condition that the value of goods received is going to be equal to the value of goods sold, exchanged with your trading partner.

And traders are nothing more here than code. So we're used to thinking they should be maximizing something. Here the objective is given. It's to achieve, if possible, the Walrasian allocation. And the question is, what information would each trader as a node need as an input into the code in order to achieve the objective? It's like algorithmic trading without the max part.

Now, things seem to be pretty special here in the sense they're matched pairwise. Maybe there could be triplets. If they're singletons, that's fine. That's going to be allowed. And the pairings are known in advance. And things are going to get even more complicated if the ordering is random and is maybe even revealed over time. This is quite simple. It's completely forecastable, and yet the problem emerges. So this is the way we want to go. We want to make it as easy as possible and still show that it's potentially hard.

So in notation, there are J traders and a finite number of commodities with this price vector p . And we're going to have a Walrasian environment. So I was anticipating this a minute ago. Each trader has an endowment vector. I apologize. The notation changes across papers. This is the notation that Ostroy and Starr used. Each trader has an endowment vector, and each trader has excess demand at that price p .

So if you start with the endowment B and you get your excess demand, your final consumption or holdings would be B plus C . Remember, things can be negative here. C could be negative excess supply. And as in a competitive equilibrium, the budget constraint is this. It just says the value of excess demand is 0. We've shut down the distinction between endowments and final allocation, taken the difference, and valued it at the price p . And also the sum of the excess demands is 0.

So this is in even more succinct notation the standard definition of a competitive equilibrium starting with the objects, three of them, B , Z , and P . So any BZP satisfying this is a competitive equilibrium. Forgot to say you can't sell more than you own. Z is excess demand. Take the negative sign through both. You'll have excess supply is less than or equal to your endowment.

So the Ostroy-Starr paper refers to this block u as the underlying economy. The underlying objective. You'll see the theorems are going to-- for any blah, blah, blah in x , y , z , and for every u there exists a so-called-- so the u is going to appear in the theorems.

The mechanics of trading. There are T periods where traders are matched pairwise, unless the number of traders is odd, in which case at every date, one trader has to stand out, because there's no pairing left. And this is such that every trader meets every other trader eventually and only once.

A . A are actions or outcomes that happen for trader I , A_{IT} , trader I in period T , gets a positive number if you're acquiring a good, negative if you're giving it up. These are vectors, again, over the underlying commodities. And the restrictions on actions is they should not result in a negative holding of any good. And traders are trading pairwise. So what one guy gets, the other guy gives up for each good.

So the sum of the vectors of actions sums to 0 and trades are fair at those given prices p of the underlying environment. For each trader, they're kind of on their budget line for this pairwise trade. The valuation of the actions is quid pro quo. The value of what you give up, you get equivalent value in return. So the capital A denotes the set of all these actions with these restrictions.

So we could think about what information might be needed starting with something minimal. In fact, they proposed three types of limited information or decentralized information. D_1 , D_2 , D_3 . And then finally, which they will need, the centralizing restriction C . So D for decentralized, C for centralized. D is the minimal specifying that when traders I and J meet at day T , their actions can depend on their initial excess demands of the two on their endowment. And if this is not the first pairing, the sum of the actions of each one up to and through the previous date.

Now, I was thinking about Bitcoin with this. Bitcoin doesn't directly report the balances. You have to figure out the balances from the entire history of trades since the genesis state. And this requirement says, all you need are the balances. It doesn't matter what the individual components were over time, leading up to the current balance.

Now, this always comes as a surprise. We finally name these guys. So we have their identities, I and J . So this actually meant that whenever I and J meet, they put on blinders about who they're meeting. They don't know their names, but they do know these excess demands without knowing the name of the people. So in some sense, we have anonymity here. We drop it. D_3 is the same as D_2 with the names and excess demands, except that now we do spell out the entire history of trades of the two, I and J . And that's going to be crucial.

But it's not enough. The centralizing one is the same as D3, except that when any two get together, the code has available not just the history of trades of each of those two, I and J, but the history of trades of all the traders in the whole system. In other words, even if I and J have never met any of the other people other than I and J, they still know what those other traders have been doing. Or at least if they had that information, then the theorem is they are able to get to the competitive equilibrium in one round of trading.

So here's a formal statement of the theorem, which I called fast, efficient chains. We'll come back to that maybe. There is a trading rule that uses this information in C. It satisfies the restrictions of action set A and it completes trading in a single round for any environment u , for any p, Z, B satisfying u . Not just some of them, but all of them.

But likewise, the twin, there is no trading rule that uses the most information, but still decentralized D3 that satisfies the restrictions in A and completes trading with a single round satisfying any p, Z, B . In other words, starting from D3, you try to implement the solution for any environment. You should be able to find a counterexample. That's what the theorem is saying. You can't do it for all environments u .

But you may need some extra assumptions. And if those assumptions are satisfied in addition, then you may be able to get there. So we want to look at what I'm calling on the slides here institutional workaround. In this case, it's my language, but it's Ostroy-Starr with this paper written in the '70s about to describe to us the current set of contemporary institutional arrangements and highlighting the problems with those arrangements some 50 years after they wrote their article.

STUDENT: Can we talk about whether or not we're supposed to think of these traders as when they're engaging in trades, trying to maximize some objective function? Because you said the goal is to try and achieve some Pareto equilibrium, which kind of implies the existence of some selfish behavior of the agents. But then you also said that it kind of seemed like the objective function of the traders was to just to achieve a Pareto equilibrium.

ROBERT TOWNSEND: In this paper, it is only the latter. So they have a consensus about the social objective, and the issue is whether they can achieve it given this kind of partitioning of agents. So the maximization of individual agents, other than the fact that their preference got used in defining the social objective, is gone. We will next time weaken that and have agents maximizing utility in the way that we're used to thinking about it date by date by date. But here it's like pseudo time, and it's just the design of an algorithm that dictates what they're supposed to do when they meet.

STUDENT: So the utility, or the I guess, objective function of the traders kind of only factors into the definition of what constitutes a Pareto?

ROBERT TOWNSEND: Correct. And still there's a problem. So these workarounds have to do with money, with a big broker dealer, and with some kind of overdraft facility. So we'll do each one at a time.

Money is a medium of exchange. So we pick one security. More on that momentarily. One commodity, if you wish, that has the property that at the specified price, each and every trader can use it in pairwise meetings to buy the requisite target, regardless of the order or timing of the meetings. The intuition being that they have enough of it. So that if they happen to be buying before selling, they can buy out of the value of this security.

I mentioned on the second lecture that monetary theorists think of money as having properties, like a unit of account store of value and medium of exchange. Medium of exchange being it appears frequently in a transaction matrix. And that's what's going on here. It will be necessary, in fact, that they're using this money security a lot. It's a little more subtle than that, but we'll get to that part when we look at the theorem.

Now, there's still kind of a centralization aspect to this. Which good? Which security? If they pick it right and there's enough of it, it's great. But you may need some social consensus. And you could choose wrong. It might work for one possible choice and not for another one.

So here's the definition in theorem-- formal statement in theorem four. There is a trading rule that uses D1. And that's pretty interesting. D1 is the weakest information requirement. Doesn't even use their names. And the spirit of that is I don't care who you are, and I don't care about your history. Just show me the money. The other guys got the money. That's all I need to know.

So this satisfies the restrictions on actions. It completes trade within a single round. But we need this extra condition that there's this commodity, little m for money, such that for every trader i , the valuation at those prices of the environment u_i , the valuation of that good is not less than the demands for goods of all the other commodities. This bracket plussing operator just picks out positive components and sets negative components to 0. So it's just valuing at these prices C for commodity C 's where the excess demand is positive. Remember, Z , excess demand, is a feature of the underlying environment along with P and B . So that's the condition.

So every trader has enough of it to purchase his excess demands. This statement is interesting. Without quid pro quo, which I was arguing was a natural condition, without it, though, it would be easy to find a trading process that completes trade within a single round without these assumptions.

Now, here, please forgive me, but I'm not going to show you more of the proof. And in fact, that would be something very useful to do if any one of you are interested in it. I'm showing you the nuances, but not actually digging out the notation of the proof. So this statement is left hanging here.

However, with the restriction of quid pro quo, when two traders meet, the excess demands do not precisely offset one another. Then they would need to decide what subset of their demands would be met at the meeting if you didn't have the money. But when you have the money, they're trading what they can to match excess demands. And the difference would be filled in by this money object.

So in other words, when you're going through the theorem, it's kind of sounds like they don't do anything but give or get the money. And the theorem is actually proved, allowing some trades to happen without the money, and it's the residual that's getting filled in. I don't know if that's an essential part of the proof, and they could have done it the easy way. But the main thing here is the theorem, ample ex-ante liquidity.

So that brings us to contemporary economies, sufficient liquidity savings mechanisms. So large value payment systems used by central banks to settle the claims of commercial banks play a crucial role. This is delegated to the public sector in almost all countries. It's a natural role for a central bank to operate these settlement systems.

And the value is huge, in the US 100 times GDP. And these systems have been improved over time. It used to be something like end of day settlement, but then people didn't come up with the money. So that was deemed to be very risky. So they moved to real time gross settlement, which sounds like whenever an order is placed by one bank to pay another bank, it's kind of instantaneously settled.

That's not quite right. Why? Because then a participating commercial bank would have to come in at the beginning of the trading day with an enormous amount of liquidity put in escrow in order to honor potential demands on it, to make payments from other banks. And so almost immediately, people were searching for ways to cut back on the requisite amount of liquidity. And that's the terminology, liquidity saving mechanisms systems.

And in particular, you can imagine putting money in escrow and have it available by the algorithm to use, but that it is not used right away. And as an example, this is from the Bank of England. They have one of the best liquidity savings mechanisms you can go. It's in the syllabus if you're of the mind to invest some more. Provides details about how the liquidity savings mechanism works for the Bank of England. And these money market managers of banks follow instructions for sequestering liquidity.

It's kind of a funny argument in the sense that they don't want to mandate too much. They want to leave optionality to the commercial banks. So they basically tell them, oh, don't worry about reporting the fact that you're supposed to make a payment. We don't necessarily process it right away. We'll put it in the queue.

And odds are you might get some liquidity paid to you by other banks in advance. And you'll never have to use this liquidity in escrow. Of course, the opposite could be happening. So commercial banks went along with this, and they're basically using less liquidity overall than they would if they had to sequester the amount for all of the payments.

I find myself wondering what can you mandate? Customers go to their bank say, sellers, and they say, this customer owes me money. My customer to whom I sold goods owes me money. And so they go to the bank of the customer, effectively. And they say, we owe you on behalf of the supplier. We owe the supplier the money coming from the customer.

But banks are strategic. They don't place all the payment orders right away. I showed you that reference in the first slide of the first lecture, this tension between store of value and medium of exchange. So they conserve their liquidity and they play this game. They wait till the end of the day to submit their orders. Of course, if the orders were on a registry to begin with, there would be nothing to conceal and the central bank would just run the algorithm from the get go. But that's not the way it works in practice.

Second institutional workaround. One large trader or broker dealer. So here's another thing that would work. If a condition is satisfied, we designate one of the traders as the intermediary broker dealer. And effectively, all trade can happen through that dealer only, like a giant warehouse facility.

This key broker dealer has to have sufficiently large inventories of each of the securities so as to be able to honor the demands for any security exchange placed upon it by any customer coming to trade, regardless of who comes early, their identity, who is coming late. It's like inventory on demand. And if this condition is not met, then there is no such dealer.

What is the statement? With an embarrassing typo. There is no prices here. This is all about endowments and excess demands. And we're saying, say, this key agent as broker dealer, I , has to have enough in endowments of each commodity C such that the excess demands of all the other traders-- the typo here is that this should be a J and not an I . This is I over here.

We want to sum up over all the other guys on the right hand side. And the sum of their positive excess demands cannot exceed-- in other words, trader I here has the capacity to honor the excess demands for everybody else. So a giant facilitator of trade, you might say.

There is a trading rule. It uses a bit more than D1. It needs names, evidently, and I have no intuition for that. So again, hoping you will. One of you might take a stab at the underlying proofs. So you do have to know who's who to implement the algorithm that could work.

So what's the problem? The problem is market power. One guy who could, in principle, set prices despite the theorem, in contrast to the theorem. He or she would be acting like a monopolist. And should we be worried about market power? Yes. Again, the Kansas City Fed I referenced earlier with respect to Fed Now. And the technological difference here is another report of the Kansas City Fed documenting market structure of core banking services.

And there's something called aggregators. Basically, a retail shop or a department store could take cash. They could take credit cards. They could take debit cards. There's multiple means of payment, and you have these aggregators who basically act as agents on behalf of the stores in order to process the place, to facilitate the flow of funds to settle the trades. The three in the US are Fiserv, Jack Henry, and FIS. And Fiserv has 42% of the market of all clearing of banks in the United States.

In fact, I happened to become aware of Fiserv when I was talking to someone at the Federal Reserve. It was either the Federal Reserve or the Department of the Census, because they had this amazing transactions data proprietary. And there are two others, and it depends. So there's a lot of concentration. This is on the syllabus. The question is huge. Is this something to do with the technology or is it just market power and distortions?

This is for foreign exchange markets. So what's being traded are fiat monies. They're being traded through broker dealers such as Citibank, Deutsche Bank, and UBS. And the largest is Citibank. This table is showing you year by year by year by annual that Citibank had, in this case, was one of the top three dealers by market share, and actually, for almost every year. The others, Deutsche Bank and UBS, have less concentration in some years.

Of course, now this raises the issue. What currencies are these guys trading in? Hint. Not every fiat money in the world is being traded through these broker dealers. The US dollar, the yen, pound, those are still the dominant international currencies. If you're from Brazil and you want to trade with Colombia, a sister country in Latin America, you don't trade bilaterally in those FX markets. You have to get a hold of dollars to buy the other country's exchange.

So then you could raise this question. What if we were to implement the FX markets on a distributed ledger? Could we design it in a way that more currencies would be traded and potentially mitigate, if not eliminate, this concentration? I'm not saying it's all evil. The idea is if you're a broker dealer, you have to stand ready to trade foreign currencies.

So you have an acquisition cost. They have to be on your balance sheet and a fairly high level. So how do you finance the acquisition of currencies? Plus you're going to hold them, their exchange rate fluctuations. So there's a risk component. So the acquisition costs and the risks already mean you have to price to cover those costs. Those are real costs.

But likewise, if you could envision a blockchain chain where many more currencies could be traded, you don't necessarily have to have the security in a buffer stock if you can contract, as in a dynamic ledger, to sell things you don't currently own, but you will acquire from somebody else. And so I guess two lectures from now, we'll talk about tokenization and dynamic ledgers and atomic settlements. So we're kind of anticipating that here.

And this is icing on the cake, in the sense of finding other things that are going on in the foreign exchange markets that are really interesting and beg for an explanation. Namely, even though there are, say, a non-trivial number of broker dealers, you can categorize them as being on the periphery or in the core. And it's the few core dealers who are absorbing the risk of exchange rates. And otherwise the periphery dealers, when they get a hold of a currency, they tend to want to get rid of it right away, referring to it as a hot potato.

So why is it designed like that? That to me is a good challenging economic question. But this is the way it is. And then you would think that if it's down to a few core dealers, that the exchange rates would be determined in part by the underlying conditions of those dealers in terms of their capital ratios, for example. And people who have looked at the data are finding that that is true. I'm fascinated by this. I love economics. But anyway, if you want to explore further, this would be an interesting thing.

Now, a third workaround is overdraft credit facility. Ostroy-Starr wanted to get everything done in one round, but you could imagine two more rounds, an initial round in which traders are given overdraft facilities, credit effectively, and they can use those credits when they go shopping. And the requirement is that they got a payback, essentially, the amount of the overdraft at the end, one way or another, in equivalent value. So then it's like the money that we had before, but the money from before was one of the goods or one of the securities. And here it's another object. It's just a credit.

Well, then it's easy to have enough of it, because you can create it in arbitrary amounts, as long as these guys pay back. A hint of things to come. And I was thinking about this this morning reviewing the slide. This also means the amount of credit would be agent specific, because the value of the excess demands of each trader are going to be different. Some may be not trading much at all. Others need a lot of credit. So if it's repaid at the end, that's the second round, and you have three in total. So it works, but as we're already anticipating, there might be the possibility of default.

So let me make this a little more tangible and go to medieval trade fairs. So this is a credit system. In principle, bankers were sitting at a table and they had their ledgers. And in principle, you could, say, deposit your gold or silver coin for safekeeping, like a deposit. And you would be credited on the books of the bank. The bank owes you money. It's your money. You get an IOU, and you could use that for trading.

These giant trade fairs in France, they tried to minimize the amount of coin that traders would bring to the fair, because it was risky to be hauling around all that gold and silver. Not to mention the goods. So the fair would designate a sequence of trades, first in cloth, then in leather, then drugs and spices, and finally at the end for settlement. Those were the rules.

So if you basically were bringing in spices, then you don't get value from stuff sold. You got to buy first. These are big traders. So they're trying to bring goods that they have access to and get other goods that their customers would like, and spices being a great example.

So then comes the credits on the book of the banker. Instead of bringing a lot of coins, say, they brought no coin at all. Every person would be credited with a certain overdraft facility. You should still think about it as an IOU owed back to the banker, although in other ways, you can think of it as a loan that has to be repaid. It's overdraft. You're not expected to carry that overdraft balance indefinitely.

And then they come together in the settlement period, which I gather was also a couple of weeks. And I think there is some interesting stories there of exactly how that worked. It could work perfectly if the prices were all posted in advance and cleared the markets, as in a Walrasian setup. But obviously, they weren't doing that.

Just because I can't help but mention it, what was the unit of account being used on the books of the banker? There were different coins. Gold coins, silver coins. In fact, the king would occasionally call in all the pre-existing coins and melt them down and give out new coins that had less metal. So that's a classic debasement. Central banks have other ways of doing that now.

So they would continue to keep the unit of account they had before, even though it no longer corresponded to a circulating coin. In fact, that was called ghost currency. I love it. Making the point that the unit of account is one function, and it's not inextricably wedded to store value and medium of exchange.

Anyway, these guys would meet at the end of the fair, settle the accounts. But some people were owing back and they couldn't pay back the overdraft. So then they would issue an IOU, which was the order of I'll pay you later at a specified date and location when we meet up again. In fact, there was this huge international affair at Lyon where eventually it was all about clearing during these promissory IOUs rather than trading goods.

So I've mentioned in lecture one, and we'll come back to it, these IOUs were letters of the fair/*lettres de foire*. And they are *basically*-- well anyways, they will arise as circulating notes in England and other places and led to financial crises. But that's the lecture for next time, not for today.

So even in medieval times, they're working really hard to get an operational system with goods and finance to go against some of these overdraft problems. In practice, there's default. So I think arguably the largest market in the world, the US repo market, which is used by the Federal Reserve for monetary policy. They discovered, to their dismay, in around 2008 or '09 that these broker dealers were basically assuming the obligations of some of the people they were paired with, with intraday credit was being provided by these clearing banks during repo transactions greater than the US monetary stock.

So the reforms, among other things, I'll just focus on the first one, was to alter the timing in such a way that the broker dealers had more balance on their balance sheet. But these things come up. They came up the other day when we were talking about emerging markets. So this is still a potential problem. Talked about potential remedies, but institutions that would require it sometimes suffer from their own big problems. Is there a solution with the new technologies? Which brings us back to distributed ledgers on the blockchain.

The impossibility theorem is that you can't do it if you're decentralized in the information. D3 won't do it for you. You need C. So the idea is kind of intuitive that even if traders are meeting pairwise, those trades are reported to the ledger. In fact, you could have all their value sequestered in escrow to begin with to be used for the trading. With that history, it will be possible for people to coordinate to achieve this socially desirable outcome.

A little bit of a story here before I get into a limited amount of detail. Ostroy-Starr showed an example-- again, the idea is proof by counterexample. Try D3. Claim for any environment. You can succeed to get to the Walrasian outcome. All you need is an example where it isn't true, that a price, excess demand, and endowments corresponding to the underlying conditions would not be achievable. And I'm about to show you the example that proves the theorem as a counterexample to the theorem that isn't true.

However, Kyungmin was a graduate student here, was my student. And as I do sometimes, and Sam can testify to this, I ask my students to help prepare lecture notes and material. And you've already seen some of Sam's slides in earlier lectures. So I asked for Kyungmin to do this for the Ostroy-Starr paper, and he went away for a while, came back, and he said, Professor Townsend, there seems to be a problem. So their example wasn't working.

And I know Ostroy and Starr, so I wrote them an email. And they're like, oh, that was so many years ago. Why don't you get some graduate student to work it out? Which is, of course-- so we don't know for sure if this example is validated by-- it is not validated by Ostroy and Starr. They haven't acknowledged.

But anyway, how does it go? So their big picture is the most important thing. We want to find two economies such that in a certain trade, only takes one, two traders cannot decide what economy they are in. When they look at their history and their own endowments and so on, they get some information, but it's not enough to know the excess demands of everyone. Hence, they have no guidance about what to trade to whom when they're meeting. So that's the big picture.

The economy is kind of special, but that's good, because all you need is a counterexample. Again, it helps to understand too the aspects of Ostroy-Starr. Traders are meeting pairwise. Two, five meet pairwise. Three, four meet pairwise, and one's standing out alone. If you look at this, two meets five, two meets three, two meets one, two meets four. Every trader is matched with every other trader once, not just for trader two, but every other one. And of course, because there are five people at each period, one guy is kind of standing out all alone and is not trading.

Then there is the price vector, which in principle, you would think there would be different prices for different goods. But all they do is say if this good costs \$0.50, we're going to relabel the size of it to be two units. So then it costs \$1. So we just normalize the units in which everything is measured in such a way that the prices are all equal to 1. That kind of helps keep track of values. And this is the matrix of excess demands.

For trader one, two, three, four, five as the rows, and the columns represent the goods. So trader one has an excess demand of four for good one and excess supply of good two in the amount of two, et cetera. So you have these matrices with rows and columns for traders and goods. And you have their endowment vectors.

Curious thing about this example is trader five is worthless. Trader five has 0 of everything. So whenever you're matched with five and we have this quid pro quo requirement, no trade happens. So we can just focus on the one remaining pair where something active might be going on. I'm just sketching how the counterexample goes, because you'll get bogged down. I always do when I start going through the details.

Trader two meets five. Trader three meets four. At date one, five is worthless. Trader three and four have one good four. It's the same good. They don't trade. At day two, trader one meets five. Two meets three. Only two and three can be engaged in a non-trivial trade. Now, something more interesting. After trading at two, date two, trader three does not trade with anyone until he trades with trader one at date five.

So there's a backwards and forwards looking aspect to this. Whenever you trade, you naturally have to think of the social objective function. So you think about the current state where you are in endowments, but endowments get transformed. This is classic computer science coding. Let B be the endowment. Then you relabel the endowment to be what it is now endogenously after two periods of trade, and similarly for the excess demands, which are moving around if they're actively engaged in trade.

So I'm going to skip. The idea is to go through and determine at day three what traders one and three have to give to each other that you could work out. Then you change the underlying economy by switching these two rows around. It's another valid economy. And it turns out the sequence of trades would be the same up to a point. And you realize that they're going to have to do something different than what they did in the first example. But looking at what they know, there's no way that they would know at that point, given the trades that they were engaged in or that they saw, that they're in a different economy.

So that's kind of the way-- it's quite tedious. I've double checked it. Every time I read it, I get involved with going back and verifying all the claims and so on. But I think it's better just to take away the spirit of how the proof is working. And you can do this at your leisure.

So we talked about the distributed ledger and how the trades are recorded there, and everyone has access to them, that it solves this information problem that arises from having partitioned trade. And do we see something like this? Yeah, a big aspect of what we see on the blockchain, not the only thing, is tracking. So Walmart Canada is involved in shipping goods. And there are lots of problems plaguing the transportation industry. There's discrepancies in the invoices and the payments and fights about whether or not you were paid or that the goods were not there, because they could have been stolen. And so they have this tracker running on the distributed ledger to mitigate those problems.

And another example are diamonds, although I don't know, you may not like diamonds coming from the Congo. But De Beers, the big diamond company, has a tracker on a distributed ledger that allows you, in principle, to know the origin of the diamond. So I could not find on their website exactly how they do that, because this has to do with the origin of the goods, not money that's changing hands along the way.

But likewise with trade finance. Ghana has this prototype. And as Henry can tell you, Hong Kong has similar designs. Ghana is collaborating with Singapore, and they've got proof of concept for trade finance. So these very small merchants in Ghana who are exporting goods want to know that they're going to get paid in an Asian currency and that the goods are there.

There's a tracking for the shipping of the goods and so on. And the goods, the money is held in escrow and it's not released. That last part has to do with the smart contract, which we didn't talk about today, but we will talk about it next time.