[SQUEAKING]

[RUSTLING]

[CLICKING]

ROBERT M. TOWNSEND: So let me just say from the beginning, there are some connections between the lecture today and the one last time, and I will try to point that out. And then there's a missing third lecture. If time permits, I will kind of outline that at the end of today, but let's see how it goes with these slides first.

So this is stochastic financial networks, liquidity and the value of key players versus contagion dynamics. Or to put it succinctly, do we enhance or limit markets? So here's an outline. Introduction to stochastic financial networks, showing you how markets vary over time and in principle with other shocks. I'll lay out the economic environment. We'll define a stochastic financial network with examples. Anyone who participates at all is in a centralized market and otherwise isolated, and in the second, the markets fragment. They're partitioned.

And then we'll think about ex-ante injections of liquidities as buffers against shocks, and in particular try to pinpoint who would be the most valued person to receive the injection of liquidity to carry it into subsequent markets in which that agent participates. And so we'll characterize the most valued player in that sense, in the baseline environment, and also for a more general class of environments. And then we'll do some positive economics. What does this value correspond with in financial markets, go to Thai villages you've seen before and do some empirical work to see how well the theory is holding up, at least at an initial level, and end with financial centrality and contagion, disease, systemic risk. And the point is to compare and contrast the market making aspect of judiciously chosen liquidity injections with what is the current policy framework, which is to limit the interactions across players due to this concern about financial contagion.

All right. So I should say from the outset that when these slides were written and the draft of the paper, we were referring to our measure of the value of liquidity for key players as a measure of financial centrality. That has caused confusion because people associate financial centrality and network financial centrality with the contagion point of view. Our measure is not equivalent with other measures of financial centrality, and I'll show you that. We are going to change the wording to be liquidity value of a player, rather than the financial centrality of a player. But I can't change the slides that quickly, so we're kind of stuck with it for today.

The main thing here is to talk about disruption to markets. Disruptions take the form of shocks that limit market participation. There is a literature, at least three literatures, one a very famous paper by Darrell Duffie and coauthors having to do with over-the-counter markets, which has a search friction aspect to it and has to do with the broker-dealer markups and the volume of trade that can be supported.

And you may or not be familiar with a class of monetary models in which traders meet at random. There's a supplier of good and a potential buyer of a good, and then the issue is what to carry around with you to facilitate trade. Kyotki and Wright and others are born into that way of thinking, and it's quite influential. The point here is, again, we have this kind of random matching. And I should say, we've been studying partitioned setups almost from the get-go in this class. Granted, deterministic pairings, but not everybody was paired with it, matched with everybody else all the time.

And finally, we have this explicit random market participation, where there's a kind of borrowing and lending market, but some of the participants leave early, others arrive late, and this was used to model the need for the Federal Reserve to inject liquidity into the system. It's based on much earlier work of Milton Freeman and the need for the Federal Reserve to manage liquidity, which is by the logo on the door. As you go into the New York Fed, it has this quote. It's not about full employment. It's not about price stability. The Fed was set up to manage these liquidity shortages.

Financial centrality or liquidity value is the marginal social value of giving a little bit more purchasing power or goods to an agent, conditioned on that agent being able to trade with other people. So you don't carry the liquidity into autarky. It's of no use for it. You can think of it as like a financial instrument that you can trade.

It's a social point of view, and I've emphasized this before when we talk about money, what is the purpose of money. Here, quote, "The liquidity is to enhance the social value." So you give resources to an agent, which not only increases her consumption, but the consumption of other agents that trade with the recipient. And we'll formalize this.

It's also possible that markets not only suffer from these exogenous shocks that determine who's in the market, but that, in addition to that, people can decide endogenously whether to go to the market or be attentive to get online and so on. And it could, in principle, go either way. But intuitively, by subsidizing an agent who carries liquidity into the market, that may make the market more attractive to other agents who otherwise would bear a participation cost. So you can get these externality-like aspects going on, but we will formally model what we mean by community value.

So here's a actually pretty well-known picture of interbank market, federal funds markets back in 2006. So these are banks who are deficient or have excess reserves, and they're borrowing and lending with each other. But here, the point is breaking a typical day down into half-an-hour intervals, although not everything is shown here. And you can see, at the beginning and end of day, we have these classic kind of network pictures of who was trading with whom, and the market's thin. On the other hand, midday and so on, it's a much denser graph.

Now, this shows variation over time, so it doesn't really show you the shocks. But you could well imagine that if you picked a particular interview-- interval of time and looked at different days, you would see the extent of this kind of varying over time, which could be due to endogenous participation as well as potentially to shocks. A side note that I hope we can come back toward the end of class, people don't show pictures of the federal funds market anymore because it's essentially collapsed. There's so much liquidity in the system that the banks don't really borrow and lend with each other, and instead there is an active repo market where you have money market mutual funds with excess liquidity, in other words, not commercial banks, a different type of player, effectively lending to hedge funds and pension funds over a very brief amount of time. And "repo" refers to the collateral, which is backing the loans. So that's the principal. The repo market has become the principal venue for Federal Reserve monetary policy, not tracking the Fed funds rate anymore.

STUDENT: What are the colors in the-- [INAUDIBLE] different colors?

ROBERT M. TOWNSEND: Yeah, I don't know. Good question. It makes it look nice. It must have a meaning. I'll try to go back and find it. Oh, and this may remind you a bit of the network pictures that Tomaž was showing last time, but they're just nodes, and the edges here refer to trades.

So what's the underlying environment? It's a risk sharing environment. And I did show you a couple of slides on the third lecture about what is a reason for intervening or not, and how well do we do in accommodating risk. And I referred in one slide to Thai villages and another slide to Indian villages. And we'll do more of that later in 193. But here, it's all self-contained and laid out.

So there's a finite number of agents, capital I. The number of agents is little n. They are risk-averse agents with a concave utility function, strictly. Agents maximize expected utility by choices of various things, and they have random incomes, so of the n guys, this vector 1 through n of realized incomes drawn from some distribution. So that's one of the fundamental shocks in the system, exogenous random endowments. An example would be these constant absolute risk-averse utility functions, which are basically exponential, coupled with parameterized risk distributions, so that this vector of incomes has a mean mu vector in [? RN, ?] and this variance-covariance matrix, which is an n by n matrix that captures the correlation of one person's income with any other person's income.

So let's look at the other shock, and we're going to treat it like an income shock in the sense that it's an exogenous random shock. We want to enumerate all possible shocks to the economy. This xi, this guy here, is a vector of them over the n guys, and they take on binary values. They're either 0 or 1. 0 obviously means not in the market. 1 means in the market.

When they're not in the market, they're doomed to just eat their income. But when they are in the market, they may be able to share risk, the income risk, subject to the resource constraint that agents who are in the market cannot eat more, in the aggregate, than the amount of income that's available in the aggregate. So those xi's are just zeros or ones, so we're adding up over everyone. But the guys who aren't in the market are getting a 0 xi, so they're really not in that constraint.

So what's the community objective here is to maximize expected utility to maximize lambda weighted expected utilities, where the expectations over the underlying income shocks as well as over the market participation shocks. And the choice variables are consumption, and again, if you're alone, you're doomed to eat your income. But otherwise there is this resource constraint repeated out here that applies for all possible configurations. No matter what the realized income and market participation shocks are, you cannot violate the resource constraint. So there's a whole series of these guys depending on the xi vector and the y vector.

Yeah, so the idea of choosing consumption is to pool risk, right? So ex-ante, for example, and we'll look at it. It's a very special case, but it's intuitive. Everybody's all alike. They have concave utility. Incomes could be high or low. They're drawn at random. Ex-ante, you don't know if you're going to be on the high side or the low side, so you enter into a mutual agreement that if one of the realizations is one person is high and the other is low, then the high person is going to give up purchasing power to the low guy, and vice versa. If they're, in this case, both high or both low, there's not much you can do, and that's kind of the aggregate state.

So let's start talking about the network aspect of this, refer to a centralized market you're in or not, and I'll come back to these bilateral connections when we get to the relevant slide. So let's look at an example. We have a network given, and these participation shocks are modeled as a process. So it's as if one node in the network, one trader, is randomly chosen to be the host with some probability, maybe just 1 over n, that everyone is equally chosen, or it could be pretty special, and then that host is sending out messages to other traders. And if they get the message, let's just say for simplicity, they join with the host and have a party. They're in a market.

You do see this process in actual markets. You'll see, especially now with the internet, you have someone with something to sell, and they'll send out invitations. It's good for 10 minutes, and whoever shows up can potentially bid for the asset. It's like a mini auction and it's very focused in time. And obviously, not everyone participates. So you start to get into these shocks.

What is it? It could just be latency. We haven't even said what the frequency is here. I mean, these could be markets that are occurring every 10 seconds, but with latency in computers and so on, there's a "technological reason," quote unquote, that not everybody receives the invitation, or at least they don't get it on a timely basis.

STUDENT: If you receive an invitation and someone else also receives an invitation, can those two trade with each other?

ROBERT M. TOWNSEND: Yeah.

STUDENT: Or do they all just trade with the one-- they don't just trade with the one trader, then [INAUDIBLE].

ROBERT M. TOWNSEND: Yeah, the slide I skipped, you'll see it in the graph. They all trade with each other. Exactly how they do that, we'll come back to it. So with what probability do you get an invitation? Well, something to do, say, for example, with the distance of the two traders, in the sense of whether or not they're adjacent in the network map. Remember, we're taking as a primitive here the network.

So let me show you the picture. Here's the network with the nodes and the edges. In this case, almost everyone's connected to everyone else. There's a few isolated things going on here. And then we see this network with the host, so that guy is chosen at random. And then he or she is sending out invitations, and the adjacent nodes are more likely, with probability q, to get the invite. If you're two nodes away, it's q squared, so it's diminishing. And this picture acts as if, after traversing three nodes, the probabilities are 0, but you can imagine that's just an approximation. Occasionally, you might get a q4, but not too many of them, and so on. So this is a market formation process emanating from some underlying fundamental network map.

The obvious thing that comes to mind here is distance. And again, at this level of abstraction, it could be-- it could be geographic, in terms of who's adjacent to you or not. More generally, there could be some underlying characteristics that we're not modeling, some kind of affinity or similarity across agents that make some people close to you and easier to communicate with and other people further away.

STUDENT: Could do equivalently to think about these distances being defined in terms of the probabilities of trading? Or is it important that the structure of the network implies these probabilities? It's not [INAUDIBLE]

ROBERT M. TOWNSEND: The former is the main-- no, it doesn't matter. This is an example to show how you could have a network and that that will generate the stochastic market participation. What we need for the theory is only that market participation is random, and I'll show you another way of doing that. Oh, yeah. This is what's left after they're connected.

I deleted some slides here. You can imagine the seed is this guy, and then you'll get a different picture of who's connected to whom. So this is the market chosen stochastically through that underlying process. Determined, I should have said.

OK. More generally, you could think, not that everyone is either in or out of a market, but that they're getting divided into fragmented or getting divided into segments. So in this case, we have capital I here. It's the set of all agents. This power notation is the set of all subsets that you could form with n agents. And so this is a cluster, one subset, say, proper subset, of the set of possibilities of all of them. And their intersection is null, so they're partitioning themselves into these clusters. And this is a typo. So if you take the population in-- you're somewhere, maybe isolated, or paired or matched with other people, you'll get just a set of agents back at the end.

So we're randomly partitioning the set of all traders here. So you could imagine we have a probability, to answer Laurel's question, we have a probability distribution over these clusters. Mu would be the probability of a particular cluster drawn from some known distribution.

STUDENT: So what is k? If you just take the previous picture, you have either a group was chosen and the rest was out, and here we're picking a partition that contains possibly more than two?

ROBERT M. TOWNSEND: Yes.

STUDENT: And then that means that--

ROBERT M. TOWNSEND: Where did it go?

STUDENT: What's the difference between the groups?

ROBERT M. TOWNSEND: Here's a network, and then we have a random variable that chooses who is active in the network. And in this case, if you're concentrating, this guy, because this guy's not active, this guy's going to be completely isolated. These two are connected.

STUDENT: Oh, those are the partitions.

ROBERT M. TOWNSEND: Yeah. So these will end up being-- and likewise you could-- oh, and so now I'll answer your question about the dotted lines here. So if they're sharing risk and they're in a group, in this case, the group of three, we will determine the extent to which they're getting or giving consumption relative to their endowments. And you could well imagine that effectively they're all connected, all three with each other. This dotted line is a way of saying, don't burden me with the math of figuring out how much this guy should give to this guy. But likewise, you could imagine that they are perfectly able to achieve the optimal risk sharing just by going through the central node.

Whereas last time, what Tomaž was emphasizing in the pictures was you're only connected bilaterally with the people who owe you money or to whom you owe money, although there was a multilateral aspect to it in the sense of forming those chains and those cycles here. We just kind cut to the chase and allow full interconnectedness among any subgroup. And it's really without loss of generality. So I skipped in order to answer [? Raffa's ?] question.

So what are we doing? We're, again, maximizing utility subject to resource constraints. This is the resource constraint for a particular cluster, [? ML ?] in the previous slide. And yeah, somehow the notation is getting a little sparse here. So the relevant state is going to be income within this partition, although the larger state is the determination of these clusters, randomly chosen. So a stochastic financial network is a network formed by a collection of complete isolated subgraphs.

So you can have the graph, which is given, characterized by the nodes and the edges. And we have links, in this case, not directed, undirected links. i and j are linked or not. There's an edge connecting them or not. And if there is, we say that the Gij is equal to 1, and otherwise it's 0.

And then I went through these examples already, although I'm not sure-- this was one randomly chosen partition and this was a second randomly chosen partition. I think I cut myself off and didn't show this comparison slide, just to illustrate how very different the thickness of the market can be and who's participating and participating in which cluster. The international markets these days are becoming much more fragmented than they used to be, partly as a consequence of trade embargoes and financial sequestering of value, and so on. Just a comment.

So financial centrality, let's think about the value of liquidity where we increase income by agent i by epsilon, where epsilon is a small number, just an infinitesimal variation around their income endowment. And we do it ex-ante before the market shocks are determined and before incomes are determined. So it's an ex-ante commitment to this named trader without knowing what exact situation she's going to be in.

So a liquidity injection, in essence, is marginally increasing income. You could write it out a little bit more formally. Income is the sum of endowments plus assets. There's some asset which is very liquid. Its price is set to 1, so it converts into consumption goods on a 1 to 1 basis. And you have that asset set and the value of it as your starting income if you're in the market, But it's a liquid asset. So if you're in autarky, you don't have anyone to trade with. And this is equivalent with agent i having this initial asset position and then getting epsilon more of it.

So you could think about it as money, except we're not really closing the model, "money," quote, from the Federal Reserve being arguably, if it's the US, a very liquid asset. When I say, not closing the model, if we gave out too much money, we'll start to influence the price level. But this experiment is starting with an infinitesimally small amount that would have virtually zero impact on the price level, with the goal of figuring out who are the valued players.

The Fed does not do this explicitly. They intervene in the repo market by, say, lending liquidity, or they have market facilities that allow banks and so on to borrow a window from the Fed. But this is like saying, no, suppose a monetary authority were proactive and could identify these key players who bridge across other players. And the frequency of the shocks is so high that you can't wait for Silicon Valley Bank to do something bad. You've got to decide in advance, and so you could very well interpret this as the value of monetary injections and who they should go to. But it is on the assumption that the traders have entered into a risk sharing agreement with each other.

So the definition of financial centrality, or our value of liquidity, is how much the community objective function is enhanced, in this case, if it's an injection to trader i, how the community value function is enhanced as we change the value of epsilon around 0. We're trying to find who is central, so we want to run this calculation over all the agents and rank them.

And again, what was that Vi of Epsilon? V of lambda was the value function given fixed weights lambda over the community of people participating in the market. Vi, i for injecting something to agent i, which will alter this resource constraint by epsilon. And we're then taking the derivative with respect to that injection.

STUDENT: So here, the consumption you choose can depend on y and-- what's this letter called xi, psi?

ROBERT M. TOWNSEND: Yeah, I'm saying xi, but I'm probably getting it wrong.

STUDENT: But it could also depend on-- or is that implied by xi what the realized partition is? Because I mean, if I'm--

ROBERT M. TOWNSEND: It's implied by xi, yes. Yeah. So let's do a policy experiment where we imagine we have more than just epsilon close to 0 amount of liquidity. Suppose it's a finite amount A that's available for injection. So in principle, we could be injecting to multiple, but not necessarily all agents. The sum of the injections cannot exceed the amount that's available. So the larger objective would be to choose the injections over potentially all the people in such a way to maximize the social value function subject to this limit on availability.

But if there are finite number of players, then we go back to our other infinitesimal measure and rank order all the players. So you can imagine that there is diminishing value, but maybe not diminishing so quickly. So if the total available is not very big, you might still want to inject it all to just one guy. If the amount available A is slightly bigger than that, then you might want to go to a second person, and so on. So that's all this said. There exist these critical values, A bar, such that if the amount available is, say, less than that, we're still going to want to give the money to just one person.

So what's the value of giving the liquidity? Well, the risk sharing effect is what we're emphasizing that propagates through the risk sharing network. There's a participation effect, which I don't have time to talk much about today. It's like a subsidy for participation in the market. And there's an income distribution effect that the injection can start changing, basically, their endowments in a model that's more articulate about how the endowments are determined. These are interesting formulas. But then again, we don't do much more with it, so I'm just going to today continue to focus on these exogenous shocks, market shocks.

So if we think about the Lagrangian, as always, the repeat of the objective function is the lambda weighted expected utilities with expectation over the states, including market participation and income enhanced by the resource constraints, where qs is the Lagrange multiplier, or the shadow price, of the particular constraint s that the sum of incomes cannot be less than the sum of consumptions. So qs is the Lagrange multiplier, and not too surprisingly, it as a shadow price reflects how much the objective function could be increased if you were to slacken or weaken this constraint by a tiny amount. So it's like a price of resources in that state.

In fact, when you're thinking about what prices are in a decentralized competitive equilibrium, they're essentially marginal utility prices. You have to pick a numeraire, and then it's really only relative prices. But a decentralized competitive equilibrium is using marginal utility pricing. So there's a proposition. What is the liquidity value for agent i, this financial centrality for agent i? It's just the expected product of i's participation shock with those shadow prices.

If agent i is not in the market, then xi i, under a given realization of shocks, then xi i is 0. So this goes away. And if xi i is 1, agent i is in the market, and what is i's contribution? The epsilon is small, so it's just the value of liquidity in that market, which is reflected in the shadow price, the qs price. So that is the measure of liquidity value of player i, and now we can go through some examples. How does that value vary as a function of the underlying environment?

So one example is everyone has an equal weight. This weight could be 1 over n so that it sums to 1, but that's entirely arbitrary because we can rescale the objective function. So they're all going to get equal weight in this example. They all have a common utility function, no i, and their incomes are drawn from a distribution, which is independent across all the agents. And it has a mean mu, which is common, and a variance sigma squared. So if we were filling out that variance-covariance matrix, there would be sigma squares on the diagonal and 0 everywhere else. And these participation shocks are orthogonal to the income shocks.

So in a given market size state xi, which is summing up over all of them, but many could be 0. So we get the total number who are participating is just the sum of those across the objects of the vector. And what is the solution to the risk sharing problem? Essentially, you give everyone the average of the incomes of the people participating in the market.

The market's random here, so just kind of a word of caution, but whatever x is take the-- oops. I should have had a xi in here. So we just add up over the incomes of those people in the market and divide by the number of people in the market, and that's going to be the average income. And that's going to be the target for these people. Their consumption is going to be the average. So it's like the two agent example, high-low, low-high, and so on, except now the number of agents in this mutual fund is random. But whatever it turns out to be, with identical preferences in IID shocks and so on, they just share all the risk because they're all risk averse to the same degree.

And so the marginal utility, the shadow price of that resource constraint is just the marginal utility of the common utility functions evaluated at that average income, because they're all alike and they're all getting the same thing. We can do a little more, generally, especially with approximations. This we had before. Financial centrality is the expected product of participation with the shadow price, and that's going to be equal to mu, the mean of the income distribution plus the marginal utility at that mean plus the expected value of this guy here. So it's the product of xi i with this.

And this gamma has to do with the third derivative of utility. It's called prudence, and the idea is the more prudent you are, the more you will want to save when you're facing a risky situation in the future. Of course, not all utility functions display prudence, but some of the familiar ones do, like the power functions, C to the omega or something. So we have prudence times this. So the higher is the variance, then the bigger is the ex-ante expected utility of the liquidity injection, because there's more risk, more diversity across realizations of income shocks, so there's more gain to pooling. On the other hand, it diminishes with the number of agents in the market. Well, that's kind of like a diminishing returns. I mean, you're getting closer and closer to the average. Having one more person doesn't really help much.

If we had the segmented market, rather than you're in or you're out, we get a familiar almost equivalent expression, except that now we're not dividing by the number of agents in the centralized market. We're just dividing by the number of agents in the cluster in which you're randomly located. And again, you're taking expectations over all these guys, over all these clusters, because it's entirely an ex-ante measure.

If you think about these xi's, xi i, it's binary. It's either a 0 or 1. So you could imagine writing two branches here, the probability that it's 1 plus the probability that it's 0. But if it's 0, this whole thing goes away, so you're left with the branch where xi i is 1. And so then we're able to pull that out here conditioned on xi i being 1. That measure of liquidity value is just this probability. I mean, the intuition is pretty brutal, so why am I giving liquidity to a guy that doesn't show up very much in the market? It's kind of worthless. So you're weighting positively in the liquidity formula someone who is around a lot in the market.

STUDENT: So this is an ex-ante measure, I understand. But everything we did, also consumption and the injection, all depends on realization-- can depend on realization. So it's all exposed. So then why do we care about an exact measure? If we can make our injection depend on the realized clusters, and we can make the consumption that they have depend on the realized clusters, and [? y's. ?]

ROBERT M. TOWNSEND: Yeah I mean, that is what we're doing mechanically. We're taking realizations of income, realizations of shocks, and then allocating across the agents. The ex-ante part starts to enter in when you say, well, look, if we waited to trade, and I'm high and you're low, I'm not giving up stuff. So it has to be mutually beneficial, and they're in a smart contract that commits them to share the risk.

And then this is ex-ante for sure. It's a probability object. Yeah, so we're starting with your point but then integrating back before anything happens. But it is crucial here that the injection has to-- the site of the injection has to be determined ex-ante also before anything happens, like a buffer stock.

STUDENT: Oh, but I thought the injection was allowed to depend on the realization of the fragmentations.

ROBERT M. TOWNSEND: No. Yeah. It's all about hedging the income risk and hedging the market participation risk. And this guy, the second term, looks very similar to what we had before. It's just that you're conditioning on this branch.

STUDENT: So you're not injecting liquidity to the guy with the highest shadow price, but you're injecting liquidity with the guy with the highest expected shadow price or [? something. ?] I thought it was a shadow price or something that determined the [INAUDIBLE].

ROBERT M. TOWNSEND: It's the expected thing.

STUDENT: Expected, [INAUDIBLE].

ROBERT M. TOWNSEND: Here it is. Yeah, so this is the shadow price conditioned on a given draw of the income shocks and the market participation shocks, and you're taking expectations over that.

All right. So another example. We're back to the host. Suppose you pick agent i as the host with probability 1 over n, and then the markets are formed, as I said, by this invitation. In fact, a version of that is that agent can connect only to adjacent nodes, so that's like his or her neighborhood. A neighborhood of i is just a set of adjacent nodes in the network, and di is the number of people in that to which agent i is directly connected.

But there's a bit more to the formula than that, even though it's additive. This other component has to do with the variance, the prudence. And a bit counterintuitively, perhaps it's divided by the extent of the neighborhoods of the other traders j to whom agent i would be connected as the host.

So let me show you a picture two different networks. Here's agent i immediately connected to 1, 2, 3, 4 guys. Here's an agent-- it's kind of bad notation. It should be i prime, some other configuration, and he's connected directly to 4. But now these guys, like 5, is connected to three people, not to two people. So in every instance, the number of connections of the periphery guys is larger on the right than it is on the left.

You might have thought that this guy would be connected because he's directly and indirectly connected to more people, and this is the disease contagion, systemic risk point of view, that this guy is dangerous because he's quite connected to almost everybody, at least indirectly. But our measure of liquidity is about market making. So why is this guy more valued in terms of liquidity? The answer is that if these guys come in, they come in with their own players. And the more players you have in the market, the less the incremental value of the liquidity injection.

And you'll see this summarized in a minute. Valued players are players who are in the market when the market is thin, on average, in expectation. And when people in the market don't have a lot of income, when the risk is high, when they're valued a lot in the sense of lambda weights, yeah, so we changed the environment to have a more general variance-covariance matrix, not necessarily identical Pareto weights, not necessarily the same mean incomes. And they're not IID, heterogeneous preferences. The utility functions don't have to be all alike. And there are shocks to these, quote, fundamentals.

So this is what I was already beginning to enunciate. Agents are more central if the market is small, if the average volatility is high, or income shocks are positively correlated. I mean, when income shocks are positively correlated, there's a lot of aggregate risks. So you can't use the IID aspect adding up over independent draws to get the mutual fund benefit. So even if the market is large but the incomes are highly correlated, there's a lot of risk, and hence an agent could be more central if he's likely to be in those clusters. This has to do with the Pareto weights, the average Pareto weights of those participating in the market. The average agent is poor, the incomes are low, or the average of risk aversion is high. So all those things I think are intuitive.

So let's talk about positive counterparts to financial centrality. And there are two pieces, and one piece of it is really sparse but important. In a constrained-- means constrained by the environment with the shocks-- complete market implementation of the optimal risk sharing contract. OK, so what are we doing? We're taking the planning problem in the fragmented environment, solving it. So with a smart contract, we implement the risk sharing. They're committed to do what they said they're going to do. And we've been focusing on the normative part of solving the planner's problem, which is equivalent with determining the constrained Pareto optimal allocations.

We could talk about a market where there are financial assets. We can decentralize the risk sharing environment where they're entering into risk contracts, and those contracts are priced, contract determining premia and indemnity and so on, and the price being what you would pay to enter into that contract. And there could be things like bonds. So a bond would be a financial asset that pays off, maybe all the time, but maybe it's contingent. So you could imagine a bond that pays off when some event happens and 0 otherwise. And the event could be a hurricane or whatever, but here the event could be having something to do with the market, and in particular having something to do with whether agent i is in the market.

So you can imagine a bond that pays off if you're the holder only when agent i is in the market. And that bond, the price of that bond, is exactly our measure of the value of liquidity of trader i. And maybe you could imagine how to derive this because we did those qs's, the qs's, the shadow price of the state contingent resource constraint, and we were summing up over all those prices. It's like if you're familiar with arrow indexation of securities, you can price any security you want. You just pick the states over which it pays out, and then you add up over the pure bonds that pay 1, 0 depending on the state. You add up over all the states that are relevant, and that's what-- it's just this somewhat weird state here is that agent i is participating. You could call it a personalized bond.

Obviously, here the exogeneity matters. We don't want to enter into people entering into a trade. If i shows up and then i gets to choose whether he's showing up or not, that could be a very bad outcome. It could be all kinds of moral hazard considerations. But mainly this is to share conceptually what our measure of the value of liquidity, or our measure of financial centrality is not an alien concept. It's basically standard asset pricing formula.

This one is another interpretation, and it's if agents are getting together ex-ante, before anything happens, to bargain over who's going to get what, when. So we can think about-- I think it starts on the next slide. Let me-- not really written out, so let me just say it in English. The Nash bargaining solution would be determined by what people would threaten. You have a higher bargaining position if your threat point is high. So in this case, the threat is you're not going to participate, and the utility value of that is autarky. So the terms in the objective function in the Nash bargaining objective function is the product of the difference of each agent's utility relative to autarky. It's that raised to a power and multiplied. So that's the Nash version of bargaining, and there's this-- I don't think it's [? Kaliaj ?] Minsky, but there's a guy beginning with K, whose name I'm forgetting, that also works.

So let's actually spell it out a little more. Suppose we have those constant absolute risk averse preferences. Then we determine a Pareto efficient allocation. So this is a risk sharing formula. Because it's absolute risk aversion, they're all going to get the mean income plus or minus this intercept. And this intercept is high for agents with a high Pareto weight, lambda i, relative to the average weight of the others. And again, it's all conditioned in this case on xi, the market participation, because you're only sharing risk in these segmented markets.

So this is like an intercept. If you were going to run a risk sharing regression, and we do all the time-- I'll show you that this term is going to be like a fixed effect. Oh, that would have been a better slide. I talked you through the English. So you run this regression in data over time of the consumption of agent i with an intercept. You allow for the counterfactual, which should be 0, that individual income matters. You allow a time-varying fixed effect. This captures the aggregate risk. And epsilon is an orthogonal error term, could be measurement error or something else.

So we want to focus on this intercept. We ran this regression in my Thai data, and so we know what these intercepts are. And yeah, intuitively, the higher is your lambda weight, the higher is this intercept relative to the lambda weights of others. And if you do this Nash bargaining thing, you get an expression for the lambda, except it's a bit more complicated because lambda appears on both sides of the equation.

This is like financial centrality. The higher is the value of liquidity, the higher should be lambda. Why? Because if you're financially central or you have a high liquidity value, that's socially beneficial. That's what the number represents. So if you're bargaining ex-ante, you should get more out of the pie. You have more sway in determining the social contract. But as you've seen, these formulas for financial centrality depend on the lambda's too, because we're maximizing a lambda weighted sum of utilities. So we really, in principle, have to come to grips with the fixed point aspect of the lambdas.

Instead, what we do is go back to the basics. As I was saying, who are the valued players? They're in the market when the market is small, and they're in the market when there's a lot of risk and so on and so forth. And so in the Thai data, we did that. We know who's borrowing and lending and giving gifts to whom in a given month. Those are the edges. We see those things. And we know how many people were doing that as well as-- but let's imagine they're all connected to each other or some people just aren't there. That's the way we deal with the data.

So this is the covariance of the participation shock being there, say, with the number of agents who are also there. So these are objects that are varying, each one of them over time. They're random variables. We take the covariance of that. That's this rho xi for the participation shock. And this one is having to do with the variance of these two random variables, the participation shocks, and the average variance of the players who are there.

For each player, we have their incomes, each household, so we can use all of the data to determine the variance-covariance structure of the incomes. And then we look at those variance-covariance numbers, the sigma ij's, conditioned on the participation shocks. And that becomes this random variable, and then we look at the covariance with h and i's participation to get the variance value of agent i. And then we run this regression. These [? Ai's ?] were the intercept terms from before, and we know how to run that regression. So this isn't deep science. There's even disagreements or questions we have about whether we should include this, but I can't help but be wanting to share it with you because it turned out quite remarkably well in the sense that the higher is the intercept-- the higher are these two terms, the participation term and the variance term, the higher is the intercept.

And what is it saying? It's saying, somehow or other, it's as if this community is aware not only of the income shocks, but the participation shocks and the fact that not everyone is around all the time. So if someone is a steadfast moneylender or gift giver, that's a good thing, and they're valued. And they get a higher average consumption than other people do, as if they're getting a premium. I mean, de facto, of course, they're not entering into an explicit insurance contract, but they are, we know, in the data getting higher consumption pretty much along the lines that the theory would predict.

Now, I would love to do something like this in the New York financial markets or in the Swiss markets. I've gone fishing. Those data sets are pretty confidential, and it's not obvious either exactly which variables we would want to measure. We should probably have to change the theory, but we do think that the theory is capturing something important about stochastic markets, and that that's a fundamental risk, and liquidity injections could help. So yes.

STUDENT: If some policy like this are implemented in some markets, will that create some more [? harder? ?] [INAUDIBLE] those who are [INAUDIBLE]

ROBERT M. TOWNSEND: So we have partial answers to that in the paper. We have a model of market participation where it's costly, and you can choose to go or not, depending on who else you think is going. And we characterize a Nash equilibrium and the impact on effort, or whether or not you're making an effort to be in the market is the moral hazard part.

But on the one hand, you're not the first to ask us this question. That's why we include in the paper a discussion about endogenous participation. On the other hand, I don't think we got far enough along with it. It always seems to be very special to depend on the particular structure of a model. And you can imagine lots of different models.

OK, so the other branch, contagion, that's the title of the book. It actually has an exclamation mark in the title, *Contagion! Systemic Risk in Financial Networks.* This is a book reviewing the literature and contributing to the literature that financial contagion is analogous to disease. Damaging financial crises are better understood by this conceptualization of the problem, and the policy implication would be manage the systemic risk, as you would in epidemiology, to identify situations where the danger is high and make targeted interventions or ex-ante policy rules to try to mitigate systemic risk.

So this book, it's on the reading list. It would always be tempting to choose as something you might want to write up. The book presents a mathematical framework for the transmission channels of damaging shocks that cause instability in [? financial. ?] This all really took shape with the great financial crisis. And then somehow they made this leap to networks, and it stuck. They're still stuck on it.

There's a lot of disclaimers. So there's the systemic risk basics in the introductory chapter, the features of past historical financial crises, the characteristics of banks, their balance sheets, how they're regulated, and then these cascade models, network effects, like default contagion, liquidity hoarding by banks and through their interbank exposure. For example, if there's wind in the air about a crisis brewing and it starts to impact, people may sell their assets to mitigate the problem, to gain some value, because the markets are thin. And that exacerbates the problem because it's like a fire sale, and so the asset prices are going down and everyone else has fewer buffers. And there are other kinds of cascade mechanisms.

So it becomes very real when you read the book. I'm not taking away from it at all. Instead, what I want to compare is this notion of, let's make sure traders have limited exposure, especially the bigger ones, because if something were to go wrong, they're going to spread their problems around the entire financial market. So it's kind of an ex-post measure, but it's guiding ex-ante policy. And it leads you to different conclusion. The spirit of what we're doing is to try to allow ex-ante policy to mitigate adverse shocks, to keep the market thick or as functioning as well. Now, that's not to say you have to choose one or the other. There's a whole branch here of ex-ante thoughts about risk sharing versus ex-post, unexpected shocks, and maybe deriving some hybrid policy that is some kind of policy that reflects both strands. There seems to be nothing like that in the literature that we are aware of.

In the few minutes that remain, I want to live up to my hope and promise to take you through a third model we really don't have time to do, and that's back to the US repo markets. So just a word about the repo markets. As I was saying earlier, you've got these money market funds looking for with, say, excess liquidity, and they're happy to lend it out for a return. They're restricted by regulation and who they are allowed to deal with. They're not allowed to trade directly with hedge funds, but they can trade with broker-dealers, especially if broker-dealers are associated with large banks. And they take treasuries as the collateral for their loans.

The other side of it are the hedge funds and pension funds, and they're involved in all kinds of security trades and derivatives and so on. And they want to finance that. So they're short of liquidity, and they're very eager to borrow, even overnight, in order to adjust their asset position. So that's the basis of the repo market, and the dealers are in between. They are the nodes in the network that are connected to these clients, the mutual funds or, say, the hedge funds.

Well, one thought, how are they coordinating if they're not in touch with one another? How thick will the market be depends on the arrangements that money market funds have with their clients or risk managers have with their clients, so there's a bit of an anticipation problem. And indeed, formally, you can show there are multiple equilibria in the repo market. That should remind you of the paper I went through in lecture 3 with Neil Wallace, having to do with the circulating IOUs and the fact that there were multiple perfectly valid configurations, but it requires some kind of coordination, in that case, among who was issuing assets, or in this case, how much of a position to take on behalf of your client before you enter the inter-dealer repo market.

Second aspect has to do with the regulation and this financial contagion. As a result of the great financial crisis, it was judged that simple risk-weighted portfolios are not enough to allow a true measure of the risk of financial players. You just want, in addition, to look at the size of their balance sheets. So what's going on with these broker-dealers? And then you should be thinking about Tomaž's diagrams.

So they're broker-dealers, so if they're in touch with a mutual fund, that means they're kind of accepting money to pass along in return for passing the securities in the other direction. But they are connected to another broker who's taking the same position. It's called rehypothecation. Basically, you're borrowing and lending simultaneously, but that expands your balance sheet. The assets and the liabilities match. But under Basel regulations, that's a bad thing, that you've expanded the balance sheet. So the policy directive was that these broker-dealers are limited in the amount of their participation in the repo market, and then you get these shortfalls of liquidity.

And the big banks are saying we can't intermediate more because it would violate-- if you take it at face value, they may be getting what they want in the sense that they succeed in lobbying the Fed to inject more liquidity into the market. In fact, they don't need to lobby. The repo rate goes through the roof. It hit 10% at one point when the Federal Reserve policy rate was, like, 3% and the repo rate was 9% or 10% in October of 2023, I think. I don't remember exactly the date. And the Fed enters as a node, providing liquidity to the system as a lender, reverse repo.

The third thing I want to say is there a policy to get around to mitigate this balance sheet regulation? Well, it's kind of weird because you're both a borrower nor a lender, and you're passing it right along the chain. So like Tomaž's paper that he presented, you can imagine being explicit about the interconnectedness and only worry about the source and the sink as the source and point of the flow. So you can, and we are, taking the data from the repo market. And fortunately, we have collaboration with the US Treasury on this. We are able to do the flow decomposition algorithm, one of them, and look at all these interconnected edges among the dealers, and separate out cycles, closed cycles from chains.

And the chains gives you a way to order the source and the sink, but we don't eliminate the underwriting. We're saying even if this stuff isn't on the balance sheet, they've still entered into a bilateral agreement to absorb the risk if the person they passed the loan to does not pay it back. So they're willing to take that risk, but that's still a change relative to-- so I had thought for quite a while about taking through the details about all of this because it links up closely, although not identical with what Tomaž is doing. But enough said about that for today. All right. Thank you very much.