Hindawi
Complexity
Volume 2019, Article ID 8436505, 16 pages
https://doi.org/10.1155/2019/8436505

Research Article

Liquidity Hoarding in Financial Networks: The Role of Structural Uncertainty

Stojan Davidovic,[1] Amit Kothiyal,[1,2] Mirta Galesic,[1,3] Konstantinos Katsikopoulos,[1,4] and Nimalan Arinaminpathy[5]

[1]*Max Planck Institute for Human Development, Center for Adaptive Behavior and Cognition, Berlin, Germany*
[2]*Indian Institute of Management Udaipur, Udaipur, India*
[3]*Santa Fe Institute, Santa Fe, NM, USA*
[4]*Southampton Business School, Southampton, UK*
[5]*Faculty of Medicine, School of Public Health, Imperial College London, London, UK*

Correspondence should be addressed to Stojan Davidovic; stojand@mpib-berlin.mpg.de

Received 15 May 2018; Revised 21 August 2018; Accepted 4 September 2018; Published 8 January 2019

Guest Editor: Bernardo A. Furtado

The dynamics of confidence affect a plethora of financial phenomena including liquidity hoarding. We present a multiagent model of a financial network in which confidence dynamics are shaped by *structural uncertainty*—that is, the lack of knowledge about the network of interbank cross-exposures. During a financial crisis, structural uncertainty makes it difficult for banks to assess the risk of financial contagion and their own health. Under such conditions, banks are more likely to behave conservatively and quickly act on information they receive from their local environment. A sudden financial shock, therefore, can be characterized by high-intensity local impact on confidence. We find that such local impacts quickly spread throughout the network, causing more damage than a shock that evenly affects all localities in the system; for example, a complete breakdown of the system occurs with a higher probability. The results are explained analytically by linking system performance to the speed of decrease in confidence.

1. Introduction

The "freeze" of the interbank market in the recent financial crisis denied financial institutions (banks, for short) access to liquid assets when they needed them most. In late 2007, the U.S. and European markets experienced simultaneous runs on asset-backed commercial papers [1]. The second large shock occurred in the fall of 2008 when failures of AIG and Lehman Brothers set off defaults of money-market funds (e.g., the Reserve Primary Fund), which subsequently triggered runs on the repurchase agreement market. Banks responded with precautionary liquidity hoarding, causing interbank market to dry up [2, 3]. The resulting difficulty of borrowing money and the increase in interest rates produced a series of adverse consequences for financial and real sectors. First, pressure to obtain liquidity through sales of long-term assets led to fire sales that further deteriorated banks' asset positions. Second, facing the growing prospect of illiquidity, banks struggled to maintain their lending activities. Third, given that the price of money in interbank markets is a benchmark for the interest rates in the economy, the real sector experienced difficulty obtaining funding under reasonable conditions. Together, these factors further aggravated already existing symptoms of the crisis.

There are two common explanations for the interbank market collapse: the increase in counterparty risk and liquidity hoarding [4]. The uncertainty about the network of cross-exposures between banks, also known as *structural uncertainty* [5], however, made both of these factors more effective. The fear of counterparty risk—a risk that a business partner cannot meet its obligation—can be linked to the subprime market crash that led a large fraction of banks holding mortgage-backed securities to experience financial difficulties. The resulting increase in liquidity demand and

the reduction in the number of liquidity providers set off basic conditions for banks to withhold liquidity from the market—liquidity hoarding. Furthermore, structural uncertainty made it difficult for liquid banks to assess the risk of events farther in the network and identify risk-free parties. For the same reason, liquid banks could not rule out the possibility of suffering a sudden loss through already existing cross-exposures. In such a context, they were more likely to anticipate increased liquidity needs in the future, as well as the possibility of limited access to the interbank market. The loss of confidence in the interbank market and liquidity hoarding were hence the result of a number of interrelated factors, but it was the structural uncertainty that made each of them more consequential.

In this paper, we explored a model of a banking network in which the confidence of banks is shaped by uncertainty about the network of interbank lending. This is a realistic assumption since the network of interbank cross-exposures is not known due to the over-the-counter character of interbank transactions. We found that when confidence is sensitive to local information only, the probability of systemic failures increases substantially.

We adopted the framework of Arinaminpathy et al. ([6]; hereafter the AKM model) in which a bank's confidence is modeled to reflect the severity of the financial situation in the banking system. A bank's confidence is directly linked to its decision to roll over or withdraw previously established interbank loans—the lower the confidence, the higher the possibility of precautionary withdrawals. The model also includes the effects of fire sales and asset price contagion, as they are in a close relationship with market liquidity. Here, we assumed that banks are responsive to information received from their direct interbank counterparties but not from the banks that are further away in the network. The intuition is that in the face of structural uncertainty, banks rely on actions that take place in their locality. As was done in the AKM model, we also considered a case in which banks receive information from all banks in the system, but we made the information either noisy or delayed relative to the distance information needs to travel in the network to reach the receiver.

The remainder of this paper is organized as follows. We start with the discussion of related literature in Section 2, followed by the model of a banking network in Section 3. Sections 4 and 5 are devoted to our model of uncertainty and the application of initial shock(s) to the system, respectively. In Section 6, we detail our simulation procedure; our main results are presented in Section 7 and additional analysis in Section 8. We end with the discussion in Section 9.

2. Related Literature

Our paper is related to the network models of financial contagion [7]. Network approaches have proven useful for understanding contagion processes in biological, social, and financial systems (e.g., [8–12]). In financial systems in particular, individual institutions are linked to each other through a complex system of interbank lending [13] and holdings in common assets [14, 15]. Such a system lends itself naturally to being modeled with a network approach.

Most models of financial networks have treated contagion as being directly transmitted between institutions (e.g., [16–19]), leaving the mechanism of market panics largely unexplored. One reason may be because the outbreak of herding behavior is not well captured by a cascade that spreads through the network of cross-exposures. Instead, it is predominantly driven by a collective change in expectations, which does not exhibit simple cascade-like spreading patterns. A simultaneous drop in banks' confidence, for instance, can be a mediator of collective withdrawal of liquidity and fire sales. This is because the interbank market relies on the collective confidence in its service as a safe resort in case of unforeseen liquidity needs. Furthermore, the functioning interbank market attenuates banks' liquidity buffers by allowing them to operate with minimal holdings of low-profit liquid assets. Thus, the steady reduction in banks' liquid reserves in the decades prior to the financial crisis reflects the increasing market efficiency and the growing confidence in its reliability (Figure 1). However, the recent crisis demonstrated that such a scheme is not resistant to large financial shocks, which proved to be capable of undermining the collective confidence. Taken together, these all point to the importance of understanding of how confidence decay spreads in the banking system.

Hansen and Sargent [20] studied the sensitivity of beliefs to uncertainty, although they did not look at how such beliefs spread in the financial system. Gai et al. [21] introduced a network model of liquidity hoarding where the propensity for precautionary withdrawals, a proxy of the collective confidence, was exogenously decided. Such a setting left the process of confidence loss out of consideration. By contrast, in the AKM model [6], confidence was endogenously determined as a function of the severity of the financial situation in the interbank market but with the unrealistic assumption that banks have complete information about other banks in the system. This implied that confidence shocks were well distributed among all banks, leaving the impact of structural uncertainty and heterogeneously distributed confidence in the network unexplored.

Our paper is also related to the literature focused on the relationship between liquidity hoarding and asset prices (e.g., [4, 22–24]). For instance, Gale and Yorulmazer [4] and Diamond and Rajan [24] showed that in certain conditions, privately optimal decisions can lead to hoarding behavior and fire sales. These authors also considered speculative hoarding, when a liquidity shortage stimulates liquid buyers to withhold liquidity in expectation of high returns from potential fire sales. Nevertheless, the main difference from our work here is that their studies were not concerned with the complexity of interbank cross-exposures or accompanying structural uncertainty.

There have been several studies connecting different sources of uncertainty to market liquidity. Caballero and Krishnamurthy [25] argued that capital immobility and liquidity hoarding can be explained by the reactions of decision makers to the Knightian uncertainty embedded in

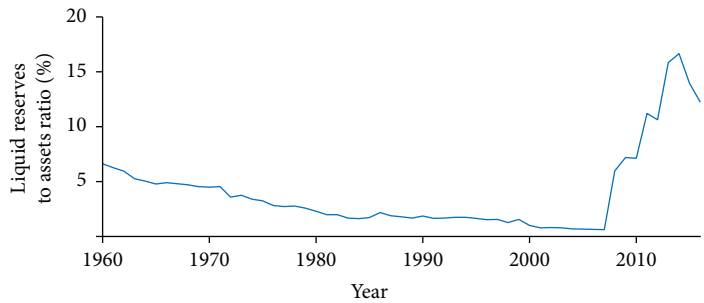

Figure 1: Ratio of bank liquid reserves to bank assets in the United States (1960–2016), that is, the ratio of domestic currency holdings and deposits with the monetary authorities to claims on other governments, nonfinancial public enterprises, the private sector, and other banking institutions. The recent 2007–2008 financial crisis led banks to hoard liquidity. Source: the World Bank, retrieved from http://WorldBank.com and http://NationMaster.com.

the financial environment in the form of sudden events and untested financial innovations. Routledge and Zin [26] showed how derivative pricing under uncertainty produces diverse effects on market liquidity. A link between uncertainty and market complexity was tackled by Zawadowski [27], who found that the layers of financial intermediation amplify uncertainty about the availability of funding, causing a cascade of liquidity withdrawals. However, all banks were familiar with the underlying network of lending.

In a recent paper addressing the same problem as our study, Caballero and Simsek [28] considered structural uncertainty in a model of liquidity hoarding and fire sales. They argued that ignorance about the underlying network of interconnections, a dormant factor in normal times, becomes relevant in a crisis when liquid banks have to consider who will be next affected by the cascade of failures. This shifts their preference toward keeping liquid assets instead of investing long term, causing liquidity shortages and fire sales. While sharing the assumption that banks reliably know only information about their counterparties, the Caballero and Simsek study differs from ours by assuming that banks know the outline of the network of cross-exposures. For simplicity, they also assumed that the banking network forms a circular shape. In our model, the network of interbank lending approximates features of real-world banking networks, and its shape is not known to the banks. Unlike the circular arrangement, the small-world feature of real financial networks makes all banks relatively close to each other [29], which has implications on the spread of information and financial risk in the system.

3. Model of a Banking Network

In the AKM model, banks are connected by lending relationships and holdings in common assets. Our analysis is focused on a short-term horizon in which banks can decide to roll over, shorten loan maturity, or terminate already established lending contracts. Their decisions depend on their confidence, which is expressed as a function of the level of assets and interbank loans in the system. When the system faces financial difficulties and bank defaults, the value of its assets and interbank loans shrinks, lowering confidence. This in turn leads to more preemptive actions of banks, putting pressure on their counterparties and causing more defaults. In this way, the model captures positive feedback between the severity of the financial condition in the system and individual behaviors of banks.

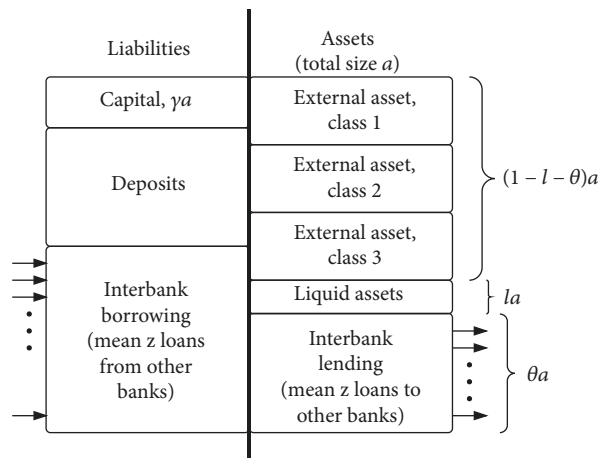

Figure 2: A balance sheet representation of a bank (adapted from [6]). a = total assets; γ = capital ratio; l = liquidity ratio; θ = interbank loans-to-assets ratio; z = average number of incoming and outgoing loans.

In addition to liquidity hoarding, two other contagion mechanisms take place in the model. One relates to the propagation of counterparty credit risk, which affects lenders when borrowers are not able to repay their loans. The other is asset price contagion, which occurs when liquidation of assets of failed banks pushes the corresponding asset prices down. All banks holding the affected assets suffer from the price drop, which is modeled according to Cifuentes et al. [30]. Correlations between assets are not included in our model.

3.1. Nodes and Edges. For simplicity, nodes or banks in the network can only be large or small, where the size disparity is fixed by the size ratio q(q = large bank assets/small bank assets). Banks are represented as simplified balance sheets with properties listed in Figure 2. The liability side contains capital (also known as owner's equity), retail deposits (money in the accounts of banks' customers), and interbank borrowing (assets borrowed from other banks). The level of capital represents the amount of asset loss that a bank can

withstand before becoming insolvent and going bankrupt (insolvency and illiquidity are separately discussed in Section 5). Retail deposits are taken to be external to the system and do not play an active role in the model. Interbank borrowing represents loans received from other banks, and the number of incoming loans represents the in-degree of an individual node. On the asset side, there are n external asset classes (investments in assets that are external to the banking system), liquid assets (e.g., cash), and interbank lending (assets lent to other banks). External asset classes are distributed among banks from a fixed number of distinct asset classes contained in the system (see next section). This means that multiple banks will share the same asset class, which can lead to asset price contagion—the drop in value of one asset class will affect multiple banks. Liquid assets are a small fraction l of the overall assets that banks keep in the most liquid form to meet immediate needs. They are mostly composed of cash or any cash equivalent, such as central bank reserves or high-quality government bonds, which are easily convertible to money. Finally, interbank lending corresponds to outgoing loans to other banks in the system, thus giving rise to a lending network, as described below. Parameters γ and θ (Figure 2) determine the initial proportions of capital and interbank loans in the total assets a, respectively. The balance sheet's parameters reflect the values observed in the banking sector before the crisis [31]: the proportion of total assets initially determined to be held in interbank loans $\theta = 0.2$, the proportion of total assets initially liquid $l = 0.01$, and capital-to-asset ratio $\gamma = 0.04$.

3.2. Network. The network is a directed random graph with $N = 120$ banks (Figure 3). The default value of the size ratio q is 10 (The main pattern of results is insensitive to changes of q and N as long as they are large enough.), which given total number of banks results in a network with $N_b = 11$ large and $N_s = 109$ small banks. The in-degree and out-degree of banks are determined by a Poisson distribution with parameter $z = 5$ for small and $q \times z = 50$ for large banks. That is, small (large) banks on average have five (50) incoming and five (50) outgoing loans. Each edge in the network is a loan with direction from lender to borrower. The default value of each single loan is normalized as 1. The maturity of interbank lending is also simplified: a random half of interbank loans are assigned to be "short-term" and the rest to be "long-term" loans. The short-term loans can be withdrawn immediately by a lender in a single decision, whereas long-term loans have to be shortened first. The banks are also interrelated by sharing the same external asset classes. These relationships are the basis for the asset price contagion. Small banks have 10, and large banks 20 external asset classes $(n_s = 10, n_b = 20)$. Given that on average 10 banks share the same asset class $(g = 10)$, this implies 131 distinctive external asset classes $(G = (N_b n_b + N_s n_s)/g)$.

The difference in the connectivity of large and small banks and the random assignment of their connections result in the core-periphery structure of the network. That is, large banks with many links are densely interconnected—forming the core, and small banks with few links are loosely interconnected—forming the periphery. The resulting structure is

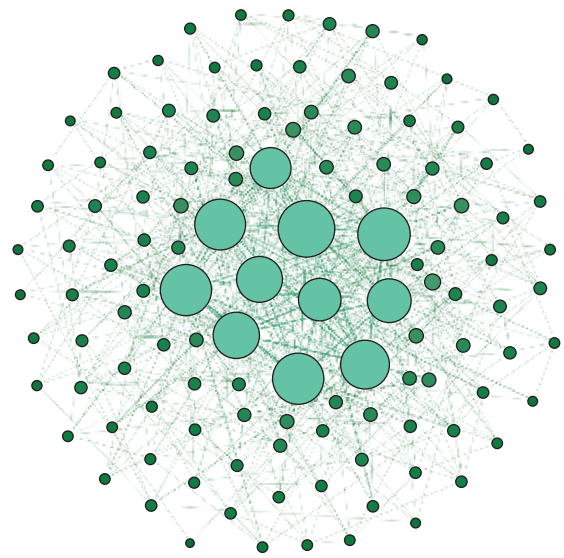

Figure 3: Banking network with 120 banks. The node size is weighted by its degree. Large banks are colored in light green, whereas small banks are colored in dark green. The network is constructed using the default values of the model parameters (Section 3.2) and has a core-periphery structure: large banks with many links are densely interconnected—forming the core, and small banks with few links are loosely interconnected—forming the periphery.

"shallow," (This is in agreement with the core-periphery structure of the global banking network [13]. Roughly speaking, any two banks at the periphery are a few connections away since they are linked via the well-connected core.) meaning that the average path length in the network is relatively short due to the well-connected core.

3.3. Confidence and Individual Health. Confidence C is the first important determinant of a bank's behavior. In the AKM model, confidence is calculated as a function of A and E, which are measures of solvency and liquidity of the system, respectively:

$$\begin{aligned} C &= AE, \\ A &= \sum_{i=1}^{N} A_i, \\ A_i &= \frac{a_i}{\sum_{i=1}^{N} a_i^0}, \\ E &= \sum_{i=1}^{N} E_i, \\ E_i &= \frac{e_i}{\sum_{i=1}^{N} e_i^0}. \end{aligned} \tag{1}$$

At a given point in time, A denotes the total value of all remaining assets in the system as a proportion of its initial value; E is similarly the fraction of interbank loans not withdrawn from the system; A_i and E_i are the remaining assets and interbank loans of bank i as the proportion of initial value of total assets in the system; a_i and e_i are the absolute values of remaining assets and interbank loans of bank i;

and a_i^0 and e_i^0 are the initial absolute values of assets and interbank loans of bank i.

To calculate C as defined in the AKM model, and to take any action, banks have to know the current and initial values of assets and interbank loans of all banks in the system. To explore how the network behaves in a more realistic setting, especially in times of crisis when the system is changing rapidly, we consider several uncertainty scenarios, described in Section 4.

Unlike C, which is a systemic parameter, h_i denotes the individual health of bank i and is calculated as a function of its indicators of solvency c_i and liquidity m_i:

$$h_i = c_i m_i, \quad 0 < h_i < 1,$$
$$m_i = \min\left[1, \frac{\left(A_i^{\mathrm{ST}} + l_i\right)}{L_i^{\mathrm{ST}}}\right], \tag{2}$$

where c_i is the capital of bank i defined as a proportion of its initial value; m_i is the fraction of i's short-term liabilities that the bank can settle immediately, through its liquid and short-term assets; A_i^{ST} is the total value of i's short-term interbank assets; L_i^{ST} is the total value of i's short-term interbank liabilities; and l_i is the amount of liquid assets held by bank i.

3.4. Decision Rules. The dynamics in the model are determined by decisions that banks make in discrete simulation time. For each of its outgoing connections, a loan provider can decide whether to shorten a long-term loan and whether to withdraw a short-term loan. A lender can withdraw short-term loans in a single decision, resulting in the connection between the two banks being removed immediately; long-term loans can only be "shortened" in a single decision, resulting in the connection between the banks becoming a short-term loan. Thus, an eventual withdrawal of a long-term loan requires an additional decision, i.e., time step (see Appendix B). Depending on its maturity, a loan between two banks i and j is, respectively, shortened and withdrawn when

$$h_i h_j < (1 - C),$$
$$h_i h_j < (1 - C)^2. \tag{3}$$

If C is high (1 or close to 1), these conditions are satisfied only under extreme conditions for h_i and h_j. In contrast, a drop in C can cause liquidity hoarding, as both decision conditions are more likely to be satisfied for all banks in the system. In addition, the shortening condition is easier to satisfy than the withdrawing condition, which means that banks resort to withdrawal only in relatively urgent situations.

4. Model of Uncertainty

In interbank markets, business partners trade privately, which often leads to a relationship with preferential treatment and repeated transactions [32, 33]. Accordingly, in our main uncertainty setting—the local information (LI) scenario—we restricted information availability to the nearby, that is, "local" banks in the banking network. This makes banks highly sensitive to local events and insensitive to events that take place further away in the network. Their responsiveness is calibrated so that they do not take action unless they notice signs of trouble in their locality; but when they do, then their reaction is intense. As a reference to models of complete information such as the AKM model, we also consider delayed information (DI) and noisy information (NI) scenarios, in which banks receive information from all banks in the system, but information is either delayed or noisy. For simplicity, details of the DI and NI scenarios are presented in Appendix A. As described in Section 3.3, the assumptions used in the AKM model are equivalent to what we call the complete information (CI) scenario.

To model uncertainty and determine the amount of information that is included in the calculation of confidence C, we rely on the distance between nodes in the network. The distance $d(i, j)$ is the shortest path length between information user i and information source j. If banks are directly connected, the distance between them is 1 and we call them neighbors. All neighbors of a particular bank constitute its neighborhood. The distance between neighbors of neighbors is 2 and so forth. The main principle for modeling uncertainty is that information availability and/or quality deteriorates when the distance from the information source increases. Once uncertainty is introduced, instead of one common estimate of confidence for all banks ($\forall i : C^i = C$ in the AKM model), each bank has its own individual perception of confidence C^i.

We use the following notation template of any model parameter $P : P^{\text{time step (optional)};\text{observer}}_{\text{observed (optional)}}$. For example, a_j^{0i} denotes bank i's judgment of j's initial (0 time step) absolute value of assets. Absence of the time step indicator implies the current value of a parameter. The indicator of an observed bank is omitted in the case of aggregate parameters, such as C, which are not based on information of an individual bank.

In the LI scenario, information is available only up to a certain "interbank" distance. That is, bank i calculates C based on the information about itself and all banks placed within the fixed value of distance $d_{\max}$. This is our general definition of locality where the value of $d_{\max}$ determines whether the locality includes only neighbors or also neighbors of neighbors and so forth. For instance, if $d_{\max} = 1$, then only i and its immediate neighbors contribute information to C. Now we can express the confidence of bank i as

$$\begin{aligned}
C^i &= A^i E^i, \\
A^i &= \frac{a_i + \sum_{j \in J_i(d_{\max})} a_j^i}{a^{0i}}, \\
a^{0i} &= a_i^0 + \sum_{j \in J_i(d_{\max})} a_j^{0i}, \\
E^i &= \frac{e_i + \sum_{j \in J_i(d_{\max})} e_j^i}{e^{0i}}, \\
e^{0i} &= e_i^0 + \sum_{j \in J_i(d_{\max})} e_j^{0i}.
\end{aligned} \tag{4}$$

A set $J_i(d_{\max})$ contains all banks that i considers for estimation of C, except for i itself, and is a function of $d_{\max}$. It is useful to think of $d_{\max}$ as a parameter that determines the reach of i's perception. To define $J_i(d_{\max})$, we first define the set $J = 1, 2, \ldots, N$, which contains all banks in the network. Then, its subset $J_i(d_{\max})$ is defined as

$$J_i(d_{\max}) = \{j \in J | d(i,j) \leq d_{\max} \& j \neq i\}. \quad (5)$$

We consider two versions of the LI scenario (Figure 4): LI1 in which $d_{\max} = 1$ and LI2 in which $d_{\max} = 2$. Since the network is quite shallow (average path length is barely above 2), LI2 contains almost the full graph, and LI3 is equal to CI. Thus, LI2 will provide a useful sanity check in respect to CI.

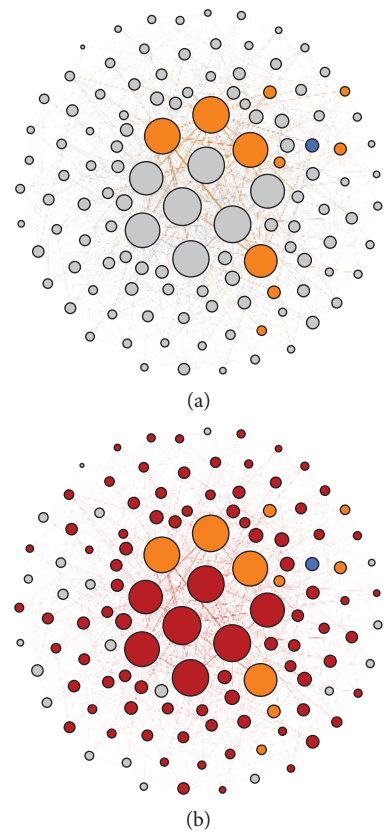

Figure 4: Information availability in the banking network of 120 banks in two local information (LI) scenarios: LI1 and LI2. In the LI1 scenario (a), information is available up to distance 1 in the network. That is, the bank in blue receives information only from the orange banks—its immediate neighbors. In the LI2 scenario (b), information is available up to distance 2. That is, the bank in blue receives information not only from the orange banks but also from the red banks—its neighbors of neighbors. The remaining banks are colored in grey. The node size is weighted by its degree. The network is constructed using the default values of model parameters (Section 3.2).

5. Model of Shocks and Bank Failures

Under each of the conditions described above, we simulated the response of the system to an initial shock. We explored two types of initial shock: (i) a *concentrated* shock (or a single-bank shock), applied by randomly selecting a large or a small bank and forcing it to fail by setting its capital to zero, and (ii) a *distributed shock* (or a multiple-bank shock), applied by forcing multiple small banks to fail simultaneously. The multiple-bank shock is designed to involve a number of small banks whose aggregate assets are equivalent to the assets of a large bank. Therefore, comparing these two treatments can be informative about how the system responds when the same shock is concentrated in a single bank or distributed among multiple banks.

A bank can go bankrupt for both liquidity and solvency reasons. A bank is *illiquid* if its liquid assets and interbank loans are insufficient to meet the demand of other banks to repay the loans previously taken from them. A bank is *insolvent* once the asset devaluation (from an external asset price decrease or counterparty default, for instance) exceeds its level of capital.

6. Simulation

In our simulation, each replication is a computational experiment with two phases. The first phase is to form the network and apply the initiating shock. The second phase is to simulate the propagation of the shock through the network, which unfolds in several iterations, here called time steps.

To form the network, in- and out-degrees that determine the numbers of banks' incoming and outgoing links were drawn from a Poisson distribution. (To design a network, we first drew out-degrees from a Poisson distribution and used this draw as weights for random sampling of corresponding in-degrees. If both in- and out-degrees were drawn directly from a Poisson distribution, the procedure would require the random draw to be repeated until the sum of all in-degrees is equal to the sum of corresponding out-degrees. As a result, draws with nonmatching degrees would have to be discarded, which is computationally expensive and problematic for the purpose of the analytical analysis.) We used a zero-truncated version of a Poisson distribution to ensure positive values of interbank assets and liabilities, which provided more balanced initial liquidity of banks. Once in- and out-degrees were determined, it was possible to reconstruct the rest of the bank's balance sheets based on the parameters of the model (see Section 3.2).

After the network was formed, a shock was applied. The shock hit one or several randomly chosen banks, depending on the type of shock to be applied. Then, the remaining simulation procedure entailed iteration of actions that take place in discrete time (see Appendix B).

7. Results

The results are based on 1,000 simulation replications per scenario, with each replication lasting until the system came to rest. Figure 5 depicts the probability distribution for the total number of failed banks after a shock is applied to a single small, a single large, and multiple small banks. In the

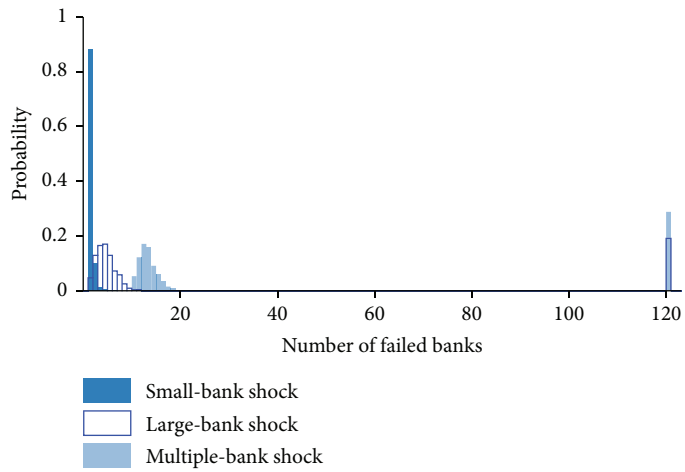

Figure 5: Probability distribution of number of failed banks in the complete information (CI) scenario after a shock is applied to a small bank, a large bank, and multiple small banks. The systemic breakdown occurs with sizable probability of nearly 20% and 30% after a large-bank shock and multiple-bank shock, respectively.

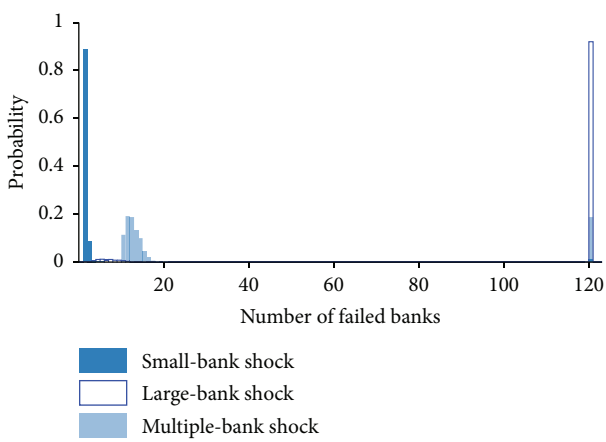

Figure 6: Probability distribution of number of failed banks in the LI1 scenario after a shock is applied to a small bank, a large bank, and multiple small banks. In comparison to the CI scenario, the probability of systemic breakdown increases after both single-bank-shock treatments, whereas it drops after multiple-bank shock. LI1 is a scenario in which a bank has access only to information from its direct neighbors at distance 1. LI = local information; CI = complete information.

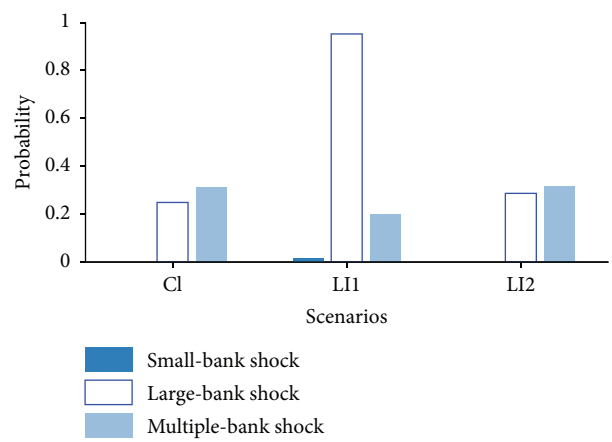

Figure 7: Probability of whole-system breakdown in complete information (CI) and local information (LI) scenarios. The LI1 scenario obtains the highest probability of systemic breakdown in both single-bank shock treatments. LI1 and LI2 = scenarios in which banks have access only to information from banks up to distance 1 and 2, respectively (Figure 15 in Appendix A displays all scenarios).

cases of a large-bank shock and a multiple-bank shock, the fat tail of the distribution indicates that the entire system collapses with a probability of nearly 20% and 30%, respectively.

When uncertainty is introduced, the highest impact on the probability distributions of number of failed banks is realized in the LI1 scenario. Figure 6 shows that the probability of systemic breakdown (i.e., all banks in the system fail) after a large-bank shock is now more than 90% and that even a small-bank shock results in a nontrivial probability of systemic breakdown (In high-resolution data of 10,000 repetitions, the probability of whole-system breakdown after a small-bank shock increases from 0% in the CI scenario to 0.16% in the LI1 scenario. This is easier to see in Figure 7.), whereas the same probability drops to under 20% in the case of a multiple-bank shock.

Figure 7 displays a comparison of probabilities of systemic breakdown across different scenarios in all three shock treatments. The probabilities are consistently higher in the multiple than in the single-shock treatment, except for the LI1 scenario. In fact, the LI1 scenario, which is associated with the largest probability of systemic collapse in the case of a large bank's shock, is at the same time associated with the smallest probability of systemic collapse in the distributed shock condition.

The results of the LI1 scenario are particularly striking as they show that a limited information flow further intensifies the contagion dynamics observed in the AKM model after a

large-bank failure. At the same time, the LI1 scenario mitigates the impact of a multiple-bank shock when compared to the CI scenario, in which this treatment in fact yields the highest probability of systemic failure (Figure 7). While this illustrates that the LI1 scenario does not merely amplify the contagion dynamics, it also shows how an alternative assumption about information availability can flip the conclusion about which event is more likely to trigger catastrophic failures.

In the following section, we mainly focus on explaining the difference between the results obtained in the CI and LI1 scenarios. Given the insignificant change in results produced in the other manipulations (NI and DI scenarios; Appendix A), we discuss those results very briefly.

8. Explaining the Results

In the CI scenario, confidence is assessed over the extent of the whole system: although this captures the notion of a generalized psychological context, it also has the effect of diluting the local impact of a shock. In the LI1 scenario, by contrast, we have introduced the notion of "locally perceived" confidence that can vary with the neighborhood of different banks. The local impact of an initiating shock is therefore more intense than in a CI scenario but limited to the neighborhood, leaving the confidence of the remaining system initially intact. Yet, this local impact is subsequently transmitted through the system (analogous to the dynamics of crack propagation in a solid medium), resulting overall in a higher risk of system collapse than in the CI scenario. The similarity of the results of the LI2 and CI scenarios (Figure 7) provides a useful sanity check, as the portion of the system taken into account for the confidence estimation is minimally different in the two scenarios (Figure 4).

A similar rationale applies to the results of the distributed shock treatment, which involves a failure of multiple small banks. While the impacts of the small-bank failures on confidence "add up" in the CI scenario, irrespective of their placement, in the LI1 scenario, they independently harm confidences of the disparate localities in which they randomly fall. As a result, the probability of whole-system failure after a distributed shock in LI1 is noticeably reduced when compared with the CI scenario (Figure 7). That the "adding-up effect" is less prominent in the LI1 scenario can also be seen by contrasting the results obtained from CI and LI1 after small idiosyncratic and multiple-bank treatments. For this purpose, it is useful to interpret a multiple shock as adding extra instances of small shocks to a small shock. The resulting pattern is somewhat counterintuitive. While a small-bank shock alone leads to a higher probability of systemic failure in the LI1 scenario (than in the CI scenario), after multiple shocks, the system fails with a higher probability in the CI scenario.

To better understand how the dynamics of C affect the discrepancy in the results between the CI and LI1 scenarios, we conducted a further analysis to assess the portion of the system that is initially affected by the shock applied in LI1 scenario, the magnitude of C drop that corresponds to the applied shock, and the sensitivity of the system to different manners in which C can deteriorate. Finally, we compared the time course of C obtained in computational simulations in the two scenarios and designed tests to assess if the observed difference can explain the results.

8.1. Dynamics of Local Confidence. The only distinction between the CI and LI1 scenarios corresponds to the difference in confidence contexts in which banks' decisions are made. Given the complexity of confidence dynamics, we aimed to compare the two scenarios in terms of initial confidence effects caused by the applied shock. Then, we analyzed if the initial difference could account for the results. We focused on the impact of the large-bank shock, which is associated with the most striking contrast between the two scenarios, but the same rationale applies to the small-bank shock. In the CI scenario, to calculate C^1 (the level of C in the immediate aftermath of the shock), one needs to know only the number of small N_s and large N_b banks as well as the size ratio q. Then, from $C = AE$, it follows that

$$C^1 = \left(\frac{qN_s + N_s - q}{qN_b + N_s}\right)^2. \tag{6}$$

For the default values of the model parameters, $C^1 \approx 0.92$, which corresponds to the drop of C for approximately 0.08. Unlike CI scenario where confidence effects are uniformly distributed among all banks, in the LI1 scenario, the shock initially affects only the confidence of banks in the neighborhood of the shock. The assessment of the fraction of banks affected by the shock, therefore, requires the estimation of the size of "average-bank neighborhood"—that is, the expected number of unique large and small banks that are connected to a given bank, including only its borrowers and lenders (Figure 8).

For default values of model parameters, a large bank is on average connected to 40 small and 10 large banks, whereas a small bank is connected to 4 large and 5 small banks (for the proof see Appendix C). This implies that the large-bank shock on average affects approximately three-quarters of the system assets; and among affected banks, large banks experience a decline of confidence to $C_b^1 \approx 0.87$ and small banks to $C_s^1 \approx 0.61$. The faster decline of confidence of small banks is both intuitive because the average-bank neighborhood of a small bank is relatively smaller, and realistic, because it is to be expected that a smaller bank suffers larger confidence loss when faced with a shock of a given size. In what follows, we describe a test designed for the assessment of system sensitivity to the steepness of C decline.

8.2. Test 1—System Sensitivity to the Loss of Confidence. How does the system behave under different C regimes? To investigate this, we externally enforced different time courses of C while keeping the assumption that all banks share the same C. (This means that C is no longer endogenously determined from the fluctuation of asset levels.) The independent manipulation was designed to test the resistance of the system to a variety of hypothetical confidence contexts. The goal was to explore how the system responds if only the perception of decision makers is manipulated, while keeping all remaining

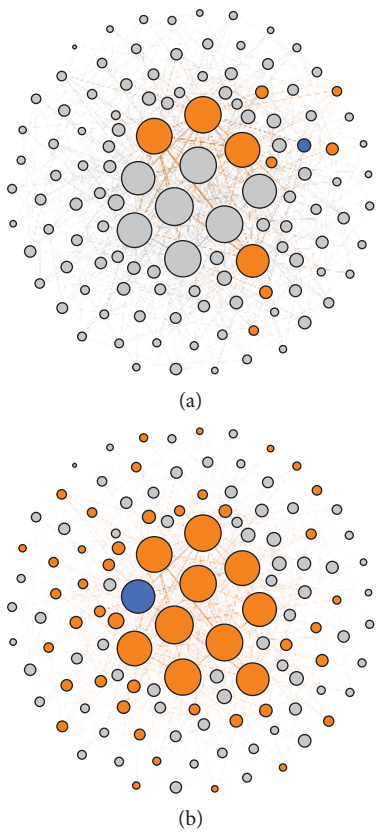

Figure 8: Neighborhood of a small bank (a) and a large bank (b) in the banking network of 120 banks. The small (large) bank is colored in blue, whereas its immediate neighbors are colored in orange. The remaining banks are colored in grey. The node size is weighted by its degree. The network is constructed using the default values of model parameters (Section 3.2).

processes endogenous. We considered different magnitudes of C drop and for each of them we also manipulated the slope of decline. We used the exponential function $f(x) = e^{-rx}$ to model the slope manipulation, where the value of parameter r determines the slope (Figure 9).

Figure 10 shows that the sudden drop of C ($r = 100$) has by far the highest impact on the system. This is a strong indication that the quicker loss of confidence in the LI1 scenario as compared to the CI scenario played a role in the increase in systemic risk. In addition, we compared standard deviations of confidence across the scenarios (Figure 11). The standard deviations of the end-state confidence (when the system is at rest) were calculated across 1000 simulation replications, taking into account only the surviving population of banks. Two results stood out. First, standard deviations of confidence were consistently higher after the large concentrated shock than after the distributed shock. The immediate implication is that the outcome of a large concentrated shock is less predictable. Second, there was a large difference between the standard deviations of confidence in the CI and LI1 scenarios. To determine if this contributed to the difference in the corresponding results, we carried out an analysis of the variance of confidence, described next.

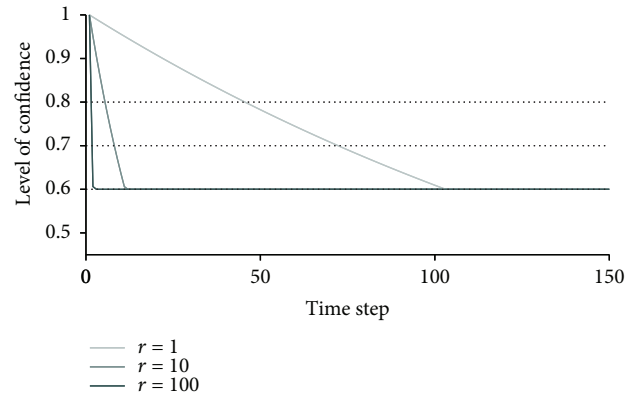

Figure 9: Manipulation of confidence (C) change over time: C drops 20, 30, or 40%, with the slope of decline determined by the parameter r of the exponential function $f(x) = e^{-rx}$; the higher r, the steeper the slope of C decline. Time step refers to discrete simulation time where each step iterates a defined set of actions (see Appendix B). Note: the observed time horizon is extended from 100 to 150 time steps, as additional time steps were needed for the system to reach the steady state when the slope of C decline was low ($r = 1$).

8.3. Test 2—Variance of Confidence. The goal of this test was to assess the sensitivity of total assets A and total interbank loans E to the change in variance of C. A realization of C in a simulation replication is in fact a vector $C(t)$, which contains values of C at different time steps. The manipulation first entailed construction of two vectors $C(t)_{i_{\mathrm{CI}}}$ and $C(t)_{i_{\mathrm{LI1}}}$ based on data from realizations of C in the CI and LI1 scenarios when a large-bank shock is applied. Two newly composed time sequences of C values were generated from a normal distribution with the same mean and two variances: $C(t)_{i_{\mathrm{CI}}} \sim N(C_i^{\mathrm{AV}}, V_i^{\mathrm{CI}})$ and $C(t)_{i_{\mathrm{LI1}}} \sim N(C_i^{\mathrm{AV}}, V_i^{\mathrm{LI1}})$. The mean C_i^{AV} was estimated by averaging the confidence from the realization of the CI scenario over simulation repetitions. The first variance, V_i^{CI}, was calculated from vectors of global confidence realized in the CI scenario and the second, V_i^{LI1}, from vectors of local confidence realized in the LI1 scenario. Finally, the two vectors $C(t)_{i_{\mathrm{CI}}}$ and $C(t)_{i_{\mathrm{LI1}}}$ were exogenously applied to the CI setting of the simulation (Figure 12). The exogenous application of confidence implies that the calculation of confidence is decoupled from assets and interbank loans in the actual simulation and taken as given. The sequences of realized networks were controlled to be the same in both conditions by setting the same seeding of the random number generator in the simulation $C(t)_{i_{\mathrm{CI}}} \sim N(C_i^{\mathrm{AV}}, V_i^{\mathrm{CI}})$.

Even when the mean of C over simulation replications is kept constant, as Figure 12 illustrates, a higher variance of C yields a faster drop of total assets A. Given that assets determine the level of C by definition, we designed an additional test to assess the impact of the time course of C on the results.

8.4. Test 3—Time Course of Confidence. For this purpose, the mean of individual confidences of all banks in the LI1

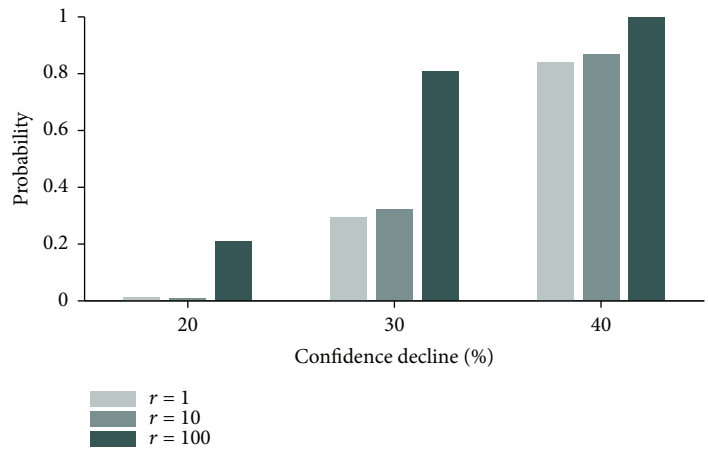

Figure 10: Probability of systemic breakdown for various magnitudes and speeds of confidence (C) decline. The magnitude of decline refers to the percentage of the initial value of C that is lost. The speed of decline is determined by the parameter r of the exponential function $f(x) = e^{-rx}$; the greater the value of r, the greater the speed of C decline. The sudden decline of C ($r = 100$) is associated with a sharp increase in the probability of systemic breakdown, especially in the cases of 20% and 30% of C decline.

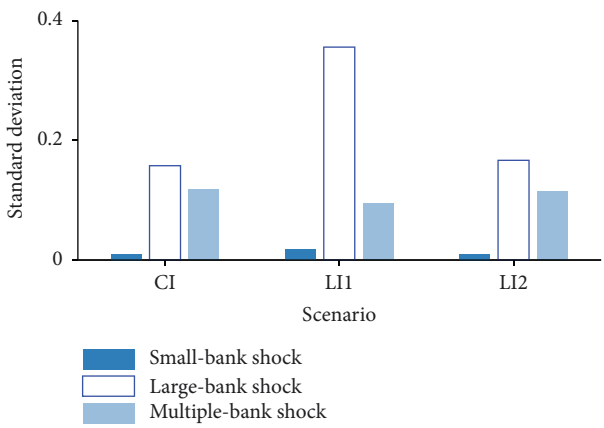

Figure 11: Standard deviation of confidence across complete information (CI) and local information (LI) scenarios and three shock treatments: to a small bank, a large bank, and multiple small banks. The LI1 scenario inflates standard deviations of confidence in both single-bank-shock treatments. The multiple-bank treatment involves a number of small banks whose total assets amount to the assets of a single large bank. LI1 and LI2 = scenarios in which a bank has access only to information from banks up to distance 1 and 2, respectively (Figure 16 in Appendix A displays all scenarios).

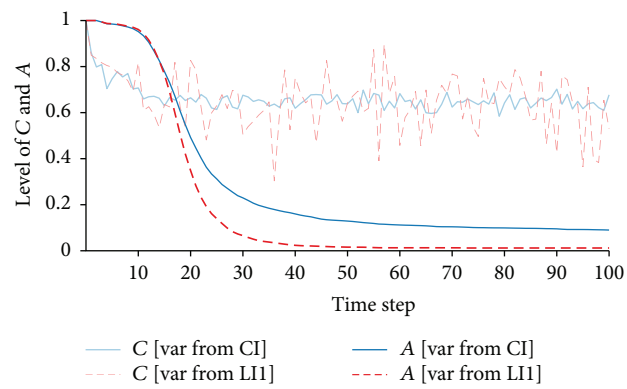

Figure 12: The impact of exogenous manipulation of confidence (C) on the level of total assets (A) in the system. Two applied manipulations of C follow normal distribution with the same mean and different variances: V_i^{CI} [derived from the complete information (CI) scenario] and V_i^{LI1} [derived from the local information (LI) LI1 scenario]. The higher variance corresponds to the faster decline of total assets in the system. CI is a scenario in which a bank has access to information from all other banks in the network. LI1 is a scenario in which a bank has access only to information from banks at distance 1.

scenario was calculated and denoted as local confidence. Confidence calculated according to the standard procedure, as in the CI scenario, was denoted global confidence. Figure 13 indicates a steeper decline of local as compared to global confidence when corresponding simulations were performed in an identical simulation setting, that is, when the identically placed large-bank shock was applied to an identical set of networks by controlling the seeding of the random number generator in the simulation.

In the next step, we estimated the impact of the observed slope difference between the two C curves by exogenous application of the local confidence to a hypothetical CI scenario together with the large-bank shock treatment. In the hypothetical scenario, as in the standard CI scenario, all banks in the system perceive confidence equally, but their perception is no longer endogenously determined. Instead, we forced their global confidence to be equal to previously determined local confidence taken from the realization of the LI1 scenario depicted in Figure 13. This procedure yielded a probability of over 90% of the whole system failing, a result similar to that in the LI1 scenario (Figure 14). The decline of confidence is therefore capable of explaining the difference in the results between the scenarios.

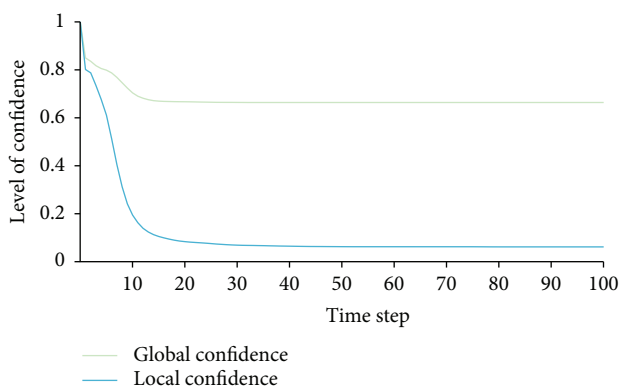

Figure 13: Time course of global and local confidence. Global confidence is calculated in a standard way as in the complete information (CI) scenario. Local confidence is an average of individual confidences of all banks in the local information (LI) LI1 scenario. CI is a scenario in which a bank has access to information from all other banks in the network. LI1 is a scenario in which a bank has access only to information from banks at distance 1.

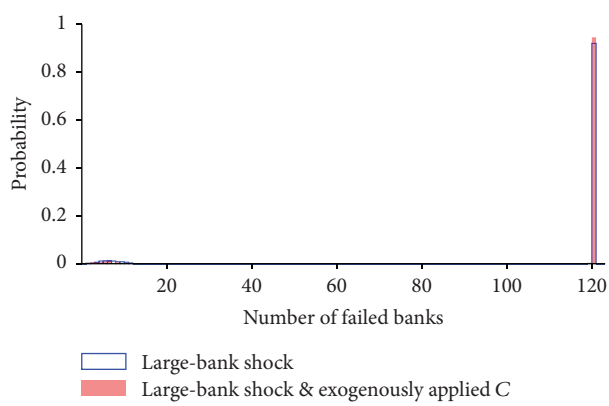

Figure 14: A comparison of probability distributions of number of failed banks after a large-bank shock in the LI1 scenario, and scenario in which confidence from the LI1 scenario was exogenously applied to the CI scenario. The probability of systemic breakdown is nearly the same in both scenarios. CI is a scenario in which a bank has access to information from all other banks in the network. LI1 is a scenario in which a bank has access only to information from its direct neighbors at distance 1.

9. Discussion

The increasing interconnectedness of the global financial network has introduced an enormous amount of structural uncertainty in the financial system. Yet, its implications went unnoticed until the recent financial crisis when the interdependencies in the network made it nearly impossible to disentangle low- from high-risk investments and partnerships. In this paper, we presented a model where structural uncertainty shapes patterns of confidence loss in a financial network. We found that a severe confidence loss in a limited part of the network carries a higher systemic risk than a moderate confidence loss in the entire network.

Our use of a multiagent model allowed us to analyze more specific aspects of confidence decline and to consider different types of shocks that can affect the system. We found that while the magnitude of decline was important, as was expected, a sudden decline of confidence was also a major factor in the increase in systemic risk; when the average level of confidence was controlled for, higher variance of confidence corresponded to higher systemic risk. We also found that under uncertainty, a shock affecting one large bank was by far more impactful than the same shock distributed among multiple smaller banks; when complete information was assumed, the opposite was true. This suggests that it is the failure of a large bank that poses a major threat to the system.

One caveat of our work is that there is a lack of understanding of how banks react to structural uncertainty during a financial crisis. Here we assumed that in a crisis, when it is infeasible to assess relevant risks, banks emphasize the information that they observe in their local environment. An alternative view explored in the literature assumes that banks maintain the practices developed prior to the crisis and rely on their own assessment of relevant risks (e.g., [28]). However, such an assessment would require the knowledge of the underlying structure of the financial network, which is exactly unavailable under structural uncertainty. Moreover, a possible extension of our work would be to explore how structural uncertainly interacts with other sources of uncertainty, such as uncertainty about the value of assets [34], in hoarding behavior.

From the view of policymaking, our observed sensitivity of the system to sudden loss of confidence suggests that interventions, such as bailouts of distressed banks or liquidity injections, should be done without delay. Our results also confirm previous findings that large banks carry disproportionate amount of systemic risk and hence are more likely to require stricter regulations than what was previously assumed. More generally, our results indicate the need for regulation designed to improve overall transparency in the financial system. For instance, policies that aim at reducing and eventually eliminating over-the-counter markets or constraining the complexity of financial contracts would be highly desirable. Only a transparent financial system would allow banks to make sound decisions from the perspective of systemic risk.

Appendix

A. Delayed Information and Noisy Information Scenarios

Unlike in the LI scenarios, in the DI scenarios banks receive information from all other banks in the system ($d_{\max}$ is no longer exogenously set), but some of the information is outdated. We modeled information delay as a function of distance—the further the information source, the longer the delay. If k denotes the time step when information originated,

Table 1: The size of delay in time steps assigned to banks at different distance in different scenarios.

Scenario	Size of delay 0	1	2
DI1	i + neighbors	All remaining banks	
DI2	i + neighbors		All remaining banks
DI3	i	All remaining banks	
DI4	i		All remaining banks

Note: DI = Delayed information; i = information user.

t the time step in which it is received, d_s the distance at which delay starts, and s the size of applied delay, then

$$a_i^j = a_i^{kj}, \quad e_i^j = e_i^{kj}$$
$$k = \begin{cases} t & \text{if } d < d_s \\ \max(0, t-s) & \text{if } d < d_s \end{cases}, \quad d(i,j) \in \{1, 2 \dots d_{\max}\}, \quad d_s \in \{1, 2\}, \quad s \in \{1, 2\}. \tag{A.1}$$

We designed four variants of the DI scenario by manipulating s and d_s (Table 1). For instance, in the DI1 and DI3 scenarios, the size of the delay is 1 time step ($s = 1$), and in the DI2 and DI4 it is 2 time steps ($s = 2$). In the DI1 and DI2 scenarios, delay starts from neighbors of neighbors ($d_s = 2$), whereas in the DI3 and DI4 scenarios, it starts immediately from neighbors ($d_s = 1$). We set the minimum value of k to 0 since negative values of time do not make sense in this context.

In the NI scenario, noise in information increases with distance. If $\in$ denotes a random error with normal distribution $\in \sim N(0, \sigma^2)$, v the size of variance in the noise term, and $d_{\max}$ maximal distance in the network, then

$$a_j^i = a_j^j + d \in, \quad d(i,j) = 1, 2, \dots, d_{\max}, \quad \sigma^2 = v a_j^j$$
$$e_j^i = e_j^j + d \in, \quad d(i,j) = 1, 2, \dots, d_{\max}, \quad \sigma^2 = v e_j^j \tag{A.2}$$

We considered two variants of the NI scenario: NI5 and NI30. In the former, $v = 5\%$ and in the latter, $v = 30\%$.

Regardless of the amount of noise, the NI scenarios yield similar results to those of the CI scenario. In the case of the DI scenarios, although the impact is very small the probabilities of systemic failure increase with delay. This is particularly noticeable in the nonzero probabilities of systemic failure after a small-bank shock. In the NI scenarios, normally distributed noise averaged out across banks, producing no difference in results compared to the CI scenario (Figures 15 and 16). Assuming an alternative distribution of noise would potentially produce more interesting results. On the other hand, the DI scenarios indicate that the delay matters. The result can be accounted for as the effect of overconfidence. Namely, in the DI scenarios, confidence at a particular moment in time was higher than what actual information would imply. This narrows the time window for the preemptive action that would enable shortening of long-term loans, which otherwise could not be used to meet the upcoming liquidity needs.

B. Simulation Procedure over Time Steps

After the application of the shock, the simulation procedure entailed five actions taking place in each time step:

(1) Recalculate health h_i of all banks. The health is used for stipulating liquidation of banks. Zero health implies that a bank needs to be liquidated.

(2) Liquidate banks that failed in the previous time step (or those that failed because of the initial shock). If bank i is to be liquidated then the procedure is as follows:

(a) Withdraw all short-term loans $A^{ST}{}_i$ that can be collected from the borrowers of i. Triggering the collection procedure means that i's borrowers will ask their own borrowers for money, and so forth. Record banks that consequently satisfy the condition of illiquidity and are to be liquidated in the next time step.

(b) Settle all short-term borrowings $L^{ST}{}_i$ of bank i that can be paid from its initial liquid assets l_i and collected short-term loans $A^{ST}{}_i$. Record the resulting shortage or surplus.

(c) Calculate the total long-term assets of i by adding long-term loans to the capital c_i. To this sum add the result from substep b. If there is a shortage of assets when the sum is compared to

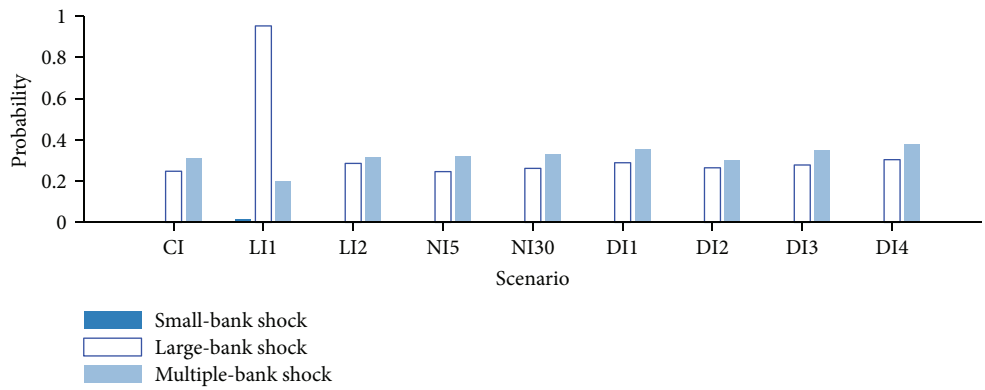

Figure 15: Probability of systemic failure across all scenarios. CI = Complete information; LI = local information; NI = noisy information; DI = delayed information; LI1 and LI2 = scenarios in which banks have access only to information at distance 1 and 2, respectively; NI5 and NI30 = scenarios in which noise parameter v is 5% and 30%, respectively; the delay scenarios are defined in Table 1.

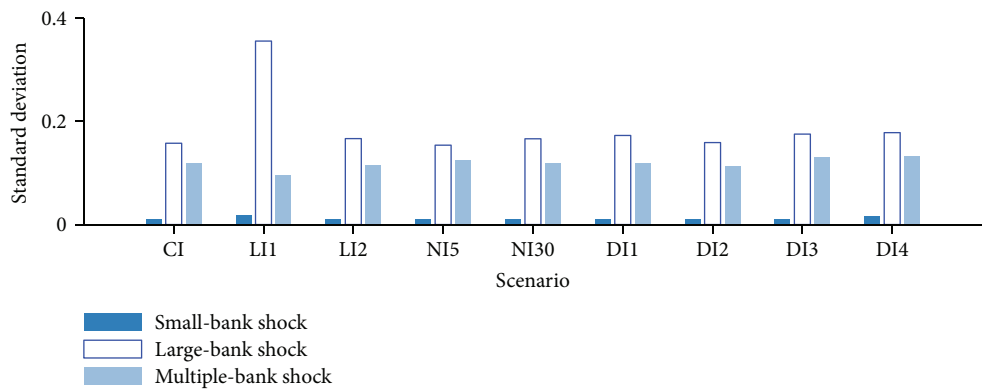

Figure 16: Standard deviations of confidence across all scenarios and three shock treatments: to a small bank, a large bank, and multiple banks. CI = Complete information; LI = local information; NI = noisy information; DI = delayed information; LI1 and LI2 = scenarios in which a bank has access only to information at distance 1 and 2, respectively. NI5 and NI30 = scenarios in which noise parameter v is 5% and 30%, respectively. The delay scenarios are defined in Table 1.

long-term liabilities of i, then i's long-term lenders suffer from this amount of shock u applied to their capital. The shock is evenly distributed among the lenders, but only up to the level of individual exposures. This ensures that the shock cannot exceed the level of individual lending amount.

(d) Sell external assets of i, applying the shock to all holders of the same asset classes that i had in its portfolio. The external assets are sold at a market that is taken to be external to the model. The price of asset w is assumed to be decreasing to a fraction $exp(-\alpha x_w)$ of its initial value (modeled as in Arinaminpathy et al., 2012), in which x_w is a proportion of the asset w that is sold by i, and α is an indicator of market liquidity that is directly related to confidence C, $\alpha = 1 - C$. If any bank suffers from the capital default based on the shocks from substeps c and d, its health once it is recalculated will be 0. This automatically qualifies such banks for liquidation in the next time step.

(3) Apply decision rule 2 (see Section 3.4) and withdraw short-term loans if condition is satisfied. It is assumed that loans are perfectly divisible and partial withdrawals are possible. Then, record all banks that become illiquid during the withdrawal in order to be liquidated in the next time step. Note that the second decision rule is applied first as otherwise it would be possible to withdraw long-term loans in a single time step.

(4) Apply decision rule 1 (see Section 3.4) and administer shortening of long-term loans if the condition is satisfied.

(5) Recalculate the network and other parameters and go back to step 1 for the next time step.

C. Calculation of the Size of an Average–Bank Neighborhood

In the main text, we stated that the size of a particular bank neighborhood is probabilistic. Here, we provide a simple

calculation of the size of an average-bank neighborhood for a small and a large bank.

Let us consider a network of N banks, of which N_s are small and N_b are large. Let L be the out-degree of a small bank and q be the ratio of average degree of a large bank over average degree of a small bank. In the main text we assumed that in- and out-degrees of all banks are drawn independently from a Poisson distribution with mean z for small and qz for large banks. Here, we make a simplifying assumption that out-degrees of all small (large) banks are equal. When tested in a simulation this assumption did not change our previously reported results. To determine the size of an average-bank neighborhood of a bank i, which includes i's borrowers and lenders, we have to calculate the expected number of unique small and large banks connected to i. For this purpose, let us define random variables:

$$X_{\overrightarrow{ij}} = \begin{cases} 1 & \text{if there is a directed connection from } i \text{ to } j \\ 0 & \text{otherwise} \end{cases}$$
$$X_{ij} = \begin{cases} 1 & \text{if there is a connection between } i \text{ to } j \\ 0 & \text{otherwise} \end{cases} \tag{C.1}$$

with the correspondence $X_{ij} = \max\{X_{\overrightarrow{ij}}, X_{\overrightarrow{ji}}\}$. From here it follows that

$$\begin{aligned} E\left[X_{ij}\right] &= 1P\left(X_{ij}=1\right) + 0P\left(X_{ij}=0\right) \\ &= P\left(X_{ij}=1\right) = 1 - P\left(X_{ij}=0\right). \end{aligned} \tag{C.2}$$

Given there are small and large banks in the system, there are three cases: one bank is small and another is large, both banks are small, and both banks are large.

Case 1. Both banks (i and j) are small.

The probability that a connection originating from i goes to j is a fraction of the in-degree of j and the total in-degree of all banks in the network except for i:

$$w_{ij}^{ss} = \frac{L}{(N_s-1)L + qN_bL} = \frac{1}{(N_s-1)+qN_b}. \tag{C.3}$$

From Equations C.4 and C.13 it follows that the probability that there is at least one connection from i to j is

$$P\left(X_{\overrightarrow{ij}}=1\right) = 1 - P\left(X_{\overrightarrow{ij}}=0\right) = 1 - \left(1 - w_{ij}^{ss}\right)^L, \tag{C.4}$$

where $(1 - w_{ij}^{ss})^L$ is the probability that none of L outgoing links of i connect to j (note that the probabilities of connections are independent).

Then, the probability that there is at least one connection between i and j irrespective of the direction is

$$\begin{aligned} P\left(X_{ij}=1\right) &= P\left(X_{\overrightarrow{ij}}=1 \cup X_{\overrightarrow{ji}}=1\right) \\ &= P\left(X_{\overrightarrow{ij}}=1\right) + P\left(X_{\overrightarrow{ji}}=1\right) \\ &\quad - P\left(X_{\overrightarrow{ij}}=1 \cap X_{\overrightarrow{ji}}=1\right) \\ &= 2P\left(X_{\overrightarrow{ij}}=1\right) - P^2\left(X_{\overrightarrow{ij}}=1\right) \\ &\quad \cdot \left(\because P\left(X_{\overrightarrow{ij}}\right) = P\left(X_{\overrightarrow{ji}}\right) \& X_{\overrightarrow{ij}} \coprod X_{\overrightarrow{ji}}\right) \\ &= P\left(X_{\overrightarrow{ij}}=1\right)\left(2 - P\left(X_{\overrightarrow{ij}}=1\right)\right) \\ &= \left(1 - \left(1 - w_{ij}^{ss}\right)^L\right)\left(1 + \left(1 - w_{ij}^{ss}\right)^L\right) \\ &= 1 - \left(1 - w_{ij}^{ss}\right)^{2L}. \end{aligned} \tag{C.5}$$

We can now express the expectation of Y_i^s, which is the number of unique small banks connected to i, as:

$$E(Y_i^s|L) = \sum_{i=1}^{N_s-1} P\left(X_{ij}=1|L\right) = (N_s-1)\left(1 - \left(1 - w_{ij}^{ss}\right)^{2L}\right). \tag{C.6}$$

Case 2. One bank is small and another is large.

Case 2a. i is a small and j is a large bank.

The probability that a connection originating from i goes to j is

$$w_{ij}^{sb} = qw_{ij}^{ss} = \frac{q}{(N_s-1)+qN_b}. \tag{C.7}$$

The probability that there is at least one connection from i to j is

$$P\left(X_{\overrightarrow{ij}}=1\right) = 1 - P\left(X_{\overrightarrow{ij}}=0\right) = 1 - \left(1 - w_{ij}^{sb}\right)^L. \tag{C.8}$$

Case 2b. i is a large and j is a small bank.

The probability that a connection originating from i goes to j is

$$w_{ij}^{bs} = \frac{L}{N_sL + q(N_b-1)L} = \frac{1}{N_s + q(N_b-1)}. \tag{C.9}$$

Similarly to Case 2a, it follows that

$$P\left(X_{\overrightarrow{ij}}=1\right) = 1 - P\left(X_{\overrightarrow{ij}}=0\right) = 1 - \left(1 - w_{ij}^{bs}\right)^{qL}. \tag{C.10}$$

From Case 2a and Case 2b we derive the probability that there is at least one connection between i and j irrespective of the direction:

$$\begin{aligned} P(X_{ij}=1) &= P\left(X_{\overrightarrow{ij}}=1 \cup X_{\overrightarrow{ji}}=1\right) \\ &= P\left(X_{\overrightarrow{ij}}=1\right)+P\left(X_{\overrightarrow{ji}}=1\right) \\ &\quad - P\left(X_{\overrightarrow{ij}}=1\right)P\left(X_{\overrightarrow{ji}}=1\right) \\ &= 2-\left(1-w_{ij}^{sb}\right)^{L}-\left(1-w_{ij}^{bs}\right)^{qL} \\ &\quad - \left[\left(1-\left(1-w_{ij}^{sb}\right)^{L}\right)\left(1-\left(1-w_{ij}^{bs}\right)^{qL}\right)\right]. \end{aligned} \tag{C.11}$$

Then, the expectations for the number of unique large banks connected to a small bank i (Case 2a) and the number of unique small banks connected to a large bank i (Case 2b) are

$$\begin{aligned} E\left(Y_i^b|L\right) &= \sum_{i=1}^{N_b} P(X_{ij}=1|L) \\ &= N_b\left\{2-\left(1-w_{ij}^{sb}\right)^{L}-\left(1-w_{ij}^{bs}\right)^{qL}\right. \\ &\quad \left. - \left[\left(1-\left(1-w_{ij}^{sb}\right)^{L}\right)\left(1-\left(1-w_{ij}^{bs}\right)^{qL}\right)\right]\right\}, \\ E(Y_i^s|L) &= \sum_{i=1}^{N_s} P(X_{ij}=1|L) \\ &= N_s\left\{2-\left(1-w_{ij}^{sb}\right)^{L}-\left(1-w_{ij}^{bs}\right)^{qL}\right. \\ &\quad \left. - \left[\left(1-\left(1-w_{ij}^{sb}\right)^{L}\right)\left(1-\left(1-w_{ij}^{bs}\right)^{qL}\right)\right]\right\}. \end{aligned} \tag{C.12}$$

Case 3. Both banks (i and j) are large.

The probability that a connection originating from i goes to j is

$$w_{ij}^{bb}=qw_{ij}^{bs}=\frac{q}{N_s+q(N_b-1)}. \tag{C.13}$$

The probability that there is at least one connection from i to j is

$$P\left(X_{\overrightarrow{ij}}=1\right)=1-P\left(X_{\overrightarrow{ij}}=0\right)=1-\left(1-w_{ij}^{bb}\right)^{qL}. \tag{C.14}$$

Similarly to the previous cases, it follows that

$$\begin{aligned} P(X_{ij}=1) &= P\left(X_{\overrightarrow{ij}}=1 \cup X_{\overrightarrow{ji}}=1\right) \\ &= P\left(X_{\overrightarrow{ij}}=1\right)\left(2-P\left(X_{\overrightarrow{ij}}=1\right)\right) \\ &= 1-\left(1-w_{ij}^{bb}\right)^{2qL}. \end{aligned} \tag{C.15}$$

The expectation of the number of unique large banks connected to a large bank i is

$$E\left(Y_i^b|L\right)=\sum_{i=1}^{N_b-1} P(X_{ij}=1|L)=(N_b-1)\left(1-\left(1-w_{ij}^{bb}\right)^{2qL}\right). \tag{C.16}$$

Data Availability

We do not use any data that are not publicly available.

Disclosure

Part of this work is based on the first author's dissertation [35].

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

This research was funded by a pre-doctoral fellowship from the Max Planck Institute for Human Development to the first author. We thank Sujit Kapadia, Jolene Tan, members of the ABC group, the IMPRS Uncertainty School for their helpful comments on an earlier draft; and Anita Todd for editing the manuscript. This paper is dedicated to the memory of Professor Alfred Hübler who inspired our exploration of complexity.

References

[1] D. Covitz, N. Liang, and G. A. Suarez, "The evolution of a financial crisis: collapse of the asset-backed commercial paper market," *Journal of Finance*, vol. 68, no. 3, pp. 815–848, 2013.

[2] V. V. Acharya and O. Merrouche, "Precautionary hoarding of liquidity and interbank markets: evidence from the subprime crisis," *Review of Finance*, vol. 17, no. 1, pp. 107–160, 2013.

[3] J. Berrospide, *Liquidity Hoarding and the Financial Crisis: An Empirical Evaluation (FEDS Working Paper no. 2013–3)*, Washington, DC, 2012 http://www.bis.org/bcbs/events/bhbibe/berrospide.pdf.

[4] D. Gale and T. Yorulmazer, "Liquidity hoarding," *Theoretical Economics*, vol. 8, no. 2, pp. 291–324, 2013.

[5] M. Pritsker, "Knightian uncertainty and interbank lending," *Journal of Financial Intermediation*, vol. 22, no. 1, pp. 85–105, 2013.

[6] N. Arinaminpathy, S. Kapadia, and R. M. May, "Size and complexity in model financial systems," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 109, no. 45, pp. 18338–18343, 2012.

[7] P. Glasserman and H. P. Young, "Contagion in financial networks," *Journal of Economic Literature*, vol. 54, no. 3, pp. 779–831, 2016.

[8] S. Battiston, D. Delli Gatti, M. Gallegati, B. Greenwald, and J. E. Stiglitz, "Liaisons dangereuses: increasing connectivity, risk sharing, and systemic risk," *Journal of Economic Dynamics and Control*, vol. 36, no. 8, pp. 1121–1141, 2012.

[9] D. Brockmann and D. Helbing, "The hidden geometry of complex, network-driven contagion phenomena," *Science*, vol. 342, no. 6164, pp. 1337–1342, 2013.

[10] S. Martinez-Jaramillo, B. Alexandrova-Kabadjova, B. Bravo-Benitez, and J. P. Solórzano-Margain, "An empirical study of the Mexican banking system's network and its implications for systemic risk," *Journal of Economic Dynamics and Control*, vol. 40, pp. 242–265, 2014.

[11] E. Nier, J. Yang, T. Yorulmazer, and A. Alentorn, "Network models and financial stability," *Journal of Economic Dynamics and Control*, vol. 31, no. 6, pp. 2033–2060, 2007.

[12] S. Thurner and S. Poledna, "DebtRank-transparency: controlling systemic risk in financial networks," *Scientific Reports*, vol. 3, no. 1, p. 1888, 2013.

[13] C. Minoiu and J. A. Reyes, "A network analysis of global banking: 1978–2010," *Journal of Financial Stability*, vol. 9, no. 2, pp. 168–184, 2013.

[14] F. Caccioli, J. D. Farmer, N. Foti, and D. Rockmore, "Overlapping portfolios, contagion, and financial stability," *Journal of Economic Dynamics and Control*, vol. 51, pp. 50–63, 2015.

[15] F. Caccioli, M. Shrestha, C. Moore, and J. D. Farmer, "Stability analysis of financial contagion due to overlapping portfolios," *Journal of Banking & Finance*, vol. 46, pp. 233–245, 2014.

[16] D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi, "Systemic risk and stability in financial networks," *American Economic Review*, vol. 105, no. 2, pp. 564–608, 2015.

[17] F. Allen and D. Gale, "Financial contagion," *Journal of Political Economy*, vol. 108, no. 1, pp. 1–33, 2000.

[18] L. Eisenberg and T. H. Noe, "Systemic risk in financial systems," *Management Science*, vol. 47, no. 2, pp. 236–249, 2001.

[19] M. Elliott, B. Golub, and M. O. Jackson, "Financial networks and contagion," *American Economic Review*, vol. 104, no. 10, pp. 3115–3153, 2014.

[20] L. P. Hansen and T. J. Sargent, "Fragile beliefs and the price of uncertainty," *Quantitative Economics*, vol. 1, no. 1, pp. 129–162, 2010.

[21] P. Gai, A. Haldane, and S. Kapadia, "Complexity, concentration and contagion," *Journal of Monetary Economics*, vol. 58, no. 5, pp. 453–470, 2011.

[22] V. V. Acharya and D. Skeie, "A model of liquidity hoarding and term premia in inter-bank markets," *Journal of Monetary Economics*, vol. 58, no. 5, pp. 436–447, 2011.

[23] M. K. Brunnermeier and L. H. Pedersen, "Predatory trading," *Journal of Finance*, vol. 60, no. 4, pp. 1825–1863, 2005.

[24] D. W. Diamond and R. G. Rajan, "Fear of fire sales, illiquidity seeking, and credit freezes," *The Quarterly Journal of Economics*, vol. 126, no. 2, pp. 557–591, 2011.

[25] R. J. Caballero and A. Krishnamurthy, "Collective risk management in a flight to quality episode," *Journal of Finance*, vol. 63, no. 5, pp. 2195–2230, 2008.

[26] B. R. Routledge and S. E. Zin, "Model uncertainty and liquidity," *Review of Economic Dynamics*, vol. 12, no. 4, pp. 543–566, 2009.

[27] A. Zawadowski, "Interwoven lending, uncertainty, and liquidity hoarding: Boston University School of Management Research Paper No. 2011–13," *SSRN Electronic Journal*, 2011.

[28] R. J. Caballero and A. Simsek, "Fire sales in a model of complexity," *Journal of Finance*, vol. 68, no. 6, pp. 2549–2587, 2013.

[29] M. Boss, H. Elsinger, M. Summer, and S. Thurner, "Network topology of the interbank market," *Quantitative Finance*, vol. 4, no. 6, pp. 677–684, 2004.

[30] R. Cifuentes, G. Ferrucci, and H. S. Shin, "Liquidity risk and contagion," *Journal of the European Economic Association*, vol. 3, no. 2-3, pp. 556–566, 2005.

[31] Bank of England, *Instruments of Macroprudential Policy (Bank of England Discussion Paper)*, London, England, 2011 http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Instruments+of+macroprudential+policy#3.

[32] F. Bräuning and F. Fecht, "Relationship lending in the interbank market and the price of liquidity," *Review of Finance*, vol. 21, no. 1, pp. 33–75, 2016.

[33] A. Temizsoy, G. Iori, and G. Montes-Rojas, "The role of bank relationships in the interbank market," *Journal of Economic Dynamics and Control*, vol. 59, pp. 118–141, 2015.

[34] A. W. Lo, "Complexity, concentration and contagion: a comment," *Journal of Monetary Economics*, vol. 58, no. 5, pp. 471–479, 2011.

[35] S. Davidovic, *The ecology of financial markets: from analogy to application, [Ph.D. thesis]*, Humboldt University of Berlin, 2016.