0:00

Markus Brunnermeier: Nice to see you all again. Thanks for coming back to another webinar organized by Princeton for everyone worldwide. We're very happy to have Ben Golub with us from Northwestern visiting Stanford at the moment. Hi, Ben.

Ben Golub: Hi. Great to be here.

Markus Brunnermeier: Ben will talk about networks in contagion and resilience. And the topic of resilience is close to my heart. And we'll learn more about how it's connected to networks. So what is resilience? Resilience is essentially about bouncing back after a shock. And it's about the characteristics of a stochastic process. And it depends how easily the underlying system or network is adjusting or adapting such that the system can actually bounce back. So it's all relative to not adapting. So you can think of an adaptive system. There's some exogenous process hitting some system or some network, and it affects different drastic processes, endogenous processes, and it depends how the system adapts and shifts and changes and alters this endogenous process relative to a system which does not adapt. And if it adapts and improves the situation and bounces back, then positive resilience. It might also get worse. If it diverges, then it's negative resilience. So that's also the question, how should we measure resilience? And that's one measure I proposed in my presidential address for the American Finance Association, is you can argue that if there's a shock, the initial shock is just a one-time permanent shock, but if then the endogenous process is actually coming back and it's bouncing back, one potential, that's a benefit, it's beneficial, and you could integrate over all the benefits over time, so a discounted present value of these benefits, like this shaded area in a discounted fashion, would be one possible measure of resilience. Now if this adaptation of the system goes the other way around and makes things even worse, then the resilience is negative. So you have this red shaded area, and it becomes negative. It could also be that actually the situation gets so bad that you get an adverse feedback loop, and things get much, much worse. And then the discounted area becomes minus infinity. So that's, in particular, if you get in some adverse feedback loops. What are these adverse feedback loops? They are typically characterized. There might be a shock which hits some person 2, let's say. The person 2 reacts, and then because of the reaction, it might spill back to person 1. And there's always a spillover and externality. And then there's a reaction which depends on strategic complementarities. With strategic substitutes, then it does not amplify. If there's strategic complementarities, it amplifies. So it can spill back, and it can spill over as an externality. Then you react, and the reaction might cause some spillback. And then it gets further reaction. You get these adverse spillover spirals going on. Now, how much these reactions will spill over and go through a network depends very much on the network structure. And we will learn more about this today. Here's just a simple picture. We have different network structures. There's a centralized network structure, decentralized, and they have a distributed network structure. And if you have a centralized network structure and you knock out, you have a shock to the center, then the whole network is going down. So the whole network is not very resilient. At the other extreme, you might have a distributed network structure. And if you knock out one node, it does not create a big damage because it is leading back. The system is resilient. It can bounce back and take care of it. If you think about a terrorist network, it's typically organized this way.

3:49
If you think about our brains, they're organized this way. If parts of the brains are knocked out, the other parts of the brains take over. So that's interesting networks to understand. Also economics, also networks to some extent. What I like about resilience is also that micro-resilience is not necessarily good for macro-resilience. So if some sub-network is very resilient, it might actually make the whole macro system less resilient. And so that's sometimes what you would like to have. You might have to shut down, make certain sub-parts or sub-networks, shut them down in order to save the rest of the network. And so the resilience of each part of it might actually not be good for the resilience of the whole of it. So there's some fallacy of composition in the resilience perspective. So all of these things, hopefully, Ben will put this in a better perspective today with regard to networks. I think there's a lot of additional research to be done and a lot of exciting stuff hopefully happening in this area. So then we have the questions, the poll questions you answered gracefully, which Ben put forward, and here are the questions and the answers you gave us. Are contagion cascades of domino effects more consequential in the real economy, the interaction between the economy and the ecology, or the financial sector? And the answers were 17 percent said real economy, the interaction between economy and ecology was 34 percent, and the financial sector was 49 percent. So the majority, almost 50 percent, is for the financial sector. The second question was, are there, there are important parts of the economy at the edge where a lack of disruption, might not be a big one, will cause a massive damage. Is this the case or not? And the answer is overwhelmingly yes, 95 percent say it's yes, and only 5 percent say it’s no. Of course it includes probably butterfly effects and other things, chaos as well. The third question was, is markets generally, do they provide adequate incentives for investments in resilience, so the markets itself, the market economy is doing that. And the answers were, yes, 15% thought the market would do it. There's no market failure. No, 30%, so double. And in some economies, but not the most, so it depends. That's the typical answer for economists. That was 55%. And finally, are networks too difficult to model? Should we give up? And 11% thought, yes, it's just too difficult. But most people are hopeful, so almost 90% or 89% thought, yes, we can do it. We just need people like Ben to plow ahead, and we will have very good insights on the network model. And with this, I pass on the mic to Ben, and we are looking forward to your presentation. Thanks again, Ben.

Ben Golub: Thank you so much, Markus. I really appreciate this opportunity. As we've discussed, I think that's an amazing kind of public good here that you have, and I am really excited to present an overview that is, so this isn't a research, a talk on a particular paper, it's, I sort of tried to think about what have we learned from recent networks research that could be a useful input into thinking about contagion resilience, and I'm excited to tell you what I think we know, but as we'll see, there's really many more challenges than firm known answers, and so we have the tools, and I'm excited to tell you why I think this is an amazing area for researchers to contribute to. So as Markus already kind of said here, and I, and very kind of completely, and I think convincingly in his presidential address, networks come up very naturally when we start thinking about the mechanics of resilience in a detailed way, because they mediate disruption, they're key to understand in thinking about which disruptions will be particularly bad and how resilience can be obtained.

8:03
And the reason networks play that role, of course, is that they describe very general arbitrary externalities as well as the kind of spillovers through prices or strategic incentives that have been at the forefront of many different discussions recently. So I'll just give you some examples. When we think about supply and demand shocks traveling through the economy, the kind of most important economic network is arguably the real production network where firms buy inputs and make stuff with them and so forth, whether it be goods or services. And we know that those networks propagate shocks, but there are also other choices that are – we maybe think of as more strategic than standard market mediated. Things like investment decisions and new technologies by firms or in innovation. In the application that many people here found compelling, there's been, as you probably know, a lot of interest in cascades of default and other financial distress through things like fire sales. So, you can have direct contagion of firms not paying and then other firms can pay their debts, but you can also have more subtle effects like fire sales and liquidity crunches. And there's been arguments that especially in over-the-counter markets, we really need to pay attention to the network structure to make good predictions and good policies for those things. So, there's a whole, excuse me. Really sorry about that. So, there's applications to price pass-through. So, in markets that are imperfectly competitive, oligopolies, strategic pricing plays a big role and pass-through becomes much more complicated because it happens at the level of a whole network of firms. And then finally, and recently, in ongoing geopolitical tensions during the Covid crisis, we had disruptions to the fabric of the world logistics network and relationships among firms caused by economic shocks. And again, it seems important to understand what role the heterogeneity that we see in relationships plays in that. So one thing, as I was watching Markus's presidential address on the importance of resilience, as he and I discussed a week ago, I noticed there were a few robust, sturdy, standard models that we all learn in our core classes. And then by putting them together, we could begin to understand various things that were important for macro resilience and the interaction of finance and the real economy and so forth. Networks are clearly, as I've argued, a big part of how we need to think about these things. But we haven't really diffused in the same way some core basic models, even static models, that are going to be tools for thinking about these things. But I think just as in with our beloved, you know, Lucas trees and so forth, simple models are a very powerful source of intuitions and sufficient statistics. So I want to give you a flavor in this talk of two particular types of simple models that are very useful for me in trying to think about what role networks can play in thinking about resilience. And they're also models that we're still studying, you know, that one reason networks are so exciting is we are at a much earlier stage of development. So even the basic questions raise a lot of open research questions that are exciting. So I'm really going to tell you, hopefully over time, it'll be very clear what I think these two core models are, or two modeling paradigms, and we'll explore applications. So let me just jump right in. I'm going to use this picture for illustration for a few things. And so I'll remind you of a formula that I'm sure you almost all know. So we can draw,

when we have a network we often depict it as nodes and links. The links can have weights on them which can be like the size of a relationship and the adjacency matrix here has a one in entry ij if nodes i and j are linked.

12:23
This picture, and I'll generally stick with this type of example in the talk for my main illustration, it's undirected so links go both ways. Of course you could use the same type of matrix to represent directed relationships. So let me give you an economic interaction and I'll talk through that in both words and symbols. This is a game, this is the kind of simplest network model that has interesting strategic structure and I'll tell it through an R&D story. So imagine you have a bunch of firms and they have to invest in some new technology like maybe blockchain and cryptocurrency and developing, you know, competencies for that. Maybe it's a research area where you have agglomeration externalities. And so basically, when more firms invest, the workers are better, and again, it's easier to innovate in that field. So the kind of model that's, I'll tell you the literature, but I'll just tell you first what the simple game is. Firms payoffs are, there's a cost of making this investment. So that's this constant, this linear marginal cost over here, quadratic cost. And then the returns look like, well, per unit of your own effort, your own investment, you get some standalone return, which is what you get regardless of what anybody else does. And then you get some return, which is a spillover term coming from interactions with others. So basically, the more others are investing, the more you're going to get, the more your marginal return is going to be. So here I'm focusing on an example where these spillovers are all weakly positive and beta, the coefficient, is also a positive number. It doesn't have to be that way, but let's start with that example to think about. And so the reason people love this simple little model is that it's very easy to solve. If you just take the first order conditions, optimize everybody's investment, you get a beautiful little equation. You have to make an assumption that the spillovers aren't too big. Technically, the assumption is that the biggest eigenvalue of this kind of total spillover matrix isn't too big, but if that holds, then you can write how much everybody invests in this kind of pure effects way. So an econometrician would call this a reduced form pure effect equation, where how much firm i invests comes from, well, their standalone incentive, so this is going to happen no matter what anybody does, but also the average investment in their neighborhood. So this is, take all their neighbors, add up how much they're investing, weighted by how much you're linked to them, and that's going to also contribute to your investment. That has endogenous things on the right and the left. When we solve that little system, we get an expression for equilibrium that some people call a structural pure effect equation. All it is, is you just move all the exogenous things, which in this case we think of as the network and the basic incentives, the standalone incentives to the right-hand side, and we get a nice simple expression for the equilibrium behavior. So let me give you, I'm going to just tell you this.

Markus Brunnermeier: So for example – Can I just ask a quick question? If the different weights of the different links, the little g's would just be smaller or larger, I guess. And it would work the same way. Generalization is very straightforward.

15:50
Ben Golub: Exactly. So this works for arbitrary weighted networks and even directed networks, all these equations hold. So I want to give you, for many of you who have played with this kind of model, for those of you who haven't, let me give you, I've reproduced the solution here. But, you know, when we look at the solution, some people have a feeling for it, but you might think, well, what is this really telling me? What is this really capturing? So let me do a little thought experiment where we imagine that we increase Mr. Five over here. We take his basic incentive and we give him maybe a little subsidy for research. So he's more into investing in this area. And the direct effect, what I'll call the zero order effect, is of course that it increases his own action directly because his incentive got positively hit. But then when we look at people who care about his effort, they're going to now see, oh, basically this guy who's complimentary to me is investing harder. Let me also invest more. And that's the first effect. And we get that by algebraically, we just take the matrix W, which here is beta times G. It's the strength of these spillovers. And we just add that. That's the first order direct effect. Now, of course, you can see where this is going. These guys are just like the first guy in a sense. They increased their action. This is sort of rounds of best response. And now other people see that. actions have increased, including our hero here, the first guy, has these effects reflected back at him because he is now increasing his action due to the increase that he caused in his friend's action. And these ripple effects, but they also go out from him, right? And when we aggregate up all these ripple effects, you see we're getting a summation of these powers of w. And so that gives us a formula, which is this standard Neumann series or Leontief type formula for the total effect of a shock to incentives. And so when we now, well, one thing, we could have done this with anyone. We did it for Mr. Five and we found this effect. We could, of course, compute it as a number. And what I've done here is I've just scaled the node sizes proportional to how much the total action would increase from a little equal size shock to their own private incentive. And so we see that Mr. Five here, he was chosen for this illustration because he's quite central. He has the biggest centrality. We call it centrality, which is – some people, I was talking to someone yesterday, who's a little annoyed by it, because you don't have to be connected to many people,  It's a centrality capturing these direct and indirect effects – and you see that proximity to Mr. 5 here is a pretty good predictor of how central you are. And that's a general phenomenon. You're going to be central if you're adjacent to central people, because you transmit a shock to them. And because they're central, they transmit it to everyone. So these are statistics that people love. And just to nail home the policy point, if you care about the sum of actions for some reason, like you're a government that just wants to increase investment, and you have a dollar of subsidy to give away, you should basically give it to Mr. 5 here, because that will most increase the total. And that's what I would call a very simple networks insight. It's one that was developed maybe around 20 years ago. It fully percolated into the economics of these things. And so I'll just emphasize the name here. I won't keep this slide up for a long time. is a thing called Bonacich centrality, which is the name, the official name for this statistic of how key are you. In a famous paper it was called the problem of finding the key player. And you can compute players' centrality or how key they are by doing the simple algebra with your matrix or using the inverse formula if you'd like. But we have a very explicit understanding of exactly how the centrality statistic works. Okay, so let me zoom out a little bit and talk about where this fits into economic

thought. So of course these ideas didn't start with network games. They go back to Leontief and applied math somewhat before that.

19:49
But Leontief was of course interested in these things in the production network, in physical supply or demand shocks propagating in very simple economic models, and here's Leontief with one of his matrices depicted behind him. It's been, it was revived and I'll, we'll talk more about why in a way micro theory came back to these ideas before macro. It's kind of interesting because you might think that, you know, you would have just been developing these in the spirit of Leontief. There are good reasons why that took a little while. And the paper that I have been sort of teaching just now is this famous Who's Who in Networks paper that was a 2006 paper in Econometrica. It's also related to a very active literature in econometrics, which develops these ideas for pure effect applications. So this is this, this key player paper.

Markus Brunnermeier: Can you say a little bit more how econometrics is related to the network?

Ben Golub: Yes, absolutely. So, I mean, basically people want to, people are interested in pure effects. like just in many applied problems, like Raj Chetty, for example, has recent work that is very exciting about how neighborhoods really affect you. And even within the neighborhood, the most closely related age cohorts, how their parents are doing in terms of employment seems to have big effects on economic mobility. So those are very big, those seem important. And there are fundamental challenges in econometrics, the Mansky problem and so forth with identifying peer effects by introducing enough heterogeneity into the model of how shocks propagate, networks have helped a lot with identification. And so that's been a power. So they've been very excited about this. So let me just, I will not do this slide justice. I hope if you're interested in the recording, you can pause it, but I'll give you, what I wanna emphasize here is that in the second generation of taking this Leontief idea and pushing it as far as it'll go, in many different applications, we. developed ways of thinking about these ripple effects in economically interesting ways. So I've just described the network game application. There is a production network application, which we'll be talking about in a second, where the firm, where the nodes can again be thought of as firms or aggregated sectors. The links now represent technological relationships. The outcomes, rather than being this abstract total activity, will be things like GDP. But there will be a very similar message that a kind of centrality statistic will affect how much shocks to you translate into this aggregate outcome. Financial networks, which many of you are interested in as an application, have been studied with very similar methods. So there the links are financial obligations, some kind of interdependencies of my balance sheet on your solvency or your cash on hand or your liquidity or things like this. And there we can look at effects that are financial outcomes. And this has been developed in two papers that I won't have the chance to talk about a lot today, but there's a paper by Osamoglu and co-authors and a paper by Matt Elliott and Matt Jackson and me, which develops these ideas for financial applications. And finally, they've even been used in a financial, in an opinion dynamics paper to think about, you know, things like rumors in financial markets with a paper that I love by, oops, by DeMarzo Bionis and Zwiebel in what's called the DeGroote model. So lots of exciting, if you like these ripple effects, a lot has been done with them, but I want to raise a question.

23:37
Markus Brunnermeier: Can I ask a quick question about, you mentioned it's the game structure that the people behave strategically. Could you also just let the number of players go to infinities or you have different islands or different sectors in the economy? Like each node is just the whole sector and everybody is essentially a network taker.

Ben Golub: Yes. So I love that question actually, that has been an idea that has been recently developed. So that's a great idea in the sense that you could hope to, in a sense, zoom out from individual strategic effects. It turns out that externalities will still matter. Even when people are network takers, externalities are still there. But the macro perspective has been a very active area of recent research of the boundary of networks and statistics. And I'll talk about that in about three or four slides. OK. So at this point, for this talk, I am going to treat the most important applications as the real economy contagion that Leontief was originally interested in, as well as financial networks. And focusing on the real economy, I would like to tell you that the real economy is about a sort of crisis that happened in this literature and kind of the reason that it took a long time for macro to come back to these issues. And that issue is why care about centrality? So you can have this great theory, and it was developed in its kind of modern form in a second generation paper, right? It's not Leontief, but building on Leontief, people put, you know, the same kind of network mechanics I've been telling you about in a canonical general equilibrium production model. So basically, if you take a model with, let's say, Cobb-Douglas production, you log/linearize everything, and you look at the equations, you get very similar equations to the ones I showed you for the A's, but instead you're working with log prices. And you get a very similar theory where, you know, it would take longer to kind of fully present, but the mathematical ideas are perfectly analogous. But a sort of bugbear for this literature is the observation that the kind of shocks I've been telling you about, where you, let's say, increase someone's productivity a little bit, end up being only second order for welfare. And so this argument was made very, very compellingly. And I encourage you to read the introduction, at least of a beautiful paper by Baqaee and Farhi, which I value most for the criticism that it kind of collected and articulated of the basic Leontief approach, which is that if you take an efficient model, a model where the market equilibrium is efficient, and you do these ripple effect studies, all the welfare effects you get are second order, because the economy starts out efficient. And so by kind of a Harberger idea, nothing really matters that much in a simple static equilibrium model. And even worse than that, in a way, even if you do care about second order effects, as this 2019 paper argues you should – actually, let me revise that. So, you know, for second order effects, everything can matter. But if you only care about understanding first order outcomes, for whatever reason, like you want to know how much GDP will change, if you give a technological shock to a firm, you don't really need to look at the network. You can read the centrality statistics that matter off of the national accounts in a very simple way, like the sizes of sectors are a very good proxy for their importance. And so this was all a little bit sad if what your goal is to do is to study networks. Of course, that maybe shouldn't be our goal, but these guys pointed out that if you want to understand where networks matter, you need to go beyond a simple, static, modern version of Leontief's theory. And that's what happened next. So I want to, you know, I'm, I'm going to, I am spotlighting in this real economy discussion work of Bakay and Fari, I think it has been very generative for pushing our understanding of this because they made this criticism and then they gave what I think is a very exciting kind of way forward, which is what I would call third generation production network research.

27:56
So what, what they've done here is they've said, you know, focus on, let's first get clear on why these things matter for welfare and study these network effects in that kind of world. And so the key insight, and I really want to pause on this insight, is that you have to start with a distorted economy for these effects to be interesting in static models. Or, and so there's sort of two directions, right? You can study distortions even statically, and that's what we're going to talk about. Or you can go think about dynamics, where distortions are much more, as Markus has emphasized in his presidential address and in his introduction, there, you know, dynamics naturally make everything a lot more interesting. So I'll talk more about the second part, but let me tell you about the static view. And a kind of reference point that I like is if you happen to know, um, work on social value of information, uh, there's a paper by Angeletos and Pavan about how information releases affect markets, whose point is exactly that the wealth, the nature of the externalities is key to understanding how information is going to affect things. So you really need to know what kind of externalities your game features in order to know whether, for example, public information is good or bad. And here, similarly, the nature of distortions is going to be key to understanding welfare effects. So what, and now let me just put up the paper. So in a 2020 QJ paper, Baqaee and Farhi take a model with sort of very flexible distortions, markups, for example, that distort the allocation in general equilibrium. And then what they try to do is kind of understand what are the first order welfare effects of shocks. So now they're not, not zero, so that's good. And they write beautiful formulas kind of explaining the welfare consequences of shocks. And I don't have, as I'll sort of be clear about in a moment, I don't really have, I'm not, going to be able to fully teach these formulas. I can give you a flavor of them, which is that basically the effect on output and also the effect on welfare comes from a technology part, which is sort of a boring part that exists even in the second generation models. But more interestingly, changes in allocative efficiency start happening. Because your economy is distorted, when shocks happen, they're going to change how these distortions work and the extent of them. And so they are able to write formulas, and I'm giving you just a flavor here, but they define centrality measures that are very much like what we've been studying. And they study them, but it becomes very important, for example, how those centrality measures change in the course of response to a shock. And so their paper is beautiful and very general. But I have to say the most, and this is sort of an important segue, I get a lot of the thing people ask me most about this paper is, you know, their formulas, when they unwrap that little formula I showed you, they're very big and they're very complicated and they involve like the Leontief matrix on steroids. They're very forbidding to look at. If you just, I was tempted to put a screenshot. I won't do that. But the question is, you know, what intuitions can we give you? People say, Ben, you're a network theorist. Can you help me understand how the network really matters, right? What kinds of shocks are best and worst? And they do have some intuitions, some good economic

intuitions, but going back to basics in network theory, we've realized we don't really have a good grip in these distorted inefficient economies, on how the network structure matters for the welfare impact of shocks.

31:49
And so I'm gonna show you, I'm gonna come back to the baby model that I've presented, strip things down and show you how welfare effects are very interesting, even without all the macroeconomic richness in this third generation literature.

Markus Brunnermeier: And then I can ask you, so does it matter how the distortion is? So could it be that because of monopoly power, there's too little production, or it could be some because of negative externalities, there's too much production, and then the externality measures a different one, which matters?

Ben Golub: Exactly. So that is, yeah, that is exactly how I would like to think of it, that in different economic situations, you know, the nature of externalities and strategic spillovers will be different. And there's not a canonical model for macro, we get some, you know, familiar specifications. But the point is exactly that, you know, the way the network matters will really depend on the answers to your question. And we need a flexible toolkit to marry the network structure and those economic fundamentals. That's kind of exactly what I'd like to deliver. And then a very related question is, you know, when you come up with the right, for example, if you get to shape the shock to some extent. If you get to pick your favorite type of change to the economy, how much can you trust your answer? Because to Markus's earlier point, sometimes we may be able to zoom out and trust that we know some aggregate patterns in the economy. But if I really ask you, how much do you trust that you know the elasticity of substitution between these two little goods? It might be quite hard to know anything with certainty. So the network literature, we're going to come back now from these very interesting macro models, back a little bit to baby land with a very simple theory model we can keep in our heads fully, and see some forces and some intuitions that are going to shed light on these things and show you that there's still very much open questions, especially this important statistical question. So I want you to keep that in your mind. But let me first tell you a little more about the mechanics of welfare. And so this is an idea I'm excited about. So I'm going to come back to this game and to Markus's point – before I told you, let's think about this beta as a positive number. And now I'm going to say, no, it can be negative too. So just for concreteness, let's keep the Gs non-negative. But now if I flip beta, I can make it a game of negative spillovers and negative externalities, right? And then I could vary the externality separately from the spillovers, which I'm not doing in this specification. So before, we talked about just mechanical spillovers and this mechanical question of if you want more action, what do you want to do? But now kind of inspired by this macro discussion, I want to ask a very natural question, which is what is the GDP of this little investment economy? If I just care about a firm's welfare according to the utilities that govern their actions, what is welfare? It turns out that the answer can be a bear. And this is an obstacle for the macro literature where you write these formulas with matrix inverses and you have trouble making sense of them. So that would be true here too, but borrowing a trick from the applied mathematicians in this paper that with Galeotti and Goyal, we have a trick for making the formula at least algebraically legible. So let's start with the algebra.

35:14
What we do is we take the network matrix and we diagonalize it, right? So first your linear algebra, what we do is we just write it as the sandwich. And that includes, that entails finding a basis of eigenvectors here for the ambient space of actions or whatever, that when you think of things in this basis, the network just looks like a diagonal matrix. And that idea is very useful for welfare, it turns out, because it gives us a formula for welfare. Let me explain what these ingredients are. So the underlying Bs here appearing in this formula are just the projections of the basic incentive vector onto different eigenvectors or principal components of the network. And the corresponding coefficients, these alphas are just nonlinear transformations of eigenvalues. It's a particular formula. It's not too important what it is, but this formula, for example, already tells us that the aspect of basic incentives that matters more for welfare under strategic complements with beta positive here, turns out to be the high eigenvalue pieces. So the parts of B that project a lot onto eigenvectors of G with high eigenvalues, those are important. And the ones with low eigenvalues are less important. And so a question I usually get here is, okay, that's great, Ben, you seem very excited, but how can we understand this? What's the economic intuition behind this? And so let me give you a bit of it. I obviously can only give you a taste of it, but to make this concrete, suppose a policymaker now comes along and thinks, well, I can reshape people's investments, investment incentives a little bit. I can subsidize some people and make them want to invest more. and I can discourage some other people and make them want to invest less, how would I want to use that power? So in this basic canonical case of positive spillovers, like the R&D problem with positive spillovers, it turns out what you would like to do, so as Markus said earlier, there's underinvestment in this kind of game, and so it should be that people are encouraged to invest more, and the right way to exploit the feedback effects through the network is to subsidize them in proportion to a certain other centrality measure, the first eigenvector, and that turns out to be the best of all the welfare interventions you can do under a kind of natural size of intervention constraint. But when we change the economics of the story and we make the spillovers negative and the strategic interactions strategic substitutes, so think of a game where you have price competition or some kind of competition where when your neighbor does better, you actually do worse and you're motivated to do less of whatever it is, then in that kind of game, targeting for utilitarian welfare is extremely different. You want to target neighbors in opposite ways. And if you get a subsidy, it should be that your neighbors get a tax. So one little story you can think of is like a public goods game, where if I'm encouraging some people to contribute more, but I also encourage the people that they're sharing a public good with to contribute more, my own intervention is crowding itself out. And so you want to arrange the intervention to be attentive to this. And what's amazing is that this is just captured by another piece of the spectral decomposition, the last eigenvector of the network. So while economically these two stories feel very different, they come out from different pieces of the linear algebra, and they really show how this kind of spectral approach to writing the welfare function gave us a tool that we can apply, in the paper, what we do is we show exactly how the shape of the spillovers and the externalities shapes what you want to care about. And it's not

necessarily the first or the last, but it kind of connects a very developed mathematical area of network statistics to these economic questions.

19:12
Markus Brunnermeier: So can I ask a more clarifying question? So I can see when there are strategic complements, you probably want to subsidize the more central guys more, the bigger green dots. But what's the intuition for strategic substitutes? I guess you always want, if you subsidize one, the neighbors you want to tax in order to, is this intuition? Why is it that one in six is going the opposite direction? Why, you know?

Ben Golub: Yeah, yeah, so I mean, basically, yeah, if I, you can think of it as, yeah, if you, if when you do more effort, it makes me less motivated, then if I target us both with subsidies, I'm acting at cross-purposes to myself. And what would be good, great, actually, is if I target someone two steps away from you, then together, we will, you know, discourage the guy in between us, but that will actually encourage each of us more. So that kind of takes advantage of the strategic spill over shape.

Markus Brunnermeier: But it could be that, you know, I don't know, this is the optimal one, that if you change the network a little bit, or if you change the situation a little bit, that all the red ones will be green, and all the green ones will be red, no, it might flip.

Ben Golub: Oh, yes, exactly. So, and that raises very, that's a very helpful comment for me, because, in fact, you know, we wrote this paper, and then Francesca Parise and Asu Ozdaglar came along, and they were interested in this current, very active area of statistics graph on theory, which is the theory of large, this network taking kind of theory where everybody is small, and they kind of effectively interact in a continuum network. And in the beautiful paper, they used our methods to do the statistical network taking version of this to think about what will things be on average. And what is true is that some of those strategic substitutes kind of problems are very sensitive to the details of local network structure. And so there's an intuition, at least that while you may be able to target these large scale patterns, if you really want to get right the crowding out at the local level, it may be much, much harder with data that you have access to. And so, that is exactly the agenda. And this is very active, I'm working on this in the IO application where firms are pricing, you know, what a policymaker can and can't reasonably target, there's the ideal targeting, if you knew everything exactly what you would want to do. And then as you're pointing out, there's the feasible targeting: what is statistically accessible. And the beauty of the eigenvalue approaches, statisticians know a lot about how the eigenvalues work under sampling. And so, for me, and this is something that I would like to sell to the young researchers, I think it's an area where you can really arbitrage ideas from statistics into, you know, economic models with very clear stakes. Okay, so

Markus Brunnermeier: I want to ask you one question again. So the eigenvalue approach, essentially, you take a very static perspective then. It's not percolating over time. Everything happens simultaneously.

Ben Golub: Exactly.

42:25
Markus Brunnermeier: It's not that the thing, oh, there's a first shock, then the next period is the next, the network, it translates over time, something, everything is in one shot. Yes. Or the second round, third round effects, everything is integrated already.

Ben Golub: Yes. And so I'm very excited to mention this. There's a paper by Ernest Liu and Aleh Tsyvinski that's just coming out in Restud now, which exactly picks up on that theme and says, look, in real life, when we think about these things unfolding over time, the best responses don't happen instantly. So for example, back to the production setting: a firm shuts down production for a while or stops placing orders, and that's quite fast to do, but resuming that relationship, spinning the factory back up, making the contracts work for the new economic environment, that can take a while. And so negative productivity shocks, and this has been studied very, you know, there's amazing evidence for this – negative productivity shocks leave scar tissue in the economy that lasts, you know, well beyond the initial shock. What these guys do, which I love, is they basically just build a kind of neutral time to resume specification, this asymmetry between shutting down and resuming, into a canonical production network model, and then they study how shocks last, exactly the question that you ask. And here they're looking at, you know, because in World War II, as you may know, there was actual economic warfare targeting key sectors to try to disrupt the German and Japanese economies. These guys took, this is in a previous draft of the paper, I'm not sure made it into the final version, but I love these pictures. What they did is they took an economic shock. Initially, it affects everything. So we're looking at how it affects, you know, how much the sourcing disruption lasts in various sourcing relationships. So we're looking at matrices really zoomed out. And you can see that what's left after you wait a while is what looks like something that looks like a low-ranked matrix, something that looks like it has low dimensional structure, right? And so they developed that point and showed that these top eigenvectors, the ones that mattered for strategic complements, are the most correlated and longest lasting in terms of the scar tissue. And they can, and even though to fully describe the shock, you need the implications of the shock, you need a lot of factors, a lot of eigenvectors, you can have a very low dimensional description of the shocks that really last. And so I think that's beautiful, you know, I'm a huge fan of that paper. And it's a very powerful idea that then you can, of course, map it back to the actual names of the sectors or firms. And you can see that more, more sectors that use a lot of produced inputs that, you know, end up taking, like motor vehicle manufacturing takes a long time to recover because it's very exposed to these long chains of scar tissue, you know, slow resumption. So I encourage you-

Markus Brunnermeier: There's always a recovery, essentially. You go back to the steady state. There's no, you never diverge. You don't- Yes. So they have no-

Ben Golub: Exactly. So their resilience, unlike you, your questions that you raised in your introduction and presidential address are qualitative. You know, are there cases where you don't resume? They set up the model so that it's gonna bounce back to steady state. But we're looking for forces, and this is a great segue to the second, to the shorter second part of the talk.

45:45
Are there forces where the shutdown is more durable and more extreme in a way, right? And so I'm gonna, I'll just put up very briefly the recap. I won't give this full recap because I'm eager to stay within a reasonable kind of time. But I want to basically say we started with these beautiful linear algebra ideas. We realized we really need to think about welfare at least as much as mechanics. In the current active research on this, both in macro and in network theory, understanding the welfare consequences of shocks is the big active research area. And right now, these papers that came up very nicely in this discussion by Parise and Ozdaglar and ongoing work that I have with a big team in the pricing setting, you know, we're very excited about using these mathematical insights to actually say what is and isn't practical. But as Markus just pointed out, you know, these theories aren't really about resilience. Even Liu and Tsyvinkski have only temporary resilience. And part of the issue with that is that in these very continuous smooth models, they have this tendency to want to come back to equilibrium, at least the way people actually work with them. And that's a limitation. So when we think about, for example, the 2008 financial crisis or even aspects of the recent Covid shock, it seems to have left some permanent scar tissue in economic relationships, or at least we want to think about that possibility. And so I do want to say, to be fair to the literature, that people have built in shutdown kind of margins, extensive margins, into these continuous propagation models. So in the paper with Elliot and Jackson on financial networks, we had bankruptcies that exogenously happen and burn a lot of value once things get bad enough. David Baqaee, his job market paper actually did a macro style extensive margin model where you have exit, which could be permanent. But these models really aren't, what Markus said toward the end of his introduction, about the fallacy of composition really bites here, you need more model than we've had so far to capture the idea that these discrete shutdowns can really be much bigger than the sum of their parts and lead to permanent scar tissue. And so it turns out that in network theory that is used to study viruses and, you know, network theory that's very important in physics, for example, they do have tools that are called percolation theory and related things, where we disable some nodes or links at random and we look at aggregate consequences. And once you make, once you include the more discrete phenomena, you get some very dramatic phase transitions and tipping points that are new and don't really come in the continuous models. So I want to sort of give you a flavor, since we only have about 18 minutes, I'm not going to be able to give you a full version of these ideas, but I can give you –

Markus Brunnermeier: take a little bit more time if you need it.

Ben Golub: Thank you. But I think I've structured it so I can give you at least a flavor, a spirit of these things. So I just want to, the application that motivated me to work on this issue is, you know, in 2021 the New York Times had a few beat reporters and other, the economists, you know, people were really tracking the supply chain disruption that affected the whole world economy. And to kind of summarize it, it seemed like there was real damage to the fabric of global logistics. Like a lot of ports were congested and they didn't have very efficient economic mechanisms for managing the queues.

49:42
And so everything was late and that seemed to really affect, you know, availability of a lot of consumer and industrial goods. And so we saw that, you know, the White House wrote reports on this. You saw huge drops in inventory, huge reports of shortages, especially in sectors like manufacturing and construction. And so Sue Helper and Evan Soltas wrote a beautiful report on the practical side of this as this was happening. And you saw huge jumps in shipping prices, suggesting that the world logistics network was getting sort of discontinuously stressed. And so, and this is just stuff we can measure on an immediate, you know, in things like shipping prices, but underneath these shipping relationships, there are also relational contracts and all sorts of more subtle things that can also be damaged when the economy gets hit by a shock. And so I felt like the models we have in the networks literature aren't really talking about this. And so I wanna tell you about a paper that I wrote, I'll tell you, I'll give you the citation in a second, but give you a sketch of some work I'm doing, I did, and I am continuing to work on regarding resilience. So what we're interested in is complex production where each good is produced in many steps of manufacturing, let's say. And the key thing is that important inputs are customized. So rather than buying stuff off the shelf, you need to have specific contracts and you rely on specific deliveries from your suppliers. And as I've already mentioned earlier, in the context of Ernest and Ali's paper, there's great empirical evidence that firms really are reliant on their counterparties, not just on the sector they're sourcing from, but the particular firms they're sourcing from. And so let me just give you a concrete example. You have a little airplane manufacturer and it's sourcing, let's say, I'm just going to focus on, I'm going to assume it has everything but two inputs, which is brakes and computers that it needs to manufacture the airplane. It needs to get them. And it may have multiple sourcing options, multiple suited suppliers that can call up and say, hey, can you ship me some brakes? But it can't go to Home Depot and just buy them there. And the brakes are in a similar situation. To make brakes, you need to buy some computer chips to control the activation of the brakes. And you might need to buy some other inputs. And these are also produced. And so you have this branching supply network that extends. Imagine it extends. I'll show you in a second, but oops, sorry. I'm sure you want...

Markus Brunnermeier: B2 is also connected to the D1, D2, or is it connected to a different disk computer?

Ben Golub: Say the question again?

Markus Brunnermeier: The B2 node, like the brakes, are connected to the same D and Cs?

Ben Golub: Exactly. So the whole, not the same ones, but different ones. Different, different ones. Yeah. So I'm going to assume, and this is a good point to raise this, that one very obvious network force that comes up here is that if there are Taiwanese microchips, upstream suppliers that everybody relies on, that's a very clear channel for fragility. But as you will see in this paper, I'm gonna focus on a maximally diversified kind of sourcing structure and show that nevertheless, there are certain other effects that we should still really worry about, okay?

53:18
And so the basic question I wanna ask is, if you know these dotted lines, they're sourcing relationships that may work or not in a given month, let's say. One way for them not to work is the shipment that they send you didn't get delivered. And so we are gonna ask just mechanically, similar to the ripple effect studies we did before, you know, if we have this shock now to this aggregate link capital, what I like to call connectivity capital, what is gonna be the aggregate effect of that? And also then, you know, firms, because firms are smart and care about this stuff, what are the externalities? And, you know, are there gonna be big inefficiencies when firms are deliberately investing in their multi sourcing and the robustness of their relationships. And so the main message is going to be that there's a big fragility where aggregate output can be extremely sensitive, much worse than we saw in the basic models in a kind of a way that can really destroy a lot of relational capital, which may then be permanently or at least, you know, for a long time lost. So I want to do this. I once had a chat about this paper with David Krebs, and he gave me a title but by then it was too late. But if I could call – if I could give it a name again, I might call it –  he suggested Livin' on the edge. It's a theory of self-organized criticality and supply networks. What we actually called the paper is supply network formation and fragility. And then, you know, I'm going to have to cut down on my discussion of context a little bit. But we wrote an annual reviews piece, which tries to summarize the broader literature. that inspired us to think about fragility. So I refer you to these for background. Okay, so let me tell you the key, let me make a key distinction that's gonna help you understand why this is different from the models we've been talking about. The first and second and third generation macro approach mostly focuses on a sectoral perspective. So we have a chip sector that supplies the cars and the car sector and also resins go into cars. And now Baqaee and Farki do have models that are big enough to think about individual firms having different relationships, but it's still certainly when these things are taken to the data, a lot of aggregation is typically done, right? But in reality, as I've already told you, a car manufacturer isn't sourcing chips from some open market for chips, they have specific suppliers. And so they're really reliant on these specific members of these sectors and relationships with them. So even if the sector is doing okay, it's really the disruption to specific links or specific nodes that you rely on that I should care about. And so that disaggregation is at the heart of this theory, taking Markus's exhortation to avoid the fallacy of composition seriously, right? We wanna study the microeconomics of this before we do the aggregation. And so we're gonna-

Markus Brunnermeier: If I rephrase this, is this a strong argument we're having some standardization? So you don't want these micro connections, but you standardize the industries that it has to satisfy a certain standard, so.

Ben Golub: Yeah, so jumping, yeah, absolutely. That's great, yeah. So jumping ahead to the, if you buy my claim that things are gonna be pretty bad in a world like this, you may say, we wanna make it easier to substitute, right? And I do think that's one of the interventions. It's not one that we actually wrote about much in the paper. But now that I think about it, it could be a very powerful one.

57:08
I guess there's a tension between that kind of thing and innovation, potentially. For example, Apple was able to extract a lot more rents by not standardizing a lot of things. And maybe that was important in how hard it was willing to innovate. But for these fragility issues, I do think standardization could be a huge help. And so let me now tell you. So I want to give you a flavor of the math. And I think we're going to be able to do this. So now this picture answers Markus's earlier question. Exactly how am I modeling this whole tree? I'm doing it in a very Mickey Mouse way, where I just assume this regular branching structure. And it's never going to come back to the same places. There's going to be lots of diversified suppliers. But the key structure that's repeated over and over again for many layers, I've just plotted three, but you could think of like five layers or ten layers, is that each firm needs multiple types of inputs, multiple types of goods that it needs to produce like brakes and computers. And each link for delivering this stuff is going to work with a certain probability, independently. So the independent part is idiosyncratic shocks, but if I vary X, that's like an aggregate shock to the fabric of the supply network. And it's important to say two things here. First, the arrows now point in the direction of orders being placed. So the downstream firms are ordering from upstream firms. You'll see why I chose that convention a bit later. And I'm doing percolation at the link level. You could also have firms disrupted. You could just have firms randomly not be able to work. They couldn't get a loan to operate or they were shut down for some exogenous reason. For shipping disruptions, we thought the link perspective made more sense, but you could do it both ways. And then once you get to the most upstream level, those firms aren't reliant on specific sourcing. They can just go by their inputs. This could be a level where things are standardized enough that you don't need specific inputs. And so the main mathematical thing that we're studying in this model is the reliability of production, which is just defined to be, for a given firm, reliability is defined to be the probability that in a random realization of these outcomes, if I choose a realization of how the links are on and off, what's the probability that this firm can produce? And I'll show you an illustration. Now, of course, reliability depends on depth, and that's why "d" here is an argument. If I looked at a 10-depth network, its reliability might be different from a two-depth. And we're gonna be interested in studying this reliability statistic, not just for a single firm, but on average, responds to shocks. So let me show you a picture of how this actually plays out in practice. What we do here is we take these links and so each of them works, let's say x is a half, each of them either works or doesn't work randomly. I realized it and the highlighted ones work and the gray ones, those shipments would not get delivered. There was port congestion and those shipments are late. And so now we can ask a simple question, can our little A1 producer here actually make his thing? Well, we're going to do it by starting from the top, sort of a backwards solution approach. So we can see immediately that this brake producer is not able to function because neither of these G inputs that he needs got delivered, right? So because he's missing a key input, he has the other input, but he doesn't have this input. He's going to turn red and be disabled. And now you can see that A1 here, this firm of interest, is dead because though this

shipping link would work, there's no brakes to put on the ship. So this doesn't work. And then this shipping link is just exogenously disabled by bad luck. And so we turn A1 red. And so that's how we work it out in a particular case. Now you could solve the little probability problem of working out how this would look on average.

1:01:30
And what we're going to see, I just want to give you a schematic. We're going to study how these reliabilities change with x. And we're going to see a striking phenomenon, which is familiar from physics that we got really excited about, which is in physics and now applied mathematics, we now know a lot about phase transitions, how fairly small changes in conditions like pressure and temperature can make a huge difference to how your system looks. Ice looks this very structured way and then water molecules look totally different, and when we cross certain boundaries in a parameter space, we can make these changes happen. So something similar is going to be true in the economics here, and let me just give you the flavor of it. So what we do is we look at the reliability of firms on average, averaged over a distribution. A distribution of what? A distribution of depths. So I told you about a depth two firm, we have a distribution of all sorts of different depths in the economy, and we're going to look at reliability on average. And mu is a distribution over numbers that are depths. We set up a standard love of variety model where the consumer is happier when more firms are able to function in a given time period. So this is just standard love of variety theory where the consumer likes more different types of goods, and the more supply networks are up and working, the more the consumer is able to eat. And so let me tell you the key force in the picture. What I'm plotting here is on the x-axis, I'm plotting this probability that links are working, what I call relationship strength. And on the y-axis, I'm plotting the reliability. For a particular distribution of depths, I think here I chose a distribution with a mean of about 12, so reasonably deep, not crazy. And you see that you have a very steep precipice, kind of, where up until a certain x, you're producing very rarely, and then suddenly you're producing a lot. And so what we do formally, mathematically, is we consider limits where we take the depth of the supply network to be large enough. What that means in practice isn't super crazy, it's numbers like 20. And we see that this picture goes to this, in the limit, it converges to this zero reliability until a certain critical relationship strength. And then, so I'll tell you the proposition quickly in a more precise way, until some critical x, the limit reliability is zero, right? And then once you cross the critical level of relationship strength, you suddenly have uptime that's way, way above one. So that's a phase transition, it's a discontinuous phase transition in the structure of this network. Now, I'm gonna go to the, just mentioned the key sort of punchline that comes when we bring the real economics to it. If you look at that picture, you know, I showed you there's a discontinuity. And if you... If you have this aggregate shock, things can be quite bad, but, you know, you might say, well, there's only one place, right? Let me, in fact, let me go back to the picture. So with complex production, you know, there's only one place here where you're on the cliff edge. And most of the time, either you're nice and safe or you're not producing anything, but maybe you're not so worried about that network.

Markus Brunnermeier: But can I make a comment? So given that the reaction is this phase transition, actually the incentives, because if I think of a world where X is chosen endogenously by different players, they might have an incentive to choose a large enough X. The incentives are very high to be beyond the critical level, I guess.

1:05:37
Ben Golub: Absolutely. Exactly. So, maybe firms invest enough to avoid them. And maybe the government, if it's smart, you know, encourages them a little bit and helps them. And so what we'll see, so what we do in the sort of main part of the paper is we build a theory where firms endogenously invest in their relationships. We model, you know, the idea here just for interpretation is that you invest in both logistics and multisourcing, and you may also work really hard to make your relational contracts deliver for you. That's a big theme in economic sociology, that firms really care about this stuff and spend real resources to firm up their relationships. So what we'll see, and this is the self-organized criticality, it turns out that not universally, but there's a powerful force where actually because of the endogenous investment, firms are in a sense more likely to end up on the precipice, and I'll explain why that's true. I just want to give a shout out here, Schenkman and Woodford have a beautiful kind of conceptual paper on self-organized criticality and kind of sandpile models, and in applied math people are very interested in these catastrophic discontinuous phase transitions. What we really do that's new in this paper with Elliot and LeDuc is we just bring basic economic investment modeling to that kind of situation exactly to address the question that Markus asked. You know, so I'll be quick here, but what we do is very simple. We let each firm choose its own X. So before we realize any of the randomness, the executives decide how much to invest in making relationships likely to function, which you can interpret either as finding suppliers or in firming up the relationships that you do have. And you choose the X with which your links work. That's why the arrows were pointing to, that's your choice, how strong those links are, the ones that you extend. And you pay costs, of course, for this. So there's a cost function satisfying fairly standard assumptions describing how much it costs you to invest. And if a firm produces, it earns gross profits where we make, you know, so the main thing that we're gonna be changing here is a number called kappa, which is the productivity of the economy. That's how much profit you make if you produce. That's like the profitability of the whole, of how good is the good that you're making, and how much can you earn from it. And then there's a function. It's like an aggregate TFP shifter for this sector. And then intuitively, we assume that you want to be more reliable when fewer of your competitors are, because then maybe you can charge some higher prices or you get more of the market. So it's another force for resilience, if you will. And what we find, and this is I'll just assert the result, we find that there's a whole continuum of equilibria on the precipice. That is, for a whole open set, a whole interval of these parameters, kappa, the equilibria are going to be found. So we're considering finite, but quite deep networks. So this vertical part is sort of like a very, very steep part. And the equilibria where the people endogenously choose the X, it's going to be somewhere very close to the critical X. That's the finding. And actually, this is the sort of paradox. With an exogenous kappa, it would be another... knife-edge outcome to end up there, but with, sorry, with an exogenous X, but with an endogenous X chosen by the firms, there's a whole open set, it's a non-generic outcome to end up there.

Markus Brunnermeier: So, just to make sure I understand, so there's a whole, but you're always at the critical level or just below or just above?

1:09:19

Ben Golub: You are, so yes, that's, I should have been clear about that. You are producing at a positive, the reliability is actually positive before the shock. So what I want you to think about is there's maybe a small probability shock that's coming and the equilibrium is going to be realized before that in anticipation of it. Where is it going to end up? It's going to end up somewhere here where reliability is actually positive, we're producing reasonably well, but if X is degraded and we don't have a chance to respond, our relationships get a little weak, like, you know, everybody's X has shifted a little bit left relative to what they expected. That's going to really bring down the network. So, what's the trick, what's the magic? The externality is that firms don't appropriate the full downstream benefits of their investment. And that's true in any standard trade or macro model that you write down. Firms when they invest more and improve their effective reliability TFP, it does help, it does improve things, but it improves things and they make some more profits, but they don't make the full social surplus. You have to write down extremely unrealistic, efficient market models to give firms a way to extract the rents when they are pivotal. And so you don't, that doesn't happen and so that's what, that's the externality that causes this.

Markus Brunnermeier: Just to make sure I understand. So you have to subsidize them, but only a little bit. Is this the conclusion?

Ben Golub: Ah, very good. So this, that's our next question. question, what if we subsidized? It turns out that if you subsidize their investment in resilience or robustness, the equilibrium forces will crowd it out basically one for one. So you have to subsidize kind of a lot to put them – now, it is possible, there are capas that are good enough that they're way above the precipice. So that's possible. But if they're over here and you subsidize, they basically will disinvest to bring you back. And so at least that's possible. So one of our messages, and I'll sort of, I'll wrap up, but one of our messages is that in models like this, resilience policy is quite subtle and you have to really think about the discontinuous aspects of this to do a good job with it. Okay. So I've tried to convey a flavor of propagation via matrices, contagion via this percolation type of theory. There's obviously, you know, there's a lot to learn, but you hopefully get a flavor at least of what the answers look like. And I think in both the propagation stuff, with eigenvalues and eigenvectors and the statistical study of them. And also in these theories of percolation, there's a ton of theoretical work to do, bringing in tools from our allied disciplines. And so, you know, I think I've already made these points. And I think the fallacy, the point that I want to end on is that if discrete firm to firm relationships are a big part of what's going on during crises, aggregating and abstracting away from that can really give you wrong answers in principle. And so when we look at people trying to measure reliance on China, what would happen if, you know, you had a degradation of commercial relationships with China, there's a good chance that aggregate style analysis, so like at Brookings, there was a paper last year on exactly that point, but very much in the old school, we want to have aggregated world. And we want to make the point that premature aggregation will lead to these fallacies of composition and potentially big mistakes. So I think figuring out a way to do that stuff well with the data that we actually have is a super interesting and exciting challenge and a paradise of good questions for ambitious researchers. So that's what I had to say.

1:13:13
Markus Brunnermeier: That's a fantastic ending. And we always end on a positive note. And I think that's probably the best positive note for all the researchers, young with a young heart. But let me ask you one final question. Think of if you were to add a generalist, like a jumper, you know, whenever something breaks down, you can move this person in and you can, would this be very complicated in your framework that you say, Oh, you've… and what's the extra value added of this generalist? How much would you be willing to pay to have such a…

Ben Golub: Yeah, the social, the social returns of that are actually immense. And and the exactly for I think the reason that you're intuiting that, you know, related to the crowding out point, if you x anti subsidize resilience, and you say, I want you guys to invest more, then you have to deal with the equilibrium crowding out that might happen, the strategic substitutes. But if instead you say, when shocks happen, they don't happen that often, I'm going to just bring in life support at that time, then that actually has hugely, you saw how steep the curve is, the social returns are huge. And because it's only happening in those rare events, it doesn't affect firms incentives that much actually. And so there's no crowding out. So a little bit, to us surprisingly, kind of reactive interventions, where maybe you send the military to unblock the ports where you relax regulations during those times, that can be much more effective theoretically than these more long run policies.

Markus Brunnermeier: So rather than subsidizing all members of the network to push their ex up, you would have some emergency firefighters or something ready. Exactly. Now what that means, what that means in the real economy, and by the way, you can imagine financial analogs of all of this, what that would mean in the real or financial economy is super, is not clear to me, but it seems like a good question. So coming back to your poll question, is it more the real economy, more the financial sector? Would you say the methods and the models you presented today, you focused very much on supply chains and input-output matrices. Would you have given the same talk if you would have used the financial sector, or would you say, oh no, this we have to take some additional perspective?

Ben Golub: Yeah, so I think what this is a good point to make a kind of theoretical networks point that what's really key about the complex, the what's called complex contagion with the discontinuous phase transitions, is the fact that it, unlike Covid, you can't decompose contagion in a simple kind of separable way across links. Like with Covid, you get it from one, as soon as you get it from one person, you got it, right? With the supply network things, the local production function has this feature that you need multiple things to go right. And in the financial setting, you may need multiple counterparties to make good on their obligations. You're fine if one of them defaults, and you're definitely not fine if all of them default, but one wouldn't be enough to make you solvent either. So it's somewhere in between there's this threshold. And it turns out that is the phenomenon that drives all the math. So what I think is interesting is in models like Acemoglu’s and Friend's famous debt clearing networks paper, you could build in this kind of you could do those kind of threshold phenomena in large networks.

1:16:30
And that just hasn't been done. I think actually that would be, that's like a free paper lying around that you could do the complex contagion theory with a debt motivation for how the local contagion works.

Markus Brunnermeier: OK, let's stop at this and leave the work for many PhD students to figure out a lot of dimensions and there's some low hanging fruit still there as well. Thank you so much, Markus. This was a lot. This was fantastic to get a great overview about the whole literature and.

Ben Golub: Thanks so much.

Markus Brunnermeier:A good starting point to do further research.