

Contagion of Financial Crises under Local and Global Networks

Alessandra Cassar and Nigel Duffy*

Final Version

Abstract

As the world economy becomes increasingly global, will the financial sector become more stable or fragile? In this paper we study how the pattern of relations linking financial institutions - the network - affects the diffusion of a financial crisis. We analyze two such networks with a computational model: the local network, in which each bank is allowed to interact only with the most immediate neighbors, and the global network, in which each bank is allowed to interact with banks located anywhere in the system.

We find that the network matters both for the amount of illiquidity in the system and for the spread of bankruptcy. When interactions are local, bankruptcy spreads slower but illiquidity hits harder. When interactions are global, bankruptcy spreads faster, but illiquidity presents fewer problems. We conclude that a global system, in which financial institutions are not restricted to interact only with close neighbors, is more efficient in collecting and allocating funds, but is more vulnerable to contagion of bankruptcy crises.

1 Introduction

As the economy becomes increasingly global, will it be more or less vulnerable to financial crises? The traditional answer is unclear: while a global economy offers diversification against shocks, leading to an increasing stability in the system, new risks arise that might not balance across agents or regions, leading to an increasing instability. In this paper we address this question by analyzing the effect of local interactions and geographically dispersed interactions during financial crises. As in

*We are particularly grateful to Daniel Friedman for his continuous encouragement and substantial help throughout all phases of this project. An earlier version of this paper was presented during a subsidized stay at the Santa Fe' Institute during the 1999 Summer School for Economists where it received helpful comments from John Miller, Scott Page and John Padgett. Finally, we would like to thank David Kaun for his suggestions which have significantly improved the paper.

Friedman (1998), we concentrate on one channel of contagion that has not been studied much by economists: the pattern of overlapping claims that financial institutions hold on each other - the network structure.

Since we are interested in isolating the effect of different networks, we do not consider two other important channels of contagion that might play a relevant role during a crisis. The first relates to the idea that bank runs are self-fulfilling prophecies, so that a crisis originating in one region of the economy may create a self-fulfilling expectation of crisis elsewhere, Kindleberger (1978), Diamond and Dybvig (1983), Chang and Velasco (1999). The second is the propagation mechanism via the international currency markets in which financial crises spread from one country to the other, Calvo (1994), Eichengreen (1996), Chang (1998), Chang and Velasco (1998).

In our model we exclude such channels, so that a crisis in one region can spread to another region only through the web of overlapping financial assets that banks hold on each other. This channel of contagion may have different effects depending on the particular network architecture, i.e. the pattern of claims linking financial institutions. We are interested in analyzing how financial crises spread under two such networks: the local network, when financial institutions are allowed to interact only with their most immediate neighbors, and the global network, when financial institutions can be linked to any other institution in the system. We employ a computational model based on Friedman's (1998) model, in which the fragility of the financial market arises from contagion between neighboring financial institutions. Internal forces and random events can push a normal functioning financial system towards a critical insolvent state. Here, an insolvency avalanche can engulf a large connected subset of banks.

Once a specific network takes the place of a degenerate one, the resulting systems behave in ways still not completely understood, Page (1999). On non-trivial architectures, like a lattice or a torus, exact solutions are difficult to find with traditional mathematical tools. For this reason, we use a computational model of an artificial world simulating a financial crisis. For this class of problems, agent-based computations are a good way to look at network interactions which have not been fully explored by economists.

We find that the network structure matters for contagion of both insolvency and illiquidity. In particular, under the system with local interactions, insolvency spreads more slowly, but illiquidity hits harder. Under the random network, the reverse is the case. This result may explain why the banking sector has historically been interested in developing an extensive network of relations.

The idea that the architecture of these interconnections can be a source of financial distress has a very long tradition. Even before appearing in the economic literature, this concept is already present in Florence during the Renaissance with its banking system (the origin of mercantile capitalism). The Medici bank had a geographically wide organizational structure, and much thought was given to the system regulating the complex and extensive relations among branches, debtors and credi-

tors, Goldwhite (1987), McLean and Padgett (1997), De Roover (1970). More than four centuries later, in the United States, the National Bank Act of 1864 prohibited national banks from branching, either within a state or across states. The McFadden Act of 1927, and the Banking Act of 1933, finally permitted national banks to branch within a state, but not beyond.

Because of the importance of the real effect on output and growth, financial crises have long been of major concern to economists, Agenor, Aizenman, and Hoffmaister (1999). Only recently however has the pattern of interconnectedness as a channel for contagion come under scrutiny.

Eisenberg (1995), studying the propagation of a bank run, analyzes how increasing the connectivity of the bank system affects the speed of default when the network of inter-bank relations is random. When bankruptcy propagation is fast relative to the maturity of the obligations, the bankruptcy propagation over the network can be less extensive than when bankruptcy propagation is slower. However, when bankruptcy propagation is slower the network tends to reach a complete shutdown. Therefore, shortening the maturity of the obligations increases the extent and speed of the propagation. The network studied by Eisenberg is random. He argues that it is a reasonable and less specific assumption to model the network as random since traders usually cannot predict their trades, especially in the foreign exchange market, and, with secondary markets in debt issues, holders of a bank's obligations do tend to be random.

Allen and Gale (2000) argue, instead, that the network should be local, since banks usually specialize in particular areas of business or have closer connections with banks that operate in the same geographical or political unit. Thus claims may tend to be concentrated in neighboring banks. They analyze how contagion depends on the pattern of interconnectedness generated by the cross-holdings of deposits. Allen and Gale find that when the inter-bank market is complete, i.e. when each region has banks connected with banks in all the other regions, the initial impact of a financial crisis might be either limited to the troubled region, or a small loss might be spread across regions through contagion. When, however, the market is incomplete, in that each region has banks connected only with banks in a small number of other regions, the feature of contagion depends critically on the degree of connectedness (i.e. whether or not any two regions are connected not directly but by a chain of overlapping bank liabilities). With incomplete but connected markets, the initial impact of a financial crisis may be felt very strongly by the neighboring regions and progressively extend to all the system. Only if the degree of connectedness is low, as in the case of disconnected market structure in which certain regions remain isolated, is there no contagion.

In the next section we present the framework for the computational model. The behavior of the model under different networks is presented in section three. Section four reports the results of the simulation and section five the conclusions.

2 The Bank Avalanche Model

In real-world financial crises, the mechanism for contagion may entail several channels, including defaults in international payments or self-fulfillment of crisis expectations. Here, however, following Friedman (1998), we consider a single channel of contagion only, the overlapping of claims that banks hold on each other. In this model, events such as public default on loans, and internal forces like the pattern of inter-linkages among banks, can push the financial system towards a critical insolvent state. At such a point, an insolvency avalanche may start, causing a large number of connected banks to default. Increasing individual rationality does not necessarily help when the crisis happens through this channel, and appropriate simple interventions are not immediately apparent.

We employ a computational model with the intention of analyzing whether the network of interlinks between banks affects how financial crises spread. In particular, we will analyze the affects of two such architectures: the local and the global network. In the local networks the interactions among banks are restricted to geographically close neighbors. In the global network banks are free to be connected with banks anywhere in the system.

The computational model has been implemented on the assumptions presented in the following three subsections: Connection Types (which discusses the financial claims available), Behavior Rules (which presents the rules followed by banks and non-financial transactors), and Network Structures (which explains the alternative pattern of ties linking financial institutions).

2.1 Connection Types (or States of the Agents)

Consider a finite network with exogenously specified nodes and inter-node connections. Some of these nodes represent banks, some others represent non-financial transactors (NFT), which can be considered to be households or business firms. The inter-node connections are constituted by financial claims, which therefore define the “neighboring” relations. Each bank is assumed to classify its financial claims into five different types.

- D = Deposits: number of deposits that non-financial transactors have at a given bank. The events which cause increase and decrease in deposits are initiated by the NFT according to the rules described below.
- L = Loans: number of loans that a given bank issues to non-financial transactors. The events which cause increase and decrease in loans are initiated by the NFT according to the rules described below.
- ED = EuroDeposits: number of Eurodollar deposits that the given bank accepts from its connected banks. Neighboring banks initiate the events of increase and decrease of EuroDeposits subject to the rules listed below.

- EL = EuroLoans: number of EuroDollar deposits that the given bank places at (i.e. loans to) neighboring banks. The given bank initiates the events of increase and decrease of EuroLoans subject to the rule listed below, which is the same rule that governs EuroDeposits.
- $Reserve$ = balance on the reserve account. These claims are on the issuer of safe assets, like the national central bank. However, since that issuer plays no active role in the analysis, at least in this paper, it is not represented by an explicit node.

2.2 Behavior Rules (or Methods)

The rules governing banks' behavior are based on the consideration that banks are profit-oriented, and therefore seek net interest revenue. In particular we make the following assumptions on the structure of the interest rates. The cost of funds is the expected return on EuroDeposits. Neglecting the small transaction costs for EuroDollar, we may presume that the interest rates on EuroDeposits and EuroLoans is equal. We assume that the yields on reserve and deposits are lower than the interest rate on EuroDeposits, while the (expected) return on loans is higher. Market conditions may shift the entire yield structure up and down while preserving the ordering. Every time a bank accepts new deposits and agrees on new loans its net interest revenue increases, and therefore it will always conclude these financial transactions whenever possible. Following the same reasoning, the bank loses net interest revenue when it holds excess reserves, so it will try to keep them as low as possible, after considering liquidity issues. In the model implementation, we simplify by omitting an explicit representation of these returns. Nevertheless each bank will operate with respect to the assumed interest structure.

2.2.1 Bank's Rules

Banks are assumed to obey the following rules:

- $B1$: The bank's immediate accommodation to any deposit or loan event is via its reserve account. Thus the immediate impact of $[D \rightarrow D + 1]$, $[ED \rightarrow ED + 1]$, $[L \rightarrow L - 1]$ or $[EL \rightarrow EL - 1]$ is $[R \rightarrow R + 1]$; the immediate impact of $[D \rightarrow D - 1]$, $[ED \rightarrow ED - 1]$, $[L \rightarrow L + 1]$ or $[EL \rightarrow EL + 1]$ is $[R \rightarrow R - 1]$.
- $B2$: Banks actively manage their reserves, with a target balance assumed here to be $R = 1$. When R is pushed above the target, the preferred use of funds are (in order of preference): accommodating loans, EuroLoans, and refunding EuroDeposits, deposits. We assume that a bank's normal contractual obligations preclude refunding its EuroDeposits or its deposits from NFTs without being asked. Thus the events $[D \rightarrow D + 1]$ etc. trigger event $[L \rightarrow L + 1]$ when there

happens to be a loan applicant not yet accommodated; typically those events trigger $[EL \rightarrow EL + 1]$.

- *B3*: The preferred sources of funds to restore R when it is below target are: accepting deposits, EuroDeposits, and calling in EuroLoans and loans. We assume that a bank can call in its EuroLoans (i.e. withdraw its EuroDeposits with a neighbor bank) but not its loans to NFTs. Thus the events $[D \rightarrow D - 1]$ etc. trigger the event $[D \rightarrow D + 1]$ or $[ED \rightarrow ED + 1]$ when there happens to be a depositor not yet accommodated; typically those events trigger $[EL \rightarrow EL - 1]$.
- *B4*: If R falls to 0 and can't immediately be restored to target then the bank is illiquid. In this state it is unable to accept requests to increment L or to decrease D . It will still accept decreases (repayments) in L and increases in D and ED , and use them to increment R , thus restoring itself to the normal liquid state.
- *B5*: Net worth or capital is $K = L + EL + R - D - ED$. We start each node at $R = 1$ and the other four gross claims at $L = EL = D = ED = 0$. All the events described so far preserve K . Loans defaults, as specified in *NFT3* below, decrease K . In our model it is assumed that K never increases - however, if retained, net interest revenues would gradually increase K . In this model we do not account for claims on K , i.e. the stock holders; we assume implicitly that these claimants receive all the net interest revenue.
- *B6*: If $K < 0$ then the bank is bankrupt. In this state all its assets (L , EL and R) are liquidated, and with these funds some of the liabilities are met. Seniority is assigned first to deposits then to the EuroDeposits. The bank cannot do anything but default on the remaining part and become insolvent. Once a bank is insolvent, it is out of the game. Note that liquidating EuroLoans and loans has consequences for neighboring banks and *NFTs*.

2.2.2 Non-financial Transactors' (NFTs) Rules

Non-financial transactors are assumed to obey the following rules.

- *NFT1*: New deposit and loan events at a given bank $[D \rightarrow D + 1]$ and $[L \rightarrow L + 1]$ are initiated by *NFTs* randomly with probability $p_D > 0$ and $p_L > 0$ per time period.
- *NFT2*: Withdrawal and repayment events at a given bank $[D \rightarrow D - 1]$ and $[L \rightarrow L - 1]$ are initiated by *NFTs* randomly with probability $p_D > 0$ and $p_L > 0$ per time period.
- *NFT3*: With probability $p_{DEFAULT} \geq 0$ per unit of time per unit of L there is loan default, causing $[L \rightarrow L - 1]$, which induces $[K \rightarrow K - 1]$, but no change in R .

2.3 Network Architectures

We implemented the framework defined above on three alternative systems of interactions on a torus (a lattice where the sides wrap around).

- *No interactions.* First, any sort of interactions between banks is assumed away. Even if this cannot be considered a proper network, since its nodes, the banks, are not connected, it will be very useful as a benchmark comparison.
- *Local interactions.* Each bank now has the possibility to interact with the four spatially most immediate neighbors.
- *Global interactions.* Finally, each bank is allowed to interact with four other banks located anywhere on the torus. This network is a particular case of a random graph in which the number of edges, i.e. links, for each bank is constrained to be exactly four, instead of an average of four. We adopt this particular kind of random network in order to make it comparable with the local network. In fact, in both networks banks start with exactly four neighbors; should systematic differences arise we interpret them in terms of the two different network architectures only. This avoids the need to separate the network effect from the effect due to differences in the bank dimensions, as represented by the number of neighbors.

The major concern of this paper is the extent to which changing the architecture of the interactions produces different results in terms of diffusion of illiquidity and insolvency. We discuss the behavior of the model under these alternative networks in the following sections.

3 Financial Crises on Networks

The social interactions we have defined induce a network structure on the banking system. Such network structures can be analyzed using tools from graph theory which, while seldom used by economists, give us considerable intuition about the behavior of the systems in which we are interested. We employ two concepts from graph theory namely: the i^{th} -neighborhood of a node and the characteristic path length.

The size of the i^{th} -neighborhood of a node measures the number of nodes reachable in exactly i steps. Thus, in our model, the i^{th} -neighborhood measures the maximum number of other banks that a bank can interact with through exactly i intermediaries. This quantity will be larger for a node when the $(i-1)^{\text{th}}$ -neighbors are not connected to each other. This number is higher in our global network than in the local network. A trivial calculation shows that the i^{th} -neighborhood for the local model is of size $4 \times i$ while for the global network it is of size 4^i (on an infinite lattice). For example,


if we allow a bank's neighbors to contact their neighbors, and we allow these 2^{nd} neighbors to do the same, a bank on the local network can reach $4 + 4 \times 2 + 4 \times 3 = 24$ other banks each time-step, while a bank on the global network can reach $4 + 4^2 + 4^3 = 84$. When applied to our model, this number should be treated with caution, first because our lattice is not infinite, and second when banks become insolvent their connections disappear changing the neighborhood structure of the network. However, qualitatively this characterization should be correct. We will see in the next subsection how this impacts on contagion.

The characteristic path length is the median of the means of the shortest path lengths connecting each node to all other nodes. In our model, it represents a measure of the average minimum number of links that must be traversed to get between a pair of banks. The characteristic path length is higher in a local network than in a random network, Watts (1999)¹. Informally, consider again the size of the i^{th} -neighborhood. The proportion of banks reachable in i steps is larger for the global network than for the local network. As a result, the average number of steps required to reach a node from any other is larger in the local network than in the random network. Summarizing, for the networks we consider:

		Size of i^{th} Neighborhood	
		Small	Big
Path Length	Short		Global Network
	Long	Local Network	

3.1 Local Network

Assume that N banks are arranged on a torus, and have relations with four other banks and with a single non financial transactor (NFT) not represented on the torus for simplicity. In the case of a local network, each bank i_0 is connected with the North-neighbor i_1 , the East-neighbor i_2 , the South-neighbor i_3 , and the West-neighbor i_4 :



The normal functioning of financial intermediation is insured every time $NFTs$ initiate a deposit, say at bank i_4 and a loan at bank i_1 (or i_3 and i_2). Bank i_4 reacts increasing its EuroDeposits, or accommodating a requested EuroLoan at the neighboring bank i_0 . Similarly bank i_1 reacts withdrawing a EuroDeposit, or requesting a

¹In fact, this result applies to a different model of random network. However, the same reasoning appears to apply in this case also.

EuroLoan, at bank i_0 . A random event, like a loan request or deposit or EuroDeposit withdrawal from a bank that does not have enough liquidity (as when i_0 does not receive the EuroDeposit from i_4 while still receiving the EuroLoan request from i_1), is able to push the financial system towards a critical state in which illiquidity can spread throughout the connected subset of banks.

The dynamics of an insolvency crisis are similar to the dynamics of a liquidity crisis, although the details are rather different. Again there are internal forces that push the financial system towards a sub-critical solvency state, as when a bank, say i_4 , receives an NFT default on its loan. A second random event, such as a second default from an NFT or from a neighboring bank, can push the system to a critical state (i_4 defaults on i_0). At that point, assuming away central bank intervention, an avalanche might start with a large subset of banks involved sequentially through a chain of defaults.


When interactions are local, a bank's neighbors are not directly connected, but these neighbors have neighbors in common. For example, two of the neighbors of i_0 , i_1 and i_4 , have one other neighbor in common beside i_0 : the North-neighbor of i_4 which coincides with the West-neighbor of i_1 . Following our previous reasoning the number of banks reachable in j or fewer steps is the sum of the size of the i^{th} -neighborhoods for $i < j$, implying that, in j steps, the actual number of banks that a bank can reach is smaller in a local network than in a global network. The effect of reaching a lower number of banks means that the pool of available reserves, from which loan requests may be satisfied, is smaller, making the local system more vulnerable to illiquidity.

While constituting a barrier to obtaining funds in case of need, a smaller neighborhood and a longer distance between banks decreases the spread of insolvency. The high degree of overlap between neighbors implies that the spread of bankruptcy may eventually die off in an isolated region, while the long path length implies that a longer chain of defaults is required to reach geographically distant banks.

In conclusion, we expect the local system to be more vulnerable to illiquidity, but less vulnerable to bankruptcy than the global system.

3.2 Global Network

Assume again that N banks are arranged on a torus, but this time we allow them to have relations with four other banks located randomly:



Bank i_0 has now its four neighbors, i_1 , i_2 , i_3 and i_4 , some located in regions far away. Considering the size of the i^{th} -neighborhood, it is now highly unlikely that neighbors' neighbors overlap as much as in the local case. As a result, in this global network, in the same number of steps each bank can now contact many more banks, increasing the probability of obtaining EuroLoans in case of need, or allocating EuroDeposits in case of excess reserve. This should result in more efficient collection and allocation of funds, with the beneficial effect of reducing illiquidity.

At the same time, however, each bank can now be reached by any other in a small number of steps, making each bank almost immediately vulnerable to defaults originated anywhere in the network.

Within the confines of the model thus far, it seems that while globalization improves the efficiency of the system in terms of collecting and allocating funds, such financial integration increases the possibility of bankruptcy contagion.

Given the difficulty of finding an analytical solution to illiquidity and insolvency for each model, we analyze them using an agent-based simulation implemented in Swarm.

4 Model Implementation in Swarm

Our results were obtained using an agent-based computational model implemented using the Swarm package, developed and maintained by the Santa Fe' Institute. The most fundamental elements (agents) in our simulation are *banks* which interact with each other through *banking relationships* (or simply *relationships*). Ideally in such a simulation the banks act autonomously and are not centrally coordinated. This is precisely the form of an agent-based simulation and is ideal for implementation in an intuitive, bottom-up manner using Swarm.

All banks in our simulation have the same repertoire of available actions, however, their choice of action is motivated by their current state. This fits perfectly with the Object Oriented Programming (OOP) paradigm embodied in such languages as Java, Objective C and C++. In this paradigm *classes* are defined which share a common behavior. This behavior need only be defined or implemented once, for all *objects* belonging to the class. In our simulation, each bank is an object in the bank class, therefore we only need to implement *bank* behavior once. Each object which instantiates a class may also hold additional state information, which is specific to that object. Each bank in our simulation has state which describes its current assets and relationships with other banks. The OOP paradigm is extremely intuitive and efficient for agent-based simulations since it allows one to define agents in terms of their interactions with the environment.

In our implementation, the two most important classes are banks and relationships. Each bank maintains a list of relationships through which they interact with each other. Banks have access to generic actions on these relationships, such as: request loan, offer loan, demand deposit, default on loan. Relationship objects maintain

the current state of a relationship between two banks, they are responsible for coordinating the offers and requests of the two parties. In addition, the relationship objects coordinate and control the activities of banks upon completion of a transaction. Embodying a banking relationship as an object has several conceptual advantages. First, it trivially allows different banks to have different numbers of relationships. Second, it removes the need to arbitrate transactions centrally as this is now done by the relationship object itself. Third, it allows banks to trade or exchange relationships easily.

Once the central players or agents in an agent-based simulation are defined, one needs to coordinate their activities and the events in the environment. This is the primary way in which we took advantage of Swarm. Swarm consists of a set of Objective C and Java libraries designed to ease the implementation of agent-based simulations. Swarm provides a range of useful tools, including pre-defined environments and agents. For our simulation we used the Swarm facilities for constructing user interfaces, but more importantly we used the event management facilities. The event management facilities are the means by which one may simulate autonomous behavior. In Swarm, one defines a schedule of events or actions and the Swarm libraries handle the details of carrying out the actions for each agent. One may control the schedule of activities to make them appear simultaneous, occurring in discrete time steps or one may easily implement asynchronous behavior. Our simulation takes advantage of both facilities. The actions of non-financial transactors are synchronized such that all banks appear to receive deposits or withdrawal requests simultaneously. The banks then act asynchronously and autonomously to resolve these requests and negotiate inter-bank transactions. The event management facilities provided by Swarm would be difficult to replicate from scratch in Java or C and greatly simplified the implementation of our simulation.

5 Results

We start with results where banks are not allowed to interact. We then introduce inter-bank claims and compare the system with local interactions to the one with global interactions. We ran the simulation on each of the 3 models with 4900 banks on a 70×70 torus and report results averaged over 10 independent runs. The means across runs of the variables Illiquidity (percentage of illiquid banks²) and Insolvency (percentage of insolvent banks³) were calculated for each time-step, for each model.

In addition to the dimension of the torus (which represents the total number of banks), the parameters of the model are: $p_{Default}$, p_D and p_L and $CountMax$, where:

- $p_{Default}$ is the probability per unit of time per unit of L with which a bank gets a default on its loans. The results we report are obtained keeping this probability

$$^2 \text{Illiquidity} = \frac{\# \text{ illiquid banks}}{\# \text{ banks} - \# \text{ insolvent} - \# \text{ in bankruptcy process}} \%$$

$$^3 \text{Insolvency} = \frac{\# \text{ insolvent banks}}{\# \text{ banks}} \%$$

constant and equal to $p_{Default} = 0.01$. Increasing $p_{Default}$ speeds up the process toward the final state of total insolvency.

- p_D and p_L are the probabilities per unit of time with which a bank gets a unitary increase or decrease in its deposits or loans. These probabilities are assumed $p_D > 0, p_L > 0$, and for simplicity are set identical and equal to $p_D = p_L = 0.1$.
- *CountMax* is the maximum number of times that a bank can negotiate with each neighbor per unit of time. Effectively it is the maximum number of intermediaries that can be involved in the fulfillment of a request or an offer for an inter-bank loan. For the model with no interactions we set $CountMax = 0$, for the models with interactions we set $CountMax = 3$. Values of *CountMax* greater than 1 mean that when a bank makes a request to its neighbors for an inter-bank loan (or an offer of an inter-bank deposit), that cannot be accommodated, these neighbors can ask their neighbors, and so on for *CountMax* steps. $CountMax > 1$ therefore, means that each bank can actually reach a number of banks higher than its immediate neighbors, a concept closely related to the sum of the size of the i^{th} -neighborhood (see above). Increasing *CountMax* implies that, for each unit of time, a higher number of banks can be reached. Since our network is finite, increasing *CountMax* (up until the point where all banks can be contacted in a unit of time) has the effect of decreasing the differences among the local and the global network.

Next we report the results of the simulations. At first the possibility of interactions is assumed away, then the rule for neighbor relations is activated, and we observe whether the network actually matters for the spread of financial crises.

5.1 Illiquidity Crises

Here we analyze the behavior of illiquidity in the three alternative models, see Figure 1 and Figure 2 (for each model we plot the mean values and standard deviations of illiquidity).

When interactions between financial institutions are not permitted, banks become illiquid if unable to accommodate a loan request from NFTs through their reserves. Reserves, in the no interactions model, can increase only through unsolicited NFT deposits. Since it takes time for the banks to accumulate enough reserve, illiquidity appears very soon, it reaches its peak after a few periods, after which it begins a steady decline. The continuous decline of illiquidity is due to a positive bias between the rate of growth of deposits and the rate of growth of loans (loan requests or deposit withdrawals can be satisfied only after the accumulation of deposits).

***** INCLUDE figures 1 and 2 HERE *****

We proceed by allowing banks to accept EuroDeposits and EuroLoans from neighboring institutions (by increasing *CountMax* from 0 to 3). The behavior of the local

and global systems are presented in Figure 1 and Figure 2. Allowing banks to initiate and accept EuroDeposits and EuroLoans, no matter which network is considered, leads to a drastic reduction in illiquidity. Once inter-bank relations are possible, banks can accommodate more financial requests, either from NFTs or from other institutions. This improves the liquidity of the system.

In both the local and global network, the percentage of illiquid banks becomes indistinguishable from the no interaction case after approximately 150 periods. In fact, when roughly 80% of the banks are insolvent (see Figure 4 and Figure 5) enough links are broken that the system behaves as if there were no interactions remaining.

***** INCLUDE figure 3 HERE *****

The comparison of illiquidity between the local network and the global network is presented in Figure 3, which reports the difference between the means of illiquidity between the local and the global network. This difference is positive until about period 120. Tests for the statistical significance of this difference in terms of parameter values and empirical distribution functions, are summarized Table 1. Since we cannot exclude correlation in the behavior of illiquidity from one time-step to the next, we averaged the variables of interest over periods of 15 time-steps. Taking a conservative approach we ran the Wilcoxon test and the Kolmogorov-Smirnov test on the 15 time-step means. As a result, the reported P-values overestimate the correct ones. The difference between the two systems is only around 1%, but even with these conservative tests, we can reject the hypothesis of equal behavior between the local and global system for the central periods.

***** INCLUDE table 1 HERE *****

For the initial periods, the difference in illiquidity is insignificant, because both networks require time for banks to establish relationships with their neighbors. After 50%-60% of the banks are insolvent, again the difference is not significant because the remaining connection structure is too sparse to make a difference between the two networks.

In conclusion, the global network seems more successful in reducing the systems illiquidity, as expected from our previous discussion.

5.2 Insolvency Crises

Let us now turn to the behavior of insolvency. Insolvency hits a bank when its capital goes below 0, and this happens when it receives a double default on loans or EuroLoans either from an NFT or from another connected bank. Since we assumed $p_{Default} > 0$, all systems will end up in a complete state of insolvency: we have no replacement rule in the model, so that an insolvent bank can never go back to an alive state.

***** INCLUDE figures 4 and 5 HERE *****

Once interactions between banks are ruled out, the only event that can trigger a bank's insolvency is a double default on loans from an NFT. Therefore, in a system

with no interactions, contagion (as we define it) cannot arise as a means of spreading insolvency. Only when interactions are allowed, can insolvency spread through a chain of defaults on EuroLoans to neighboring institutions. Figure 4 and Figure 5 show the difference between the number of insolvent banks with $CountMax = 0$ (no interactions case) and $CountMax = 3$ for the local and the global networks (again we plot the mean values and standard deviations for each model). We notice that for either networks insolvency appears sooner and spreads faster with interactions than without. This difference does not only represent the number of bankruptcies caused by neighbors' defaults. In fact, connected systems can accommodate more NFT loan requests, increasing the probability of receiving NFT defaults. Even if the difference in insolvency between $CountMax = 3$ and $CountMax = 0$ does not only estimate contagion, it remains a useful indication of the network effect on insolvency.

***** INCLUDE figure 6 HERE *****

In Figure 6 we compare the proportion of insolvent banks between the two networks. We find that the global network constantly suffers more bankruptcies than the local network. As in the case of illiquidity, this difference is only around 1%, but it is statistically significant for the central periods, as we can see from Table 2. Again, the reported P-values are overestimated because the tests have been conducted on the 15-period means. Once more, after 60% of the banks are insolvent, the difference between the two networks is no longer significant.

***** INCLUDE table 2 HERE *****

In conclusion, despite our conservative measures, we cannot reject the hypothesis that the global network is more vulnerable to insolvency than the local network.

6 Conclusions

In a financial system in which banks are not allowed to interact, available funds are scarce, and bankruptcies cannot spread by contagion. When banks are permitted to hold claims on each other, illiquidity can be greatly reduced by banks via the new access to inter-bank assets. However, these linkages with other institutions cause problems as well, constituting the means through which insolvencies spread and contagion arises as a separate cause of insolvency.

Once we allow for interactions, the particular pattern of inter-bank claims affects how illiquidity and insolvency crises spread. We analyzed, with a computational model, two alternative systems: a local network and a global network. The system with local interactions appears to be more vulnerable to illiquidity, while the system with global interactions suffers more from insolvency. We offer an explanation in terms of two concepts from graph theory: the i^{th} -neighborhood size and the characteristic path length. In a global network the number of banks a bank can reach in a specified number of steps (the i^{th} -neighborhood size) is higher than in a local network. In addition, in a global network any two banks of the system can be connected through a number of intermediaries lower than in a local network. These charac-

teristics provide an intuition for the reason a global network has fewer difficulties satisfying loan requests but it is less protected against a chain of defaults.

This result may contribute to the explanation of why, since the beginning of mercantile capitalism, the first financial institutions, and the Medici Bank in particular, attempted to move further than local business and develop strategic alliances geographically dispersed around Europe. Beyond historical interest implicit here, this research has an important lesson for contemporary economies as well. In this electronic age, where physical impediments to banking relationships are being removed, the benefits derived from an increasingly efficient allocation of financial assets must be weighed against the accompanying increase in the risk of bankruptcy contagion.

References

- [1] Agenor, P. R., Aizenman, J., and Hoffmaister A., 1999, *Contagion, Bank Lending Spreads, and Output Fluctuations*, Working Paper World Bank.
- [2] Allen, F. and Gale, D., 2000, *Financial Contagion*, Journal of Political Economy, Vol.108, pp.1-33.
- [3] Calvo, G., 1994, *Varieties of Capital Market Crises*, Inter-American Development Bank Working Papers.
- [4] Chang, R., 1998, *The Asian Liquidity Crises*, NBER Working Paper No. 6796.
- [5] Chang, R. and Velasco, A., 1998, *Financial Fragility and the Exchange Rate Regime*, New York, C.V. Starr Center-Working Papers.
- [6] Chang, R. and Velasco, A., 1999, *Liquidity Crises in Emerging Markets: Theory and Policy*, NBER Working Paper No. W7272.
- [7] De Roover, R., 1970, *The Rise and Decline of the Medici Bank (1397-1494)*, Harvard University Press.
- [8] Diamond, D. and Dybvig, P., 1983, "Bank Runs, Deposit Insurance and Liquidation", Journal of Political Economy, 91, pp.401-419.
- [9] Eichengreen, B., 1996, *Contagious Currency Crises*, NBER Working Paper No. 5681.
- [10] Eisenberg, L.K., 1995, *Connectivity and Financial Network Shutdown*, Santa Fe Institute Working Papers.
- [11] Friedman, D., 1998, *Bank Avalanche Model*, Private Communication.
- [12] Goldthwaite, R., 1987, "The Medici Bank and the World of Florentine Capitalism", *Past and Present* 114, p. 3-31.
- [13] Kindleberger, C. P., 1978, *Manias, Panics, and Crashes: A History of Financial Crises*. New York: Basic Books, 1978.
- [14] McLean, P.D., and Padgett, J. F., 1997, "Was Florence a Perfectly Competitive Market? Transactional Evidence from the Renaissance", *Theory and Society* 26, p. 209-244.
- [15] Page, S., 1999, *Network Structure Matters*, Unpublished Manuscript.
- [16] Watts, D., 1999, *Small World: the Dynamics of Networks between Order and Randomness*, Princeton University Press.