UNIVERSITEIT GENT

FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE

TWEEKERKENSTRAAT 2
B-9000 GENT
Tel. : 32 - (0)9 – 264.34.61
Fax. : 32 - (0)9 – 264.35.92

WORKING PAPER

Network Models of Financial Contagion: A Definition and Literature Review

Gavin Wims

[1] University College Ghent, Ghent University, Department of Management & Informatics,
Voskenslaan 270, B-9000 Ghent, Belgium
{Gavin.Wims, Manu.DeBacker}@hogent.be

David Martens

University of Antwerp, Department of Environment, Technology, and Technology Management
Prinsstraat 13, B-2000 Antwerp, Belgium
{David.Martens}@ua.ac.be

Manu De Backer

[1] University College Ghent, Ghent University, Department of Management & Informatics,
Voskenslaan 270, B-9000 Ghent, Belgium
{Gavin.Wims, Manu.DeBacker}@hogent.be

July 2011
2011/730

D/2011/7012/35

Network Models of Financial Contagion: A Definition and Literature Review

Gavin Wims[1], David Martens[2], Manu De Backer[1,3,4]

[1] University College Ghent, Ghent University, Department of Management & Informatics,
Voskenslaan 270, B-9000 Ghent, Belgium
{Gavin.Wims, Manu.DeBacker}@hogent.be

[2] University of Antwerp, Department of Environment, Technology, and Technology Management
Prinsstraat 13, B-2000 Antwerp, Belgium
{David.Martens}@ua.ac.be

[3] K.U.Leuven, Department of Management Informatics,
Naamsestraat 69, B-3000 Leuven, Belgium
{Manu.DeBacker}@econ.kuleuven.be

[4] University of Antwerp, Department of Management Information Systems,
Prinsstraat 13, B-2000 Antwerp, Belgium
{Manu.DeBacker}@ua.ac.be

Abstract

Determining the risk of contagious failures due to credit exposures between organisations is a problem that has been the subject of a growing body of literature in recent years. The network model has become a commonly used tool, applied to both theoretical and empirical studies of financial contagion and systemic risk. The purpose of this paper is twofold. First, we propose a definition of the 'Financial System Network' which may be used to define the characteristics of any specific implementation of a network model in this field. Secondly, we evaluate the network models created by other researchers and compare and contrast various aspects of these implementations. We conclude by exploring avenues for future research in the area.

1. Introduction

This paper seeks to provide a definition of the 'Financial System Network' (FSN) which may be used to define the characteristics of any specific implementation of a network model of a financial system characterised by credit linkages between financial agents. We then evaluate the network models created by researchers in the field, comparing and contrasting aspects of these implementations.

All of the network models discussed here share certain fundamental elements. There is a shock to the system (or, if not a shock, an immediate requirement to settle all outstanding loans) affecting a single financial agent (typically a bank), a group of agents or indeed the whole financial system. The shock forces at least one bank to default on loans from other banks in the system, and this may in turn cause other banks to default on their borrowings if the losses they experience on their exposure to the defaulting banks exceed their capital.

The beginnings of the literature in this field can be traced back to research in the area of bank runs. In a micro-economic analysis, Diamond and Dybvig (1983) explore the idea that bank runs were caused by the self-fulfilling expectations of depositors. This research sparked others to investigate the area of informational contagion and begin to consider the extent that the patterns of interbank exposures would affect the risk of contagion. Allen and Gale (2000) start with a model of contagion similar to that of Diamond and Dybvig but expand it to consider the effect on contagion if four 'regions' (which can be considered analogous with banks) are linked to each other in different ways. Their finding that the pattern of interconnections does indeed have implications for how shocks are propagated resulted in other researchers seeking ways to explore these patterns in a manner more applicable to the complex real-world networks of interbank exposures. The network model of interbank exposures was the tool that many of these researchers adopted, applying techniques from network theory to build models that are increasingly complex, yet sufficiently tractable and comprehensible to be useful in exploring the nature of real-world banking systems and the implications of these findings.

Researchers are interested in exploring the phenomenon of financial contagion due to interbank lending because of the serious economic hazard that bank failures represent. The near-collapse of banking systems across the world in the wake of the credit crunch that began in 2007 (Brunnermeier, 2009) served both as a reminder of the high levels of interconnectedness and interdependence of these systems and a fresh warning of the potential fragility of these systems in times of economic adversity. The coordinated actions of central banks and governments have prevented the ‘doomsday’ scenario of contagious bank failures, but nonetheless there were many banks that had to be nationalized (e.g. Northern Rock, Anglo Irish Bank), sold to stronger banks (e.g. Fortis, Bear Stearns, HBOS) or liquidated (e.g. Lehman Brothers). Although the economic damage could have been far worse, the cost to the world economy of this financial crisis is still enormous, with the IMF estimating that by June 2009, governments had spent $425 billion *on capital injections alone* to support the financial sector (IMF, 2009).

The knowledge that bank failures can be costly for economies and governments is not new, and for this reason the authorities seek in the first instance to avoid such failures by regulation and in the second instance by rescuing banks that are in danger of collapsing. With regard to regulation, the Basel Accords have been a key factor in the development of regulatory frameworks in recent years. The Basel I banking accord (1988) represented an attempt to introduce international standards regarding the levels of bank captialisation, focusing principally on credit risk. The subsequent Basel II accord (2004) addressed regulatory capital not only in the context of credit risk, but also operational risk and market risk (together these represent the first 'pillar' of Basel II). Pillars two and three dealt with issues of how regulators and management supervise and react to the risk assessment, and the levels of disclosure of their calculations and methodology that are required by the banks respectively.

In spite of the work that has been done to improve bank regulation, bank failures still occur. When such a failure is anticipated, the authorities have the choice of allowing nature to take its course or stepping in to rescue the bank. Even if financial contagion is not a likely consequence of allowing a given bank to fail, there are still considerable social costs. Losses to shareholders and creditors of the failed institution are the most obvious example - James (1991) cites losses of 30% of book value for assets held by failed banks, and shareholders can see the value of their holdings wiped out entirely. Seeking to avoid these losses, the temptation for central banks and governments is to 'bail out' banks that are in danger of collapsing, but this introduces the issue of moral hazard. Bank management, acting in the expectation that they will be rescued if they get into trouble, may start taking greater risks than they otherwise would. Depositors, confident that their deposits are safe because the government will intervene, will not scrutinize or monitor the bank to ensure the safety of their money. In addition to the case where financial contagion is not an issue, there is the 'Too Big To Fail' (TBTF) case, where a bank is deemed to be sufficiently large and interconnected that financial contagion is a genuine possibility if it fails, and the social costs of such a failure would be very large (Stern & Feldman, 2004) . In such a case, the authorities have little choice but to bail out the bank. The trend towards consolidation in the banking industry since the 1980s has resulted in a landscape of fewer, larger banks, with the consequence that the number of banks that are in the TBTF category is greater than ever. The moral hazard problem that this presents means that to avoid expensive bailouts or even more expensive failures, better regulation is of critical importance. Better regulation, in turn, depends in part on better models of the banking system, such as the network model - the subject of this paper.

The structure of this paper is as follows: in the next section, we propose and discuss a definition of the financial system network. Section three looks at how current models can be understood in the context of the proposed definition. Section four looks in greater detail at these models and the assumptions that they make, the implementation of the shock to the system, the data used, and the results that are derived from these models. Section five considers the limitations of existing models and section six suggests some avenues for future research. The final section discusses our conclusions.

2. The Financial System Network

In this section we will propose a definition for the financial system network which will be used to examine existing network models. We will explain the definition, and look at each of its elements in greater detail.

2.1 Definition

We define a financial system network $N = (\Omega, L, SP)$ as a set of nodes Ω, representing entities that participate in the financial system, a set of links L between these nodes representing financial relationships, and system parameters SP that determine the characteristics of the system. More formally:

- $\Omega = \{e_i\}$, $i = 1,2,\ldots n$, is the set of nodes that represent the entities that make up the financial system. Each node has both quantitative and qualitative characteristics:
  - $EQN : \Omega \rightarrow \Omega.Quant: e_i \mapsto (eqn_j), j=1,2,\ldots,a$
    - $\Omega.Quant$ is a set of a quantitative characteristics: $(eqn_j), j=1,2,\ldots,a$
  - $EQL : \Omega \rightarrow \Omega.Qual: e_i \mapsto (eql_j), j=1,2,\ldots,b$
    - $\Omega.Qual$ is a set of b qualitative characteristics: $(eql_j), j=1,2,\ldots,b$
- $L \subseteq \Omega^2$ is a set of directed links between entities (l_o), $o = 1,2,\ldots,p$. Each link has both quantitative and qualitative characteristics:
  - $LQN : L \rightarrow L.Quant: l_o \mapsto (lqn_k), k=1,2,\ldots,c$
    - $L.Quant$ is a set of c quantitative characteristics: $(lqn_k), k=1,2,\ldots,c$
  - $LQL : L \rightarrow L.Qual: l_o \mapsto (lql_k), k=1,2,\ldots,d$
    - $L.Qual$ is a set of d qualitative characteristics: $(lql_k), k=1,2,\ldots,d$

- $SP = (sp_1, sp_2, \ldots sp_{s-1}, \ sp_s, \ sp_{s+1}, \ldots, sp_t)$, is a set of t system parameters with s static parameters and t-s dynamic parameters

2.2 Explanation of the Definition

The Financial System Network $N = (\Omega, L, SP)$ is defined by three concepts - the set of nodes Ω, the set of directed links L and the set of system parameters SP.

2.2.1 The set of nodes *Ω*

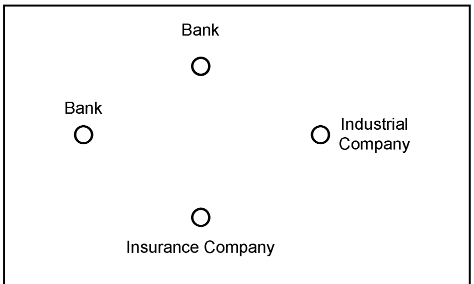

Figure 1: The set of nodes Ω

The first element of this definition is Ω, the set of nodes that represent the entities that participate in the financial system. The total number of nodes in the system is n. Each of these entities may be one of many different types of institution – examples of types of institution would include banks, insurance companies, industrial companies, and central banks. Each of the entities modeled may have quantitative and qualitative characteristics. Quantitative data is data that may be measured on ordinal, interval or ratio scales, while qualitative data is measured on a nominal scale and can only be classed in categories.

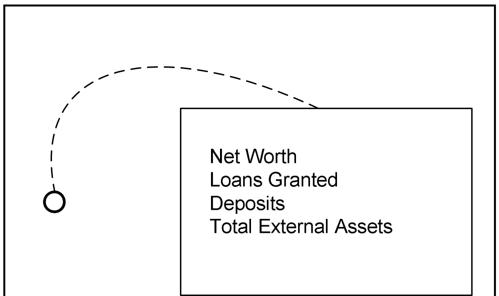

Figure 2: Ω.quant

The quantitative characteristics of these entities are captured in the definition by the set ***Ω.quant***. The function *EQN* is a map from Ω to Ω.Quant such that each element is mapped onto a tuple of <*eqn*>, where a is the total number of quantitative characteristics. Quantitative characteristics would include for example the balance sheet information of the entity, Profit & Loss data and so forth. This is an example of a tuple in a case with four characteristics - Deposits, Net Worth, Total External Assets, Interbank Loan Assets: $EQN(e_1) = (400,150,600,300)$

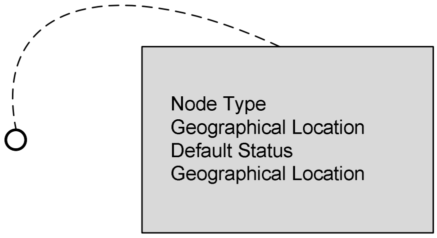

Figure 3: Ω.quant

Ω.qual is the set of qualitative characteristics that pertain to each entity where *b* is the total number of qualitative characteristics. These are mapped from *Ω* to *Ω.Qual* by the function *EQL* so that each entity is mapped onto a tuple of <*eql*>. Examples of qualitative characteristics would include the entity's type, credit rating, default status, geographical location. An example of the tuple *Ω.Qual* with these characteristics would be: EQL(e_1) = (Insurance Company, AAA, Not in default, France)

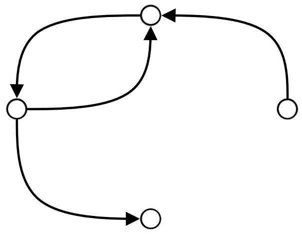
Figure 4: The set of links *L*

2.2.2 The set of links *L*

The second element of the definition is *L*, the set of directed links between entities. *L* is a subset of all the possible combinations of the elements of *Ω* . *p* represents the total number of links in *L.* These links represent the financial relationships which may exist between elements, such as outstanding loans or credit lines offered. Each directed link connects two entities in a specific type of relationship, with the direction of the link indicative of which of the entities is the borrower or the lender in respect of a loan, the extender or the potential recipient of a credit line, and so forth. An entity may be linked to itself. Note that between two entities e_i and e_j several different relationships may exist at the same time – for example one or more loans may have been granted by e_1 to e_2, or by e_2 to e_1, or both, each loan with different characteristics.

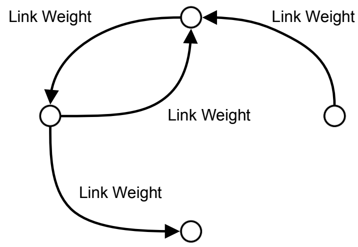

Figure 5: L.quant

L.quant is the set of quantitative characteristics of these links, where *LQN* is a function that maps each link l_{ij} to a tuple <*lqn*> and *c* is the total number of quantitative characteristics. The monetary value of a loan or credit line is probably the most important example of a quantitative characteristic of a link. Such a tuple with the single characteristic 'link weight' would be: LQN(l_{12}) = (2500)

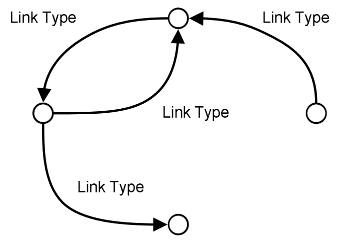

Figure 6: L.qual

L.qual is the set of qualitative characteristics of the links, where *d* is the total number of qualitative characteristics and the function *LQL* maps each link to a tuple *<lql>*. Examples of the qualitative characteristics of links would include the term of a loan, the type of collateral involved and so forth. An example of such a tuple with the characteristics 'type of collateral' and 'term of loan' would be: LQL(l_{12}) = (Bond, 3 months)

2.2.3 The system parameters *SP*

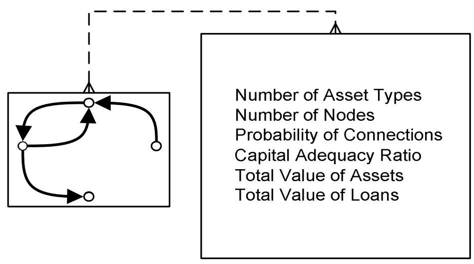

Figure 7: System Parameters, SP

The final part of the definition pertains to the system parameters *SP*. There are a total of *t* system parameters. These system parameters define the characteristics of entities in the model or are used by the algorithm which is applied to the model to explore the consequences of credit default scenarios. The definition distinguishes between two different types of system parameters – dynamic parameters and static parameters. There are *s* static parameters, and the remainder of the parameters, *t-s*, are dynamic parameters. Dynamic parameters are parameters whose initial values will change as the algorithm models the effects of a default scenario. Examples of dynamic parameters could include the total value of assets in the model, or the percentage of total assets represented by the total net worth of the entities. Static parameters are the second type of system parameter. As the name implies, the initial values of these parameters do not change as the algorithm changes the state of the model. These will typically be parameters that are exogenous to the entities that are modeled – for example the capital adequacy ratio, or the demand function for external assets. A set of system parameters with values for two static parameters – capital adequacy ratio and number of asset types - and two dynamic parameters - percentage of total assets represented by total net worth and percentage of total assets represented by interbank assets – would be: SP = (.08, 5, 8%, 30%)

2.2.4 The Contagion Algorithm

A contagion algorithm *C* defines the consequences of a default of a node or nodes in terms of the financial system network. More formally:

$C: (\Omega, L, SP) \rightarrow (\Omega', L', SP')$

- *C* is a function that maps Ω to Ω', *L* to L' and *SP* to SP'.

The algorithm C takes as an input the initial state of the network as composed of *Ω, L,* and *SP* and transforms this into a new state. The default of a node will result in changes to the characteristics of some members of *Ω* and *L.* For example, the default of a node will result in changes to the balance sheet of nodes that have lent to it, which is a characteristic of members of *Ω*. It will affect the value of the loans to the defaulting node, which is a characteristic of members of *L*. It may also have an effect on the dynamic parameters of *SP*, such as the percentage of total assets represented by interbank loans. The static parameters of *SP* such as the capital adequacy ratio will, by definition, not be changed. Hence, the original state of the network *(Ω, L, SP)* is updated to a new state, (Ω', L', SP').

3. The Definition Applied to Existing Models

We will now explore individually each of the elements of the FSN definition *N = (Ω, L, SP)* in turn as they are realized in each model implementation. We will name each network model after the first author named in the paper where the model in question is described.

We will also discuss as a type a kind of network model that has – with some variations – been widely used in empirical studies of the risk of contagious default in national banking systems. Due to their fundamental similarity, we will consider these models as a single type which we will refer to as the 'Matrix model'. At the core of these models is the creation of an *N* x *N* matrix where *N* is the number of agents participating in the system and each of the elements of the matrix, x_{ij}, represents the exposure of bank *i* to bank *j*.

$$X = \begin{bmatrix} 0 & \cdots & x_{1j} & \cdots & x_{1N} \\ \vdots & \ddots & \vdots & \iddots & \vdots \\ x_{i1} & \cdots & 0 & \cdots & x_{iN} \\ \vdots & \iddots & \vdots & \ddots & \vdots \\ x_{N1} & \cdots & x_{Nj} & \cdots & 0 \end{bmatrix}$$

The sum of each **row** represents the total exposure (usually loans) of agent *i* to all other agents the system, and the sum of each **column** represents the total liabilities (usually borrowings) of agent *j* to all other agents in the system. These models use empirical data sources to build this matrix, and typically represent a specific national banking system.

Upper (2007) provides an authoritative overview of the methodology of these models. Sheldon & Maurer (1998), Furfine (1999), Wells (2002), Upper & Worms (2004), Lelyveld & Liedorp (2004), Lublóy (2005), Amundsen & Arnt (2005), Degryse & Nguyen (2007), Krznar (2009) and Mistrulli (2010) represent a selection of the studies that rely on this type of model.

3.1 The Set of Nodes *Ω*

Ω is defined in the FSN as a set of nodes that represent the set of entities that participate in the financial system. It should be noted that the source and sink nodes described in these models can be regarded as a modeling convenience, as it is possible to model losses suffered by agents in the system and the asset holdings of those agents using only the quantitative characteristics of the agent nodes.

Eboli (2007) and Nier et al. (2008) specify three distinct types of nodes – source nodes, nodes that represent 'financial intermediaries', and a sink node. In Eboli's model, source nodes are used to represent the external assets belonging to the agents in the system, the financial intermediaries, while the single sink node is used to model losses to the assets of the agents by acting as the point where the money exits the system. The external assets are defined as assets that are not issued by agents in the system, which distinguishes them from interbank loan assets. Each asset type must be owned by at least one agent in the system – in terms of the network model, this means it will appear on the balance sheet of at least one agent. The source node in Nier is similar to the source node in Eboli, but Nier has one only asset type, while Eboli allows more than one.

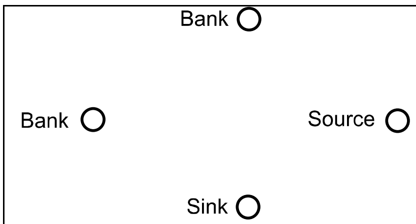

Figure 8: Nodes in Eboli, Nier

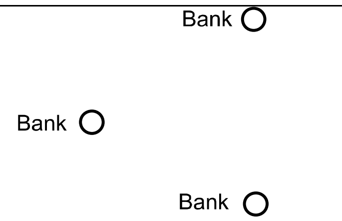

Figure 9: Nodes in Müller, Matrix model, Elsinger, Cifuentes.

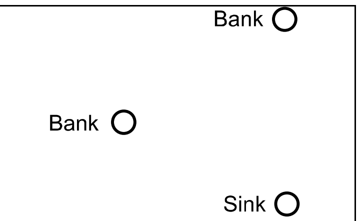

Figure 10: Nodes in Eisenberg

Canedo and Jaramillo (2009) use a distinctly different type of network to visualize their model. A source node represents the initial shock to the system, and it is linked to a set of agent nodes representing each of the banks. This is what the authors refer to as the 'shock phase' of the model. Each of these agent nodes is in turn connected to another set of nodes representing the same set of banks, but this time the nodes represent the banks during the 'contagion phase' of the model. The contagion phase may contain many rounds of defaults, and in each round each bank is represented by a single node. The nodes representing the final contagion stage are in turn linked to a sink node, where the losses from the banks that have failed at the conclusion of the modeling exercise flow.

The model described by Eisenberg & Noe (2001) is simpler than those of Nier and Eboli, featuring only two types of node. One type represents 'economic entities' and one type represents a sink. The sink node is optional in the Eisenberg model and is introduced to allow the modeling of loss-making companies. Eisenberg's model requires that each node has a positive 'operating cash-flow'. By allowing a node to have a positive cash-flow but an even greater 'operating cost' (a liability to this sink node) the authors argue that loss-making companies can be successfully modeled with no loss of generality in the model.

The Matrix model, the model of Müller (2006) and the model of Elsinger, Lehar, & Summer, (2006) make no mention of any source or sink nodes, and feature only one type of node which is used to represent agents (banks). This is also the case with the Cifuentes, Ferrucci, & Shin (2005) and Georg & Poschmann (2010) models; however, agent nodes are the minimum that would be required to visualize this or any other network model.

3.2 *Ω*.quant

Ω.quant is the set of quantitative characteristics of the nodes in *Ω*. Both Eisenberg and the Matrix model use a single quantitative characteristic of the set of nodes – net worth. Eisenberg calls this 'cash inflow' and describes it as a quantity of money that is paid to the node due to its business activities, and as mentioned previously it has to have a positive value. This ensures that the modeled business has some value that it can pass on to any creditors that it might have in the system. If a node's cash inflow plus the money it recovers from debtors is less than what it owes to creditors, then that node is insolvent. This quantitative characteristic is present in every model – usually called 'net worth' - and can be considered a fundamental requirement of network models. In the case of the Matrix model, some implementations use additional quantitative characteristics. A notable example is Degryse and Nguyen (2007), who introduce several new characteristics required to describe in greater detail the types of assets held by each agent in an extension where they derive an endogenous figure for the LGD of each bankrupt bank.

Elsinger's model can be considered to consist of two parts. The first part is a network model based on Eisenberg, while the second part is a portfolio model that is used to simulate the asset portfolio holdings of each bank (excluding interbank loans) at two points in time. At t = 0 the portfolio of each bank is observed. These holdings are subjected to simulated market and credit risk and a new value is derived for each bank's portfolio. This new value becomes the 'net worth' value for the first part of the model, and the model then proceeds in the same way as Eisenberg. For the purposes of this paper, we will focus principally on the network model element of Elsinger's model, treating the output of the portfolio model element as a quantitative characteristic of each bank.

Eboli's network represents additional agent balance sheet information with these quantitative characteristics. The balance sheet information of each agent consists of the value of external assets held, the sum of loans granted by that agent to the other agents in the system, the sum of loans received from other agents in the system, and the net worth of the agent.

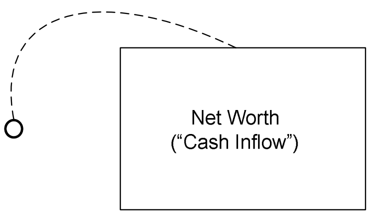

Figure 11: Ω.quant Eisenberg, Elsinger, Matrix Model

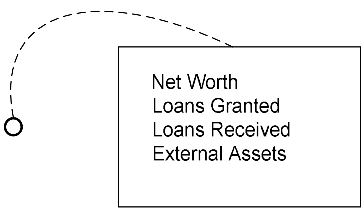

Figure 12: Ω.quant Eboli

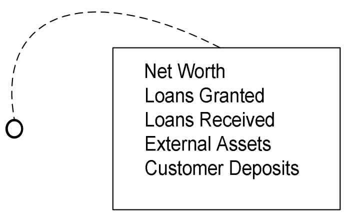

Figure 13: Ω.quant Nier

Nier uses the same set of quantitative characteristics as Eboli, with the addition of customer deposits. This creates the possibility of exploring the degree to which the losses caused by a default affect the different stakeholders in the bank system – shareholders (net worth), industry creditors (interbank loans) and retail creditors (customer deposits). Later we will briefly discuss whether the distinction between these two types of creditor is justifiable in practice. Cifuentes is similar to Nier, but divides external assets into 'liquid' and 'illiquid assets', and adds 'cash'. These additional characteristics are required by Cifuentes to model market risk.

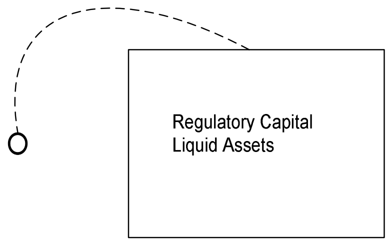

Figure 14: Ω.quant Müller

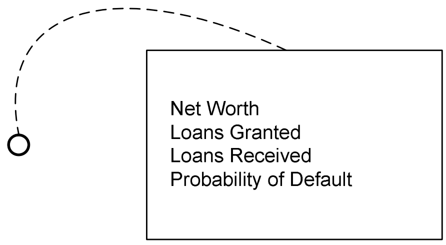

Figure 15: Ω.quant Canedo

Müller's model has nodes with two quantitative characteristics – 'liquid assets' and 'regulatory capital'. 'Regulatory capital' is equivalent to the 'net worth' characteristic in other models. The addition of 'liquid assets' is required because Müller's model makes a distinction between insolvency and illiquidity. If a bank's capital exceeds the loss on its claims on other banks, it is solvent. However, it may be illiquid if the money coming in from its debtors plus its liquid assets are insufficient to repay what it owes to its creditors. This plays an important role in Müller's model as illiquid but solvent banks are allowed draw on available credit lines to avoid defaulting on debts.

Canedo and Jaramillo introduce a new quantitative characteristic, the probability that each node will default. This is used to determine the likelihood of different loss scenarios and, in the implementation, relies on estimates calculated by the Mexican central bank. The value of loans granted is used in conjunction with the net worth figure (or 'threshold') to determine which banks are overexposed by lending more to other agents than their net worth.

Georg's model requires more quantitative characteristics than any other model examined here in its basic form; due to its complexity, characteristics not seen in other papers such as 'loans from central bank', 'investment maturity', 'investment value' are added.

3.3 *Ω*.qual

Ω.qual is the set of *qualitative characteristics of the nodes* in *Ω*. There are fewer qualitative than quantitative characteristics in the models examined here. Nier, Eboli, Eisenberg, Elsinger, and

Cifuentes and the Matrix model each feature ‘default status’ and ‘node type’ as qualitative characteristics of nodes.

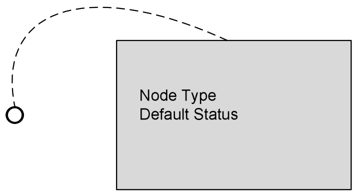

Figure 16: Ω.qual Nier, Eboli, Eisenberg, Elsinger, Cifuentes, Matrix model.

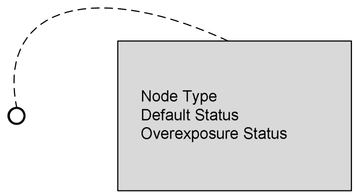

Figure 17: Ω.qual Canedo

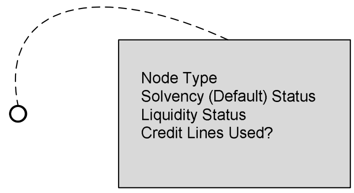

Figure 18: Ω.qual Müller

Canedo adds ‘overexposure status’ as a qualitative characteristic. As mentioned previously, this reflects whether or not a given bank has lent more than its net worth, or ‘threshold’ amount, making it vulnerable to insolvency. If it is not overexposed, it can be ignored as a candidate for defaulting when the contagion algorithm is run.

The Matrix model in its simplest version requires only ‘default status’ and ‘node type’ characteristics, but others may be added as required, depending on the goals of the model. An interesting example is the model developed by Mistrulli, which allows banks belonging to the same parent group share capital. A bank that becomes insolvent can receive a capital injection from other banks owned by the same parent group, preventing its default. In order to identify which group each bank belongs to, the model requires the introduction of an ‘affiliation’ characteristic for each node.

Müller has a wider set of qualitative characteristics. The role of the default status in the other models is taken by ‘solvency status’ in Müller’s model – this reflects whether the bank is solvent or insolvent. The conditions for insolvency in Müller are similar to those for default in the other models, but in Müller it is used to make clear the distinction between insolvency and illiquidity. Illiquidity is another of the qualitative characteristics in Müller’s model, so that it is possible to model a situation where a bank or banks are solvent but illiquid. In such a case, a bank may call on its available credit lines from other banks to obtain money to meet its obligations. The fourth qualitative characteristic in Müller’s model denotes whether a bank has used its credit lines yet or not, and is referred to here as ‘Credit Lines Used?’

3.4 The Set of Links *L*

L is defined in the FSN as a set of directed links between the nodes. No node is allowed to be linked to itself in any of these models. However, such links are not forbidden under our definition and they may serve a role in future models.

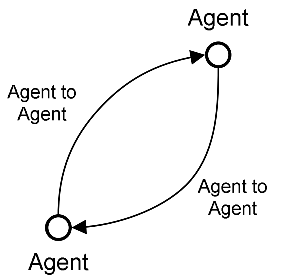

Figure 19: Links in Müller, Elsinger, Matrix model.

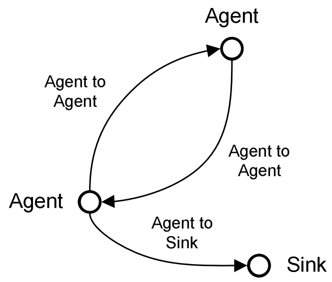

Figure 20: Links in Eisenberg

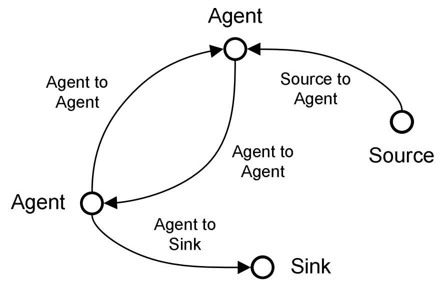

Figure 21: Links in Eboli, Nier, Canedo

The existence of various types of link is a function of the types of node present in a model. Hence, Müller and Elsinger, models with only agent nodes (banks) have only agent-to-agent links (LA). Eisenberg's model introduces a sink, which enables the existence of agent-to-sink links (LSK). Eboli, Nier and Canedo add source nodes, allowing the creation of source-to-agent links (LS). Note that source-to-sink links do not exist in any of the models presented here.

Müller's links between agents represent either loans or credit lines. In the case of all the other models examined, LA represents only interbank loans. Matrix models feature only links between agents. Eisenberg's model may contain, in addition to LA, a set of links between agents and a sink node (LSK). As discussed earlier, the sink is an optional feature of Eisenberg's model which may be introduced to facilitate the modeling of loss-making companies. If loss-making companies are not modeled, a sink node is unnecessary.

Eboli and Nier use directed links to represent the possession of different asset types by the agents (LS), the loans that each agent grants to other agents (LA), and the flow of losses to the sink node (LS). Although there is no explicit requirement in Nier's model for external assets to be modeled as 'source nodes' as in the case of Eboli, we do so here to make clear that Nier is largely based on Eboli's model. Canedo also features three types of links – those from the source to the agent, those between the agents, and those from the agents to the sink. Unlike Eboli however, the links from the source to the agents represent the likelihood that the agents will default due to the initial shock.

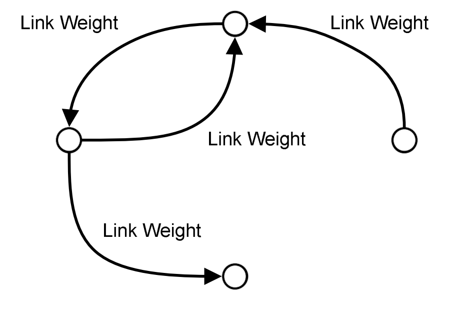

Figure 22: L.quant for all models

3.5 L.quant

L.quant is the set of quantitative characteristics of the links in L. The only quantitative characteristic of the links in the models presented here is their weight. In all models, link weights are used to indicate the value of loans in LA. Nier and Eboli also use these weights to determine the value of external assets held by the economic agents, and Nier, Eboli, Eisenberg use link weights to represent how much value is exiting the system by means of sink nodes.

Canedo uses the weights of links to express two different things; the value of loans and losses exiting the system, and as a measure of probability. When the link is between the agent nodes representing the final contagion stage and the sink node, the links represent the losses from the failed banks flowing out of the system. When the link exists between the source node and the agent nodes, it expresses the probability of default for each agent node and will have a value between zero and one.

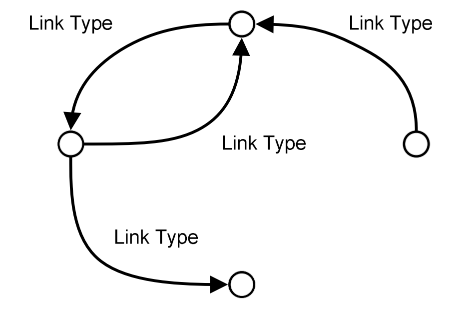

Figure 23: L.qual for all models

3.6 L.qual

L.qual is the set of qualitative characteristics of the links in L. The models we examine here have a single qualitative characteristic – the 'type' of each link. Depending on which types of nodes it is linking, a link may represent ownership of an asset type by and agent, a loan between two agents, or a loss exiting the system via the sink.

In Müller's model, there is an important distinction between two different types of link between agents. A link between two banks may be a loan or a credit line – essentially a *potential* loan that has not been drawn down yet. In the execution of the model, we will know whether a given credit line has been used by checking the 'Credit Lines Used Status' characteristic of the bank that credit line has been extended to, as Müller's model requires that a bank draws down all of its credit in one go in the event that it becomes illiquid at which point this characteristic of the node will be updated.

3.7 System Parameters (SP)

The system parameters of a model are those characteristics that cannot be ascribed to the individual nodes in *Ω* or the links in *L*, such as the number of banks in the model. Often these will be ratios of aggregate values in the model. For example, in the case of Nier, the system parameters are the net worth as a percentage of total assets, the percentage of total assets represented by interbank loans, the asset elasticity of demand, the number of banks, the probability of any two banks being connected, and the total value of external assets.

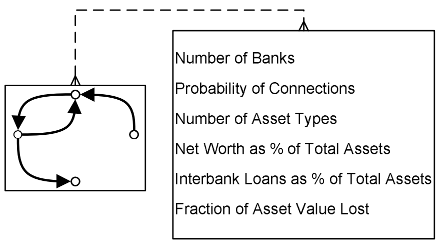

Figure 24: SP for Eboli

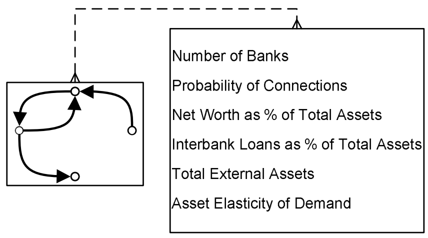

Figure 25: SP for Nier

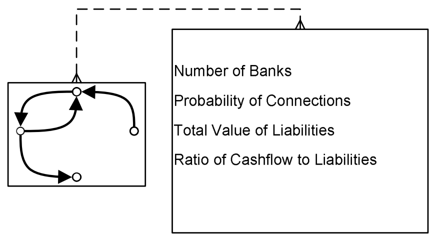

Figure 25: SP for Eisenberg

Any possible realization of Nier's model can be described by these five parameters. In the extension of the model where asset price effects are modeled, an additional parameter is required to determine the price elasticity of demand for the assets sold by defaulting banks. The similarity between Eboli's and Nier's models is clear when the sets of System Parameters are compared. Eboli's model requires a single additional parameter, the number of asset types that exist, as Eboli allows one or more, while Nier allows only one. Conversely, Nier requires an Asset Elasticity of Demand parameter for modeling the affects of insolvent banks selling assets. Cifuentes model requires a set of parameters similar to Nier, but with the addition of a capital adequacy ratio, and a liquidity ratio to determine the ratio of each banks liquid and illiquid assets.

Eisenberg's model has fewer parameters as this model does not feature the complication of external assets that we see in the other models. However, similar basic parameters are required to establish the relationship between the amount of money that the agents hold ('operating cashflow' – i.e. net worth) and the money that they owe on a system level, and the degree of interconnectedness of the network.

Elsinger, Müller and Canedo's models do not require parameters as they are populated with empirical data. The number of banks, the ratios of different asset types, the probabilities of banks being connected and so forth are derived from an empirical datasets. However, in these empirical models, such parameters can be 'read' from the model as an output rather than supplied as an input. In isolation, such parameters may not be very informative but comparisons between different national banking systems may contain some interesting information and are perhaps an avenue for future research.

Similarly, the Matrix models are populated with empirical data. However, unlike the models described above, the typical Matrix model does not derive an endogenous LGD and therefore the LGD used in each 'run' of the model is a system parameter that is predetermined by the researcher in question. Other system parameters may exist in addition to LGD; for example Lublóy extends the Matrix model to explore the consequences of applying Hungarian prudential regulatory standards. In this extension, a bank is considered to have failed if its capital adequacy ratio falls below 4%, requiring the introduction of 'capital adequacy ratio' as a new system parameter.

As with the quantitative characteristics of the nodes, Georg requires the introduction of many new system parameters not seen in other models. Georg indicates that there are 18 parameters, including such innovations as the interbank interest rate, the probability that an investment loan will be repaid, and the banks' risk aversion parameter. The risk aversion parameter of an individual bank, under our definition, is a quantitative characteristic of that node, but Georg specifies a random number for each bank between a given range, and this range falls under our definition of a system parameter.

3.8 The Contagion Algorithm *C*

Upper (2007) explores the use of counterfactual models of financial contagion, where models of actual banking systems are created and tested to explore the likelihood and consequences of financial contagion. The models that Upper examines are all network models of financial system networks under the definition outlined in this paper; consequently, the insights found in that study are also of great usefulness in this more general study.

Upper characterizes the basic algorithm used in these counterfactual models as a sequential algorithm composed of three steps. The same three steps form the basis of the algorithms used in the network models we are concerned with here. Here Upper's characterization has been slightly modified for the purposes of examining the algorithms found in these network models:

1. The initial failure: A bank Ω_i (or banks) fails by assumption.

2. Propagation: A bank Ω_j will fail if its exposure (loans) to bank Ω_i multiplied by a percentage 'loss given default' (LGD) exceeds its capital. A second round of contagion may occur if a bank Ω_k whose exposure to banks Ω_i and Ω_j, multiplied by the LGD, exceeds its capital. This step may be iterated.

3. Stopping condition: A stopping condition will be met.

This generic description is useful in understanding the basics of the algorithms used here, but each model varies in the implementation of each step, and in some cases the first step is omitted completely. To explore the algorithms used in the models under consideration here, we will consider each step in turn and discuss where a model expands on or deviates from that step.

A key distinction to note between the models discussed here is the derivation of the LGD figure used. The Matrix model typically uses an exogenous LGD that is determined as a system parameter, as does Canedo's model. Every loan exposure to a failed bank will suffer the same loss in a given 'run' of the model. Given that the LGD figure is an estimate, researchers will usually run the model using a range of LGDs. The other models examined here use an endogenously derived LGD, which means that banks will lose different amounts of their exposures to insolvent banks depending on the degree of insolvency involved – a bank may become slightly insolvent, and pay back most of the money that has been lent to it; or it may become extremely insolvent, paying back little or nothing of what has been lent. The implications of how the LGD is derived will be discussed in more detail later.

3.8.1 Initial Failure

Eboli and Nier model this step by reducing the value of assets held by a particular bank or banks to a level where they become insolvent and have to be liquidated. Nier specifically models only failures of an individual bank, whereas Eboli models scenarios where either a single bank or several banks default due to the initial shock.

The Matrix models typically begin by simulating the failure of each bank in turn by assumption, with no mechanism required to cause this initial failure (e.g. Wells, Upper & Worms, Van Lelyveld & Liedorp, Degryse & Nguyen). They may also simulate particular scenarios; for example, by starting the simulation with the assumed failure of the largest debtor in the system (Amundsen & Arnt), or the failure of more than one bank at once (Lublóy).

Müller, Eisenberg and Elsinger deviate significantly from this first algorithmic step in that there is no failure of a bank 'by assumption'. Rather, all of the nodes attempt to repay their liabilities immediately, and failures will occur only if there is some node or nodes that do not have enough money coming in to provide the amount of money required to satisfy their liabilities. In the case of Elsinger, this attempt to settle the system occurs after each bank's portfolio holdings have been revalued by the portfolio model. It may be argued therefore that these models begin without this 'initial failure' first step.

Canedo takes a different approach. Rather than modeling a single default event, Canedo seeks to derive a probability distribution of losses for the banking system by calculating not only the losses that arise given the default of a bank or a combination of banks, but also the probability for the scale of losses. At the initial failure stage, every possible combination of banks in the system fails by assumption.

Cifuentes assumes a single bank is liquidated with a given LGD.

3.8.2 Propagation

Eboli's algorithm features an 'activation function' that determines the share of each node's net worth that has been lost due to the initial shock. This function takes the form:

$$\beta_i(\lambda_i) = \min(\lambda_i/v_i, 1) \tag{1}$$

where β_i is the share of net worth lost by node *i*, λ_i is the loss experienced by node *i*, and v_i is the net worth of node *i*. If the loss a node receives exceeds that node's net worth, then the node defaults and βi is set to 1. The LGD for the initial failure is calculated with what Eboli calls the 'insolvency function':

$$b_i(\lambda_i) = \max(0, \lambda_i - v_i/l_i) \tag{2}$$

Here, b_i is the fraction of node *i*'s liabilities that cannot be recovered by liquidating *i* and l_i represents the interbank liabilities of *i*. If a node's net worth v_i is greater than the loss, the value of b_i is zero. If the loss λ_i is greater than net worth, then we get a positive value for b_i.

The first propagation step is to calculate each bank's total losses by adding their losses from other banks (which will be zero for the first bank to fail) and their losses due to the fall in the prices of assets that they hold. Having calculated this for each bank, each bank is checked with the activation function to determine whether it is still solvent or not, and with the insolvency function to determine how much of its creditors' money has been lost if it is in default. Defaulting banks are added to a new set, and this set is checked for cycles. If there is a cycle, there are two further checks. If the flow of losses that reaches a cycle is the largest possible (equal to the total amount of external assets held by the members of the cycle) and there are no debts owed to members outside the cycle, then it is possible to set b_i to '1' for all the members of this cycle. This simplifies the calculations there is no need to calculate a loss rate for any of these nodes if there are no losses to be passed out of the cycle to other nodes.

Nier's propagation is similar to Eboli's. Each bank in turn is checked to determine if $\lambda_i > v_i$. If it is, then the node *i* defaults. Given this default, if $(\lambda_i - v_i) < l_i$ where l_i represents *i*'s interbank liabilities, the losses are borne by bank *i*'s interbank creditors. If however $(\lambda_i - v_i) > l_i$ then the losses spill over into the customer deposits held by *i*, d_i. Nier's algorithm is simplified in that the values of all loans in the model are the same, so all of a defaulting bank's creditors will receive the same loss. Therefore the loss that *j*, a creditor bank of bank *i*, will receive is easily calculated as: $s_j\ [(\lambda_i - v_i) / k]$ where k = the number of creditor banks of *i*. This is only the case if $(\lambda_i - v_i) < l_i$ because if $(\lambda_i - v_i) > l_i$ then the loss to creditors is total. If the loss that bank *j* receives is greater than its net worth v_j, then that bank also defaults and the algorithm will check each bank again, passing this loss onto the creditors of *j*.

Eisenberg's algorithm, called by the authors the 'fictitious default algorithm', is probably the simplest of the group under discussion here. First a matrix is created that expresses each node's nominal liability to other nodes in the system. Then a 'relative liabilities matrix' *Π* is created where each of a node's loans is expressed in terms of the fraction of that node's total liabilities that loan represents. This matrix will be used to calculate what proportion of a loan to a defaulting node a creditor will receive when the system is being settled – for example, if bank *i* and bank *j* each loan \$100 to bank *k*, each of these loans represent half of bank *k*'s total liabilities. If bank *k* defaults and is liquidated with a total value of $150, bank *i* and bank *j* will each receive half of this $150 in settlement of their loans.

In settling the system, the algorithm determines what each node must pay to meet its obligations to settle liabilities to other nodes – a value which we will call *p*. Assuming that all nodes pay the full amount of *p*, the algorithm checks if any nodes default using the formula:

$$\sum_{j=1}^{n} \Pi_{ij}^{T}\, p_j + v_i - p_i \tag{3}$$

Here, Π^T is the transpose of the relative liability matrix, so the term on the left reflects the payments received by node *i* from its debtor nodes plus its cashflow (analogous to net worth), whereas the $-p_i$ on the right are the payments that node *i* owes to its creditors. If this equation evaluates to less than zero, the node is in default.

In the event that a node defaults, the algorithm then attempts to settle the system assuming only the defaults that were detected in the first round occur, with the payments from those defaulted nodes to their creditors reduced to the maximum that they are able to pay. These reduced payments are then distributed proportionately based on the relative liability matrix to the node's creditors. If these reduced payments cause any of these creditors to default in the second settlement round, a third round of settlements is carried out taking into account the reduced payments from these new defaults, and so forth. Elsinger's network model works in the same way as Eisenberg; the key difference, as stated previously, is that the v_i for each bank is updated by the portfolio model before the system is settled, allowing the simulation of market and credit risk.

Müller's algorithm is also an extension of Eisenberg's. Müller's algorithm makes a distinction between insolvency and illiquidity, but she bases this algorithm on the fictitious default algorithm described above where each iteration of the algorithm assumes that only the banks that became insolvent or illiquid in the previous round do not make full payments of their debts. To model insolvency and illiquidity, Müller's algorithm introduces two tests for each bank. The bank is tested for solvency by checking whether the loss on its interbank assets exceeds its capital. Secondly, the bank is tested for liquidity by checking whether it can completely repay its liabilities, using the following formula:

$$f_i \equiv \sum_{j=1}^{n} \Pi_{ij}^T \, p_j + k_i \geq l_i \tag{4}$$

Once again, the summation term refers to the payments received from debtors, with Π^T representing the transpose of a relative liabilities matrix. The amount of liquid assets available to the bank *i* is represented by k_i, and l_i represents the interbank liabilities of the bank. Therefore, f_i is a value that represents the total amount that the bank *i* is capable of paying to its creditors. The similarity between this formula and Eisenberg's formula above is clear, but there is an important distinction in that Eisenberg's formula adds the value for cashflow (net worth) to the incoming payments from other nodes, whereas Müller adds the value for liquid assets, which is quite distinct from net worth in her model.

In order to model the existence of credit lines, Müller adds a further formula to calculate a value for each bank's liquid assets, *k*. In brief, this formula adds to the bank's existing liquid assets the newly raised credit that a bank receives from those banks it has credit lines with and subtracts the credit that it has offered to other banks where those banks have taken up the credit. The algorithm takes the full amount of credit offered even if only a part of the credit line is required to restore the bank to liquidity. As with Eisenberg's model, the algorithm begins by assuming that every payment is made in full. If a default or defaults occur under this assumption, the algorithm runs for a second iteration, reducing the payments of those banks that are insolvent or are solvent but illiquid to the maximum that they can pay, but assuming that all other payments are made in full. Banks that are solvent but illiquid have the chance to seek more liquidity from those banks that they hold credit lines with, and if their creditors are able to supply them with this liquidity their payments to their interbank creditors will be increased. If there are further defaults in the second round, the algorithm will attempt to settle the system again assuming that only the banks that defaulted or became illiquid in the first two rounds make reduced payments, and so on.

The propagation stage of the typical Matrix model is quite simple by comparison to some of the other models examined here, principally because the LGD figure is not calculated within the model, but is instead supplied as a system parameter. Adapting Upper's summary of the typical propagation stage, any bank *j* will fail if its exposure to a bank *i*, x_{ji}, multiplied by the supplied LGD, exceeds its capital v_j. A second round of contagion may occur if there is a bank *k* for whom $LGD(x_{ki} + x_{kj}) > v_k$, and so forth for third and subsequent possible contagion rounds.

Canedo's propagation is straightforward; banks that have an exposure to a failed bank that exceeds their net worth (or 'threshold' value) will fail. In the event of such a failure, the entire value of any loans to this set of newly failed bank is considered lost. Any banks whose loans to this set of banks exceed their net worth will fail in turn. The innovation in Canedo's model is that every possible combination of failures is modeled. By starting with a given probability that any particular bank will default (using data from the central bank) and treating every initial failure as independent, the model produces a loss distribution that outlines both the likelihood and scale of contagious failures and losses to the whole system.

Cifuentes' propagation is complicated by the fact that there are two channels of contagion at work; a bank can become insolvent due to losses sustained on interbank loans when a counterparty defaults, or it can become insolvent due to the falling price of illiquid asset holdings. Cifuentes' model features a market for illiquid assets, where the sale of the assets of insolvent banks can drive down the price of the illiquid assets held by its peers. Banks that do not meet the capital adequacy ratio are also forced to sell assets in this market until they can do so. If they sell all of their liquid and illiquid assets (excluding interbank loans) but still do not meet the capital adequacy ratio then they are technically insolvent and are liquidated. Thus an organization may be solvent in the sense that it has positive net worth, but be technically insolvent as it falls below the capital adequacy ratio determined by regulatory authorities. The level of net worth required to meet this capital adequacy ratio is given by:

$$r^* \sum_{i=1}^{n} x_{ij} \qquad (5)$$

where r^* is the capital adequacy ratio and x_{ij} represent loans from bank *i* to bank *j*.

Cifuentes propagation algorithm starts by determining which banks do not meet the capital adequacy ratio or are insolvent, and then calculates the amount of the illiquid asset that is to be sold in the market as a consequence. A demand function is then used to calculate a new price for the illiquid asset and the algorithm checks each bank to determine if it is solvent given the new price of its illiquid asset holdings. If any bank is found to be insolvent or does not meet the capital adequacy ratio then the algorithm starts another iteration.

Georg's propagation algorithm is quite lengthy and complex, as this model introduces important new elements that allow the modeling of both central bank interactions and the impact of the real economy on the banking system. A detailed examination of this propagation algorithm is beyond the scope of this paper, but we will make some observations about its key points. Firstly, there are three ways that banks may become insolvent in this model; large deposit withdrawals, losses on investments, and the failure of bank loan counterparties. By modeling deposit holdings and investment returns stochastically, this model allows the real economy to cause the weakening or failure of individual banks. A second important point to note is that there is a time dimension in this model, as each iteration of the algorithm represents a new time period when investment returns and deposit holdings are updated, and the banks modify their portfolio for the coming year. Finally, a key innovation sees the introduction of a central bank which acts as lender of last resort to banks which cannot obtain liquidity from its bank peers, if the bank has sufficient assets to secure the central bank loan against.

3.8.3 Stopping condition

In Eboli's model, the total loss to banks in the system is caused by a drop in the value of assets held by those banks. At the end of each iteration of the propagation algorithm, the final check that is carried out is a comparison between the flow of losses to the sink and the value of the initial shock to the external assets of all the banks. When the total amount of losses that have been sent to the sink node is equal to the size of the initial shock to the system, all losses are accounted for and the algorithm will terminate. Nier's stopping condition is simpler, in that the algorithm keeps iterating until there is a round free of new defaults, at which point it stops.

The Matrix model algorithm and Eisenberg's algorithm, like Nier's, stop when there is a round with no new defaults. Eisenberg makes the point that when using the 'fictitious default' algorithm, because there must be a default every round or the algorithm will stop, the maximum number of rounds is the same as the number of banks – an observation that we can extend to any algorithm that uses a stopping condition of this type. Müller makes no mention of a stopping condition in her algorithm but it is probable that is the same as the condition used by Eisenberg. Cifuentes' algorithm terminates when there are no new insolvencies, and no bank that does not meet the capital adequacy ratio.

Canedo observes that given a set of initial failures, the consequent 'contagion path' of bank failures is entirely deterministic. Only overexposed banks can fail, so the contagion process will stop when all such overexposed banks are in default, or when a round of defaults occurs where there is no default that pushes an overexposed bank into insolvency.

Georg's model runs a number of 'update steps' (i.e. iterations of the algorithm) specified as a system parameter and stops when this number of update steps has been completed.

4. Model Details and Results.

In this section we will examine the key assumptions that these models make, the implementation of the shock that starts the contagion process, the types of data used in each model, and the results that each model implementation produces.

4.1 Assumptions

In order to model a financial network based on loan exposures, certain assumptions must be made about the behaviour of the financial agents and the world they operate in. The assumptions implicit in each model vary, but some key assumptions are detailed in Table 1. Krznar (2009) details nine common assumptions which are adapted here and supplemented with several other fundamental assumptions. We will briefly explain what each of these assumptions refers to and look at some of the notable exceptions among the models examined here.

Table 1: Modelling assumptions

	Assumption	Notable Exceptions
1	Debt has priority over equity	None
2	Liabilities are limited	None
3	Proportional repayments to creditors in default	Canedo
4	Moral hazard is not modeled	None
5	Contagion is isolated from macroeconmic shocks	Elsinger
6	No deposit flight due to defaults	None
7	Portfolios and asset prices remain constant	Elsinger, Cifuentes, Nier, Georg
8	No seniority of creditors in default	Nier
9	Collateralisation of claims not modeled	None
10	Bank failures are unexpected	None
11	Central bank does not rescue failing banks	None
12	Interbank claims not backed by government guarantees	Upper

13	Failing banks cannot be recapitalised	Mistrulli

Some of these assumptions are based on basic real-world laws, for example the first two; the first assumption states that debt has priority over equity so that an agent that cannot repay all of its debts will be liquidated, and the second assumption is that limited liability is in operation so an agent will only repay what it can, limited to whatever equity and assets it holds. Assumption three is simply that if an agent defaults, each creditor will be paid back in proportion to how much it lent the failed agent. For example, if a certain bank is owed a tenth of the failed bank's liabilities, then that bank will receive a tenth of the value that is realised when the failed bank is liquidated. An exception to this assumption is Canedo's model, where creditors lose 100% of whatever they lent to a failed bank. This may be considered a realistic representation of the short term scenario that this model focuses on, but in reality banks will expect to recover a large part of their exposure to an insolvent bank in the medium and long terms, so this assumption will lead to a considerable exaggeration of the probability of contagion.

The fourth assumption concerns moral hazard; nodes always repay everything they owe or as much as they can. This assumption greatly simplifies the modelling process. If models have to take account of moral hazard, not only would the behaviour of the management of the agents have to be modelled, but many new assumptions would be required about their behaviour.

Typically the models start from a position where a single bank or a combination of banks fail for idiosyncratic reasons, rather than due to macroeconomic factors. Thus there is an assumption that banks are isolated from macroeconomic shocks. However, Elsinger uses historical data on the fluctuating values of banks' asset portfolios to examine the role that market risk plays in systemic risk. This data spans a 13 year period with exposures aggregated into 26 categories, and introduces to the model the possibility that macroeconomic factors will influence the extent of contagion. It is worth pointing out that when building such a model, there is an implicit assumption that the future will resemble the past; if this is not the case, then the model does not give a true reflection of the risks facing the system in question. Thus, the way this model estimates the likelihood of contagious defaults and systemic risk in Austria may exclude the possibility of 'Black Swan Events' (Taleb, 2007) which may have a severe impact on systemic risk.

Deposit flight is a factor that can play an important role in a contagious default scenario, so it is unfortunate that in all of the models seen here it is assumed not to take place. The closest we see is in Georg's model where each bank's deposits are modelled as stochastic from each time period to the next. This means that it is possible that a bank in Georg's model may get into trouble due to deposit flight, but there is no behaviour driving the movement of the deposits from one bank to another, or indeed the movement of deposits out of all banks simultaneously.

The next assumption on the list is that portfolio and asset prices remain constant. This assumption is prevalent in the earlier models in the field, and in particular the Network Models, where market and liquidity risk is omitted completely, but subsequent models from Nier, Cifuentes, Elsinger and Georg have seen a relaxation of this assumption. Nier and Cifuentes model a simple market for assets where defaulting banks (Nier, Cifuentes) or banks not meeting a capital adequacy ratio (Cifuentes) sell a single asset type, forcing down the value of the asset held by other banks. Elsinger, as mentioned above, models portfolio holdings with fluctuating values based on empirical data, and Georg models stochastic returns for 'risky investments'.

It is generally assumed that no seniority in repaying creditors is recognised in the event of a default, but Nier instead models a situation where depositors are paid back in full by the liquidation of a failed bank, with whatever is left distributed proportionately between its bank creditors. This approach is will result in Nier's model overstating the extent of contagion in some instances as the losses passed on to creditors of failed banks will be far greater than would be the case if all creditors were treated identically.

The ninth assumption is that there is no collateralisation of claims. Such collateralisation would reduce the cost to the creditor if a counterparty defaults, reducing the risk of contagion. None of the models discussed here explicitly features collateralisation, but Degryse carries out an analysis in an extension whereby the debts between a defaulting bank and its counterparties are netted against each other, and finds that this may substantially reduce contagion risk.

The next assumption is that bank failures are unexpected. The importance of this assumption is that, if a bank is known to be in trouble, other banks may react by limiting their exposure to this bank. This would have the dual effect of reducing the impact of its failure on its counterparties, but it may also cause a bank that is in slight or temporary difficulty (or a healthy bank that is *rumoured* to be in difficulty) to quickly fail as its access to credit dries up.

The eleventh assumption is that central banks do not rescue failing banks. This is, in a sense, a prerequisite for allowing most of these models to function as many of them start with the assumed failure of a single bank. A more relevant case in the light of the recent credit crisis would be to examine the cost to a central bank and the likelihood of success in attempting to rescue several banks or indeed a whole banking system that has been negatively affected by correlated asset exposures. Georg features a central bank that acts as a lender of last resort for banks that cannot borrow enough from existing counterparties to fund their planned asset portfolio, but the central bank in this model does not attempt to rescue banks that are insolvent. Degryse also examines the effect of merging banks that are at risk of default – an action that could conceivably be forced by regulatory authorities – and finds that in some cases such mergers would prevent contagion taking place.

Another typical assumption is that interbank claims are not backed by government guarantees. This means that exposures to a defaulting bank will always suffer some degree of loss, depending on the LGD applied. Upper departs from other models here however; to more accurately model certain government guarantees that exist in the German market, Upper assumes that some categories of bank will never fail, and another category will not fail in the first contagion round. Incorporating these safety nets, it is found that contagion is still possible but is much more limited in its scope.

The final assumption in Table 1 is that failing banks cannot be recapitalised. In practice, a failing bank may be recapitalised by raising new money from investors, or by receiving a capital injection from government or other sources. Such recapitalisations can restore a bank to solvency and allow it to absorb further losses, and will obviously help prevent contagion. The only model examined here that features such a mechanism is that of Mistrulli, who allows banks belonging to the same banking group share their capital. If a member of the group becomes insolvent, it can receive a capital injection from other banks in the same group. Mistrulli finds that there is not a clear reduction in the risk of contagion if such bailouts are allowed because the improvement may be expected by bailing out these defaulting banks is countered by the addition of a new channel of contagion within banking groups.

4.2 Shock

Table 2: Shock implementation

	Recipient	Mechanism
Eboli	Both system and single bank	External asset losses
Nier	Single bank	External asset losses
Eisenberg	System*	System settlement
Müller	System*	System settlement
Canedo	System	All combinations of banks are assumed to fail with 100% LGD
Matrix	Single bank or group of banks	Bank or banks fail by assumption
Elsinger	System*	System settlement
Cifuentes	Single bank	Single bank fails by assumption with given LGD
Georg	Single banks	Banks may fail due to stochastic deposit holdings and investment loan returns. Additionally, a single bank may fail by assumption at a specific point in the simulation.

*Strictly speaking, there is no conventional 'shock' in Eisenberg, Elsinger or Müller, but the instant settlement used to test for contagious default affects the whole system from the outset.

Each model type in this group seeks to explore the phenomenon of contagious default in a financial network. Typically, a shock will be applied to the system and the consequences of the shock are explored, but the nature of this shock is not the same in every case.

Eboli presents two different scenarios – an 'exogenous common shock' and an 'idiosyncratic shock'. In the case of the former, the shock is generated by reducing the value of the external assets held by the all banks in the network. If the fall in the value of a bank's assets is greater than its net worth, then that bank is insolvent. Losses that are not absorbed by the bank's net worth are passed on to its creditors. Clearly, this scenario is useful in modelling situations where all banks in the system are affected by a common exposure to a certain asset class (for example, sub-prime loans). The second scenario, 'idiosyncratic shock' is intended to model a situation where a single bank is affected, for example in the case of fraud by an employee. The shock mechanism is identical to that in the previous case, but in this scenario only the assets belonging to the shocked bank are affected.

Like Eboli, Nier models the shock as a fall in the value of external assets. However, how Nier implements this is slightly different. Firstly, Nier does not model the exogenous common shock scenario at all. Rather, for a given instance of a banking network, each bank is shocked in turn and the consequences for the other banks are observed. In other words, an idiosyncratic shock is applied to each bank. Nier argues that 'idiosyncratic shocks are a cleaner starting point for studying knock-on defaults due to interbank exposures and liquidity effects', but also concedes that in their model an aggregate shock large enough to bring down any bank will result in every bank in the system failing due to contagion, which may point to a weakness in the simulation.

In another departure from the method employed by Eboli, the shock does not simply affect the net worth and interbank assets of the bank. Rather, the shock that is not absorbed by the bank's net worth is first absorbed by its interbank liabilities, and any loss exceeding these is finally absorbed by customer deposits. This may result in this model exaggerating the degree of contagion that a given shock may cause, as in reality customer deposits do not have any priority over interbank loans in the event of a default, but Nier allows interbank loans to bear the full burden of the loss, with the customer deposits only being affected once interbank loans are totally wiped out.

Eisenberg, Elsinger and Müller differ from Eboli and Nier in that there is no actual shock to the system at all. Rather, they employ a 'fictitious default algorithm' whereby every node (these nodes are 'financial nodes' in Eisenberg, and are banks in Müller) in the system attempts to call in all of its outstanding interbank assets and pay off all of its interbank liabilities as if the market is being wound down. Whereas asset losses are inflicted on otherwise solvent nodes in Nier and Eboli, pushing them into default, the only nodes that will default in Eisenberg and Müller are those which are fundamentally illiquid or insolvent in the state of nature. This is not the case with Elsinger however, as a node which is fundamentally solvent initially may see its portfolio of assets lose enough value between $t = 0$ and $t = 1$ to force it into insolvency.

Canedo takes a different approach to that of the other models. The shock is the failure of a bank or a combination of banks by assumption, with a total loss of any exposure to the set of failed banks. However, as this model is intended to produce a probability distribution for losses to the system, every possible combination of failures is modelled, and a probability attached to each loss scenario. Therefore Canedo models every possible 'shock' that may occur within the limits of the model.

As with Canedo, a Matrix model typically implements the shock by modelling scenarios where a bank or group of banks fail by assumption. Cifuentes similarly assumes that a single bank fails by assumption. Georg allows individual banks to fail by modelling stochastic customer deposit holdings which may fall suddenly (simulating a run on the bank) and stochastic investment loan returns (simulating losses on the asset side). In addition, this model allows specific banks to fail by assumption at predetermined points during a simulation run.

4.3 Data

Of the models considered here, Müller, Elsinger, Canedo and the Matrix models use empirical data from a real-world banking system. The remaining models use data that is generated to examine the general principles of financial contagion by means of the respective models. A comparison can be made of the 'data points' that are required for each model. The following table illustrates the data points each model uses.

Table 3: Data points used

Paper	Interbank Loans	Net Worth	External Assets	Customer Deposits	Liquid Assets	Credit Lines
Eboli	X	X	X			
Nier	X	X	X	X		
Eisenberg	X	X				
Müller	X	X			X	X
Canedo	X	X				
Matrix*	X	X				
Elsinger**	X	X				
Cifuentes	X	X		X	X (plus illiquid assets)	
Georg	X	X	X	X	X	

*This indicates the **minimum** data required for a Matrix model. **This does **not** include the data required for the portfolio model.

A distinction is made in the table between 'external assets' and 'liquid assets'. It could be argued that a liquid asset may also be an external asset. However, the distinction made here is between those papers which have generic data for ‘external assets’ and those which make a distinction between liquid and illiquid assets. In the case of Müller, the use of liquid asset data is important as these are the only assets available in the short term for banks to sell to raise liquidity. Matrix models may feature a breakdown of different asset classes to allow an accurate estimate for Tier 1 capital (e.g. Amundsen and Arnt).

Every model has data for interbank loans and net worth. These data types are the fundamental building blocks of a network model of financial contagion. In the case of Eisenberg, this data is called ‘cash flow’ in the model, but it performs the same role and may be considered indentical to ‘net worth’ in the other papers. If lending plus cash flow minus borrowings is less than zero, then a node in Eisenberg defaults. Similarly, Müller uses ‘regulatory capital’ to fulfil this ‘net worth’ role in the model. Note that Elsinger’s ‘net worth’ figure is derived from a model that simulates market and credit risk; this portfolio model requires many other types of data, excluded here to focus on the network model. The Matrix models often use ‘Tier 1 capital’ as the net worth figure.

The number of data types employed in each model is correlated to the model’s complexity. Eisenberg uses the fewest possible types of data – just interbank loan and ‘net worth’ data. Eboli introduces the concept of a shock to the financial system, and to model this shock adds the external assets data to the model. The shock to the system is thus modelled as a drop in the value of external assets held by one or more of the banks. Nier, in turn, builds on Eboli by introducing customer deposits as a sink for losses inflicted on the system by a shock and requires the addition of data about customer deposits. Cifuentes, adding a market for illiquid assets, requires the introduction of distinct liquid and illiquid asset types.

Müller builds on Eisenberg in two different dimensions – firstly by extending the model to include credit lines, and secondly by populating the model with empirical data. Credit line data is added to the model, in addition to ‘liquid asset’ data. Liquid asset data is required to allow Müller to make the distinction between banks that are insolvent and those that are illiquid, an important distinction as illiquid banks may solve their liquidity problems by calling on their credit lines to access the liquidity they need to pay their creditors.

Discussing the data used in these models allows us to make a further distinction between the model types. Müller, Canedo, Elsinger and the Matrix models use empirical data. Of the models that do not, only Nier discusses how the data used is generated. Eboli and Eisenberg do not disclose how their models were populated with data. Nier outlines in detail the step-by-step process used to populate the model, but in summary the data is generated by varying five key parameters that can be used to create any possible instantiation of Nier's model (as mentioned previously when discussing System Parameters).

Müller's empirical data originates from the Swiss National Bank and relates to 300 Swiss and international banks that participate in the Swiss banking system. The data on interbank loans is not complete – Müller indicates that 83% of interbank assets and only 58% of interbank liabilities can be assigned to particular counterparties. Furthermore, Müller indicates that the credit line data is estimated, as the size of the credit line is not specified by contract so that the bank that receives the credit does not know exactly how much credit the other bank is willing to extend. However, she states that these estimates are 'fairly good' (Müller, pp41).

Like Müller, Canedo's data is obtained from a central bank, in this case the Banco de México. The Mexican banking system is not as complex as the Swiss system, and the model uses data from only 25 domestic banks (compared to the 300 of Müller, some of which are international). The propagation mechanism in Canedo is quite simple – there are no asset price effects or external asset holdings, so the number of data points required in this model is quite small. However, there are two data points not seen in the other models. Each bank has a probability of default, estimated by the central bank, which is used in determining the probability distribution of losses. Also each bank has an exogenous 'loss given default' – defined here as the amount of money lost to the banking system in the event that that bank becomes insolvent.

Matrix models also typically derive their data from the central bank of the country the study focuses on. However, one of the key issues in creating a useable matrix of interbank liabilities is obtaining complete data. Frequently, researchers who build Matrix models do not have access to complete data, and are forced to create estimates of the missing data based on the data that they do have. Two methods are often employed to deal with this problem of incomplete data: entropy maximisation (see for example Mistrulli, Sheldon and Maurer), and cross-entropy minimisation (see Wells).

One source of information on banks' borrowings are their balance sheets, where their total lendings and borrowings from other banks in the matrix can be found. These totals correspond to the row and column totals in the matrix. Using these totals, entropy maximisation works by filling the missing elements in the matrix assuming that each bank distributes its lending as evenly as possible between all the other banks within the limits imposed by the column and row totals that are known. This method can be used in the absence of any specific lending data where the amount lent between two banks is known, but any information that is available can be incorporated, increasing the accuracy of the estimates. An important consequence of applying this method is that it is likely to distort the true pattern of lending in the system by linking banks that in reality have no mutual exposures. By reducing the concentration of interbank lending, this assumption will have a direct impact on how contagion spreads.

Cross-entropy minimisation aims to increase the accuracy of the estimates of unknown matrix elements by making use of other data sources that may contain clues as to how each bank's loans are distributed. In essence, this technique involves creating a second matrix of known elements whose values are assumed to provide an indication of the lending patterns hidden amongst the unknown elements of the matrix. For example, central banks will frequently require banks to inform them of particularly large exposures and who the counterparties are for these large exposures. Working under the assumption that this large exposure data mirrors the patterns of interbank lending, a matrix of this large exposure data is created and the unknown values in the interbank loan matrix are estimated to resemble as closely as possible the data in the second matrix.

4.4 Results

Table 4: Topics addressed by model output

	Level of Capitalisation	Size of Exposures	Degree of Connectivity	Banking System Concentration	Network Structure	Asset Liquidity Effects	Credit Lines	System Specific
Eboli	X	X	X					
Nier	X	X	X	X	X	X		
Eisenberg								
Müller					X		X	X
Canedo					X			X
Matrix								X
Elsinger								X
Cifuentes	X		X			X		
Georg			X					

Each paper explores different aspects of the contagion problem, but there are overlapping findings in certain areas. We will therefore consider the results in terms of these aspects rather than paper by paper.

4.4.1 Level of Capitalisation

The results indicate that higher levels of capitalisation reduce the degree of contagion.

Eboli finds that increasing the amount of capital held by banks results in a reduction in the degree of contagion, while Nier examines this parameter in far greater detail. He agrees with Eboli that increasing levels of bank capitalisation result in decreasing levels of contagious defaults, but in their benchmark case the relationship is non-linear. It emerges that when levels of capitalisation fall low enough to result in second round defaults there is then a stabilisation in the number of banks that default. This situation holds until net worth falls to about 1%, when a third round of defaults is seen to occur. Nier explains that the reason for this stabilisation is that, until the percentage of net worth falls below this critical level, there is still sufficient net worth in the other banks in the network to absorb the impact of both the failure of the first bank and those that fail in the second round. However, once net worth falls below this level, third and subsequent rounds of defaults occur and the number of banks that fail increases dramatically.

Cifuentes also notes the stabilising effects of higher levels of capitalisation, and observes that if banks hold more than the minimum required by the CAR, they may not need to make any balance sheet adjustments at all if a counterparty defaults, obviating the risk of asset price contagion effects.

4.4.2 Size of Interbank Exposures

The results in this area indicate that as the size of the total interbank loan market increases relative to other asset types, the risk of contagion increases.

Eboli finds that as the ratio of interbank exposures to the other assets grows, the amount of losses that defaulting banks pass on to their creditors also grows. Thus, a larger interbank loan market will result in greater contagion. Nier's model supports this result, finding that at low levels of interbank lending, there is no contagion at all, as most of the losses are absorbed by customer deposits and the loans that are defaulted on are small enough to be absorbed by the net worth of the lending banks. However, once a certain threshold is exceeded, second round defaults start to occur. Nier notes that a third round of defaults does not occur at “any realistic proportion of interbank assets” (pp13) as the net worth of both borrowing and lending banks increase as the size of interbank exposures rise, and the increasing levels of net worth serve to absorb the losses of the defaulting banks.

4.4.3 Degree of Connectivity

The results indicate that the degree of connectivity has a dual effect on the risk of contagion; adding connections increases the risk it at low levels of connectivity, but decreases the risk at higher levels of connectivity.

This is the third parameter considered by Nier and Eboli. A loan between two banks has a dual effect: it can act as a shock transmitter when it directs losses to a given bank, and it can act as a shock absorber when the bank receiving the shock is able to absorb the loss with its net worth. Eboli, examining the problem mathematically, concludes that at low levels of connectivity, increasing connectivity reduces the risk of contagion as losses are distributed more evenly between nodes. Conversely, at high levels of connectivity, increasing connectivity further increases the likelihood of cycles and closed paths occurring in the network, which will increase the losses directed at certain nodes and thus increase the likelihood of contagious default. Finally, as connectivity reaches it maximal level, the beneficial effects of connectivity reassert themselves and the likelihood of contagious default decreases. Eboli’s mathematical findings are borne out by Nier’s model. Nier observes that the dual effects of increasing connectivity are more clearly seen when levels of net worth are low.

Cifuentes finds that, with the addition of asset price effects, higher degrees of connectivity may actually increase the risk of contagion. This is because the default of a counterparty may force a bank to dump illiquid assets on the market, driving down the price. If the defaulting bank is connected to many counterparties, then many banks may be forced to sell illiquid assets, causing a greater price movement (and hence, contagion) than if only a few banks were forced to sell. However, Cifuentes finds that this relationship is non-monotonic. If there are sufficient connections between the failed bank and its counterparties that the counterparties only have to sell their liquid assets, then asset price falls can be avoided and contagion prevented.

Georg finds that increased network connectivity results in greater stability, and that this trend is monotonic.

4.4.4 Concentration of Banking System

Increasing the concentration of the banking system increases the risk of contagion.

Increased concentration means that there are fewer banks for a system of the same size, as measured by total system assets. By comparing the impact of shocks of different sizes on systems of 10, 15, 20 and 25 banks, Nier finds that the number of defaults increases in all cases as the size of the shock applied increases. The fraction of banks that default also increases as the concentration of the system increases. Nier argues that there are two factors that influence this result. Firstly, there is a larger shock when a given percentage of a larger bank's assets are wiped out than when the same percentage of a smaller bank's assets are hit. Secondly, in a more concentrated system, loans between banks appear to have an increased tendency to spread contagion.

4.4.5 Asset Liquidity Effects

Asset liquidity effects always lead to an increase in contagion risk.

In Nier's model, the number of failures in a system with asset liquidity effects is never less than the number in a system without, all other things being equal. Contagious failures, and the failure of every bank in system, become far more likely. Combining asset liquidity effects with varying levels of concentration, Nier finds that concentrated bank systems are particularly vulnerable if asset prices are quite liquid.

Cifuentes considers a case where the initial failure in the system occurs with an LGD of zero, meaning that all the subsequent losses to the system are caused by asset price effects. They find that these asset price effects can be a powerful channel of contagion, and that banks are more likely to survive such a scenario if they hold higher levels of liquid assets.

4.4.6 Network Structure

Centralised networks seem more susceptible to contagious defaults than decentralised systems, while the impact of a shock applied to large bank in a centralised system has a dual effect depending on the degree of network connectivity.

Müller observes that the banking system in Switzerland is not homogenous – there are two large banks that act as 'money centres' and there are two sub-networks, one of which is highly centralised (regional banks) and the other decentralised (cantonal bank).

This is an example of a tiered network, where there are many banks with a few interbank loans, and a few banks with many interbank loans. Müller compares the effects of a simulated default in the regional banks sub-network and the cantonal bank sub-network and finds that the potential for contagious default is higher in the centralised regional bank network than in the more homogenous cantonal bank system.

Nier also examines the implications for contagious defaults in a tiered network with a single large bank. If the large bank is shocked, the degree of contagion initially rises with increasing connectivity above the level of defaults caused by a shock to a small bank. When the large bank is connected to about half of the smaller banks, the number of defaults falls. From a connectivity level of about 70% for the large bank, the level of contagious defaults falls well below that of a shock to one of the smaller banks. The explanation of this result is simple – at low levels of connectivity, the default of a large bank is damaging enough to cause any bank that is exposed to it to fail. At a certain tipping point, between 40% and 50% connectivity in Nier's example, the shock of the large bank's default is spread between enough small banks that the small banks have a chance of surviving. As further connections are added, the shock is divided among so many small banks that most or all of them are able to survive.

Canedo makes an interesting observation on how the degree of connectivity affects the outcome of their loss probability distribution. Four cases are modelled. A reference case is modelled using data that would be considered to represent 'normal' conditions in the Mexican banking system. Then three stressed scenarios are modelled. In the first, interbank exposures are set at the highest level recorded over a two-year period from 2004 to 2006, all other parameters being equal to the test case. In the second stressed scenario, the probability of default of each bank in the system is set to a level equal to the probabilities that would have been estimated for them during a period of financial distress such as that experienced in Mexico in 1994. Again, all other parameters are equal to those in the reference case. Finally, a third stress scenario combines both the higher exposure levels of scenario one with the higher PDs of scenario two.

The findings of these stress scenarios reveal that the losses experienced by the system in stress scenario two are far higher than those in stress scenario one, indicating that changing the PDs of the banks in the study has a far greater effect on the probability loss distribution than changing the network topography (by increasing the number and level of exposures between banks).

4.4.7 Credit Line Availability

The availability of credit lines reduces the risk of contagion.

Müller compares two scenarios using empirical data on the Swiss interbank market. She first simulates a default without the existence of credit lines, and then carries out the same simulation with the credit lines implemented. In the first simulation, she finds that there are five rounds of contagion, resulting in 9% of Swiss banks become insolvent and 30% become illiquid. In the second simulation, illiquid banks are allowed to draw on their credit lines with solvent, liquid banks. In the second simulation, 5% of banks become insolvent and 25% of banks become illiquid. In most cases, banks become illiquid due to a combination of credit exposures and credit lines that they cannot call on.

4.4.8 System-Specific Results

Models that use empirical data will sometimes have findings that may have implications for systems other than those that the data was derived from, and we have already considered some of these findings when considering the general findings of Müller and Canedo. However, this is not always necessarily the case, particularly with regard to Matrix models.

Matrix models are based on empirical data from individual banking systems, usually with the stated intention of assessing the risk of contagious default in these systems. As such, the findings of these models tell us far more about the particular system they model than about financial networks in general. However, by looking at the results of the Matrix models as a group, we can still make some general observations.

Typically, Matrix models find that the likelihood and extent of contagious failures is quite low. Lubloy, Krznar, Furfine and Amundsen & Arnt find that contagion would affect less than ten percent of their banking systems as measured by total assets (for Hungary, Croatia, USA, and Denmark respectively). These findings stand even when the systems are tested with LGDs approaching 100% (in the case of Lubloy, Krznar, Amundsen & Arnt). Krznar finds that no bank in the Croatian system has interbank liabilities that exceed their regulatory capital, so contagious default due to idiosyncratic failure is not possible at all.

Upper & Worms find a higher risk of serious contagious default. With a high LGD (75%) they find a worst case scenario where 2444 banks of the 3246 in the system (75%) fail due to contagion following an initial idiosyncratic default. However, they also run simulations where they attempt to model the ‘safety nets’ in place in the German banking system. They do this by assuming that certain classes of banks never fail due to government guarantees and that cooperative banks pool their capital so that either none fail or all fail together. This time, they find that, even with an LGD of 75%, the worst case scenario sees only about 100 banks (representing approximately 15% of total banking assets) fail.

There are also some interesting findings with regard to the initial source of the contagion. Degryse & Nguyen find that the failures of foreign banks are likely to cause far more serious contagion than the failure of a domestic bank in the Belgian banking system. Conversely, Mistrulli finds that the failure of a domestic bank has more serious consequences than the failure of a foreign bank in the Italian system.

5. Limitations of Existing Research

While each study discussed here has advanced the understanding of financial contagion, there are definite limitations to existing research and future advancements in this field will require that these limitations are addressed. Here, we will briefly consider some of the more important shortcomings that exist at present.

5.1 Agent Behaviour

None of the models examined here are capable of modelling any form of behaviour on the part of the agents in the models. It could be argued that the manner in which illiquid banks in Müller’s model can draw on credit lines is a type of behaviour, but the exposures and credit lines that a bank has are determined by the initial set-up of the model – once the model starts, the outcome is totally deterministic. In the real world, banks that are exposed to a bank that appears to be in danger of defaulting can take steps to reduce their exposure to that bank, but none of the existing models allow for this. Note that it is not clear whether modelling this behaviour would result in fewer banking failures. While some banks in the system will benefit from withdrawing credit to potentially illiquid organisations, others will be the victims of this liquidity shortage and may become more likely to default in turn.

In addition, if customers are modelled as agents, we may be able to more accurately simulate bank runs, a potentially important element of financial contagion. Although it may contribute to system instability, each agent in the model should act in a self-interested manner. Behaviour requires action on the part of the agents, and action requires that the model has a time dimension. Clearly, some mechanism for modelling these types of agent behaviour would add valuable depth and realism to existing research.

5.2 Model Scope

A second limitation of existing models is that the financial system modelled tends to be rather narrowly defined – the participants are banks, with the possibility of modelling insurance companies in some models. However, the 'ecosystem' that banks exist in is rather more complex in the real world, and it is possible that modelling this ecosystem more accurately will allow researchers to make better predictions about financial contagion and the consequences of bank defaults. For example, different types of banks will have different mixes of assets on their balance sheets and will have different lending patterns – the balance sheet of a commercial bank will be quite different to that of a mortgage bank. Furthermore, banks and insurers are not the only important agents in the market – the central bank plays a key role in the market, for example in providing liquidity to illiquid banks (Freixas, Parigi, & Rochet, 2000), and Georg takes an important step forward in terms of model scope by introducing a central bank.

In addition, the financial regulator will determine the rules that the whole market operates under, governments can intervene in the market by nationalising or recapitalising failing banks, and large industrial customers or industry sectors can fail with serious consequences to those banks that have leant to them.

5.3 Modelling of Asset Types

A third limitation of the existing models is that, with the exception of Elsinger, they focus either exclusively on interbank loans, or interbank loans and illiquid assets only. Elsinger introduces the effects of market and credit risk on the asset portfolios of the banks they model, but the use of empirical data for a specific 12 year period means that the model can only explore the type of conditions that prevailed during that particular period.

The models that feature illiquid assets demonstrate that even banks that are not directly affected by interbank loan losses can be brought down by falling asset values. However, banks or other agents such as insurance companies may also be exposed to the risk of defaults by the issuers of bonds or other securities. Indeed, an organisation that has bought bonds from and lent to a defaulting organisation will be hit twice by a single default event. Similarly, a company may not have lent to a defaulting bank or be exposed to any assets that this bank dumps on the market, causing price falls, and yet take a large loss as the company defaults on its bonds. Clearly, if we seek to learn more than general principles of financial contagion, we may need to model more than just interbank exposures and an abstracted asset market.

5.4 Data Availability

Finally, a fourth limitation arises out of the second and third – the availability of data. Data pertaining to interbank lending may be commercially sensitive and subject to rapid change, so it is not surprising that it is not easily obtained. It is not a coincidence that most of the studies featured in this paper were projects undertaken for central banks, the only organisations likely to have accurate data on interbank loans. If central banks have - or could obtain - sufficiently detailed information on banks' other asset holdings, it may be possible to build more realistic models as described above where a wider range of exposures and potential contagion channels is considered. However, independent researchers without access to such data may have to continue to rely on simulated data.

6. Future Research

At present, there is considerable scope for new research in the field of network models of financial contagion. In this section we will consider some possible avenues.

6.1 Future Research Suggested in Existing Literature

Previous researchers have highlighted several areas for future research. Eisenberg and Noe have suggested allowing more than one clearing date in the models to incorporate 'true dynamics'. Allowing only a single clearing date severely restricts the opportunity to model any sort of behaviour on the part of the agents modeled in the existing systems – at present in network models, if contagion occurs it is a completely deterministic process. Allowing second and subsequent clearings would enable the banks or other agents in the models to react to the events at the first clearing by - for example - cutting their exposures to a failing bank, greatly adding to the realism. However, Eisenberg and Noe acknowledge that this would be a complex extension, and it is worth noting that nearly a decade later no researcher has published such an extended model. Canedo and Jaramillo acknowledge that the assumption of independence in the probabilities of default used in their model is quite a strong one and indicate that in future research they will seek to address this.

6.2 Agent Behaviour

The limitations of existing research are a good starting point for future research. As stated above, there is no mechanism in the existing research to model the reactions of banks to a default in the system or modelling any other behaviour by participants in the system. Developing an agent-based model that allows such reactions could lead to a greater understanding of the likely patterns of contagion that can arise in financial systems. In addition, an agent-based model may allow researchers to explore the consequences of modifying the behaviour and reactions of the banks in the system to determine whether proposed changes in regulation and monitoring would lead to improved stability or, as is frequently occurs when interacting with a complex system, harmful negative consequences for system stability (the 'law of unintended consequences').

6.3 Model Scope

The second limitation cited above describes the narrow definition of the financial 'ecosystem' that existing models encompass. A broader model that includes more real-world elements such as those identified previously – governments, regulators, central banks, insurance companies and other large non-financial companies and industry sectors – is an obvious direction for future research efforts. While some of these elements – insurance companies and large companies – could be quite easily added to existing models, modelling the role of central banks, governments and regulators may be more complex, perhaps requiring the creation of an agent-based model, as discussed previously, to capture the roles that these participants play. Insurance companies are similar to banks in that they hold huge amounts of assets (of varying liquidity) but differ in that they do not engage in interbank lending, while large non-financial companies can be modelled as depositors, borrowers and bond-issuers interacting with the banks. The central bank and the government, on the other hand, may need to play a different role in the model, reacting to a perceived threat of bank failure or contagion, and these reactions may require that they are modelled as agents behaving according to certain rules.

6.3 Credit Scoring Applications

Another interesting possibility for future research would be to use network information in the field of credit scoring (Thomas, Edelman, & Crook, 2002) to improve credit rating models. By mining the network information of a counterparty, it may be possible to add valuable information to existing credit rating models. For example, when making a lending decision, it may be relevant to consider who the potential counterparty is exposed to, the size of those exposures, and the creditworthiness of *its* counterparties. By determining the likelihood that the potential counterparty is *itself* exposed to credit risk, it may be possible to derive a new predictive variable about the credit-worthiness of that counterparty. The use of network information has proven to be beneficial in a marketing context (Hill, Provost, & Volinsky, 2006) and it may have particular usefulness in the context of the financial services industry where small improvements in credit-rating models can make a big difference in financial performance.

6.4 Stress Testing Applications

Finally, stress-testing is an area of great interest and importance both to banks and regulatory authorities, and this is another possible application of financial network models. The Basel III accords allow banks who wish to adopt the Internal Ratings Based approach (IRB) to perform stress tests to determine their capital requirements. Banks are required to model the losses that their asset portfolios could sustain in certain scenarios such as recessions or other harmful economic events (Basel Committee on Banking Supervision, 2009). Regulatory authorities also need to know the likely consequences of these stress scenarios on the stability not only of individual banks but of the whole financial system. From the perspective of a regulator, an accurate and realistic model of the specific financial network that they are responsible for would be an invaluable tool for investigating the likely consequences of economic events on the system that they are responsible for, and determining the responses that are most likely to lead to desired outcomes.

6.5 The Use of Network Models

Little is known at present about the extent to which these network models are used outside the academic world. The potential uses of these models and the utility that they could offer credit risk professionals, regulators and central banks have been discussed in this paper, but at present we are unaware of any research that explores their use for practical applications and how those who could use network models perceive their usefulness. Addressing this gap in the literature may be useful in determining how to prioritise future technical improvements in the models and also in learning to what extent the development of network models has or has not influenced how those in industry conduct their work.

7. Conclusion

Network models offer the possibility of exploring the ways that the failure of a financial agent or agents will affect other financial agents that are part of the same financial system. The ability to endogenously model the LGD for each individual default that occurs in a system has consequences both for the study of theoretical systems and those based on real-world data. In studying theoretical examples, it is possible to closely examine the mechanisms and characteristics that promote or inhibit contagious default and draw generalised conclusions. In the study of models of real-world systems, it allows a more accurate modelling of the idiosyncrasies of particular banks that have a large bearing on the role they play in increasing or decreasing the risk of contagious default.

However, existing models are limited by the absence of a behavioural dimension that would allow the modelling of the financial agents' actions and reactions, and by a rather narrow definition of the world that these agents operate in. Models based on empirical data from real-world systems are rare as it is unusual for any one organisation to have all the loan exposure data required, and estimating these exposures undermines the accuracy of an endogenously derived LGD, making the model far less useful.

In this paper we have outlined a proposed definition for a financial system network where financial agents are linked to each other by loan exposures. In addition, we have explored various existing implementations of such network models and have noted the differences and similarities of these models and the advantages and shortcomings of each. We also make some observations of possible future research in the field, noting that future research should focus on addressing the shortcomings identified here and elsewhere in the literature, while qualitative research on the use of network models would also be useful.

References

Allen, F., & Gale, D. (2000). Financial Contagion. *Journal of Political Economy*, *108*(1), 1. JSTOR. doi: 10.1086/262109.

Amundsen, E., & Arnt, H. (2005). *Contagion risk in the Danish interbank market*. Retrieved April 28, 2011, from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Contagion+Risk+in+the+Danish+Interbank+Market#0.

Basel Committee on Banking Supervision. (2009). Principles for Sound Stress Testing Practices and Supervision. *BIS*, (May).

Brunnermeier, M. K. (2009). Deciphering the Liquidity and Credit Crunch 2007–2008. *Journal of Economic Perspectives*, *23*(1), 77-100. doi: 10.1257/jep.23.1.77.

Canedo, J., & Jaramillo, S. M. (2009). A network model of systemic risk: stress testing the banking system1. *Intelligent Systems in Accounting, Finance and Management*, *16*(1-2), 87-110. doi: 10.1002/isaf.

Cifuentes, R., Ferrucci, G., & Shin, H. S. (2005). Liquidity Risk and Contagion. *Journal of the European Economic Association*, *3*(2), 556-566. doi: 10.1162/1542476054472946.

Degryse, H., & Nguyen, G. (2007). Interbank exposures: An empirical examination of contagion risk in the Belgian banking system. *International Journal of Central Banking*, 123-171. Retrieved April 28, 2011, from http://www.bportugal.pt/en-US/EstudosEconomicos/Conferencias/Documents/2005FinancialStability/Degryse-Nguyen.pdf.

Diamond, D. W., & Dybvig, P. H. (1983). Bank Runs, Deposit Insurance, and Liquidity. *Journal of Political Economy*, *91*(3), 401. doi: 10.1086/261155.

Eboli, M. (2007). Systemic risk in … nancial networks : a graph-theoretic approach. *Management*, (1993), 1-19.

Eisenberg, L., & Noe, T. H. (2001). Systemic Risk in Financial Systems. *Management Science*, *47*(2), 236-249. doi: 10.1287/mnsc.47.2.236.9835.

Elsinger, H., Lehar, A., & Summer, M. (2006). Risk assessment for banking systems. *Management Science*, *52*(9), 1301. INFORMS. Retrieved May 3, 2011, from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=423985.

Freixas, X., Parigi, B. M., & Rochet, J.-C. (2000). Systemic Risk, Interbank Relations, and Liquidity Provision by the Central Bank. *Journal of Money, Credit and Banking*, *32*(3), 611. doi: 10.2307/2601198.

Furfine, C. H. (1999). INTERBANK EXPOSURES : QUANTIFYING THE RISK OF CONTAGION Monetary and Economic Department, (70).

Georg, C., & Poschmann, J. (2010). Systemic risk in a network model of interbank markets with central bank activity. *Jena Economic Research Papers*. Retrieved May 30, 2011, from http://econstor.eu/dspace/handle/10419/36676.

Hill, S., Provost, F., & Volinsky, C. (2006). Network-Based Marketing: Identifying Likely Adopters via Consumer Networks. *Statistical Science*, *21*(2), 256-276. doi: 10.1214/088342306000000222.

IMF. (2009). *The State of Public Finances: A Cross-Country Fiscal Monitor*. *Economic Outlook* (July.). International Monetary Fund. Retrieved March 21, 2011, from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.167.6364&rep=rep1&type=pdf.

James, C. (1991). The Losses Realized in Bank Failures. *The Journal of Finance*, *46*(4), 1223. doi: 10.2307/2328857.

Krznar, M. (2009). *Contagion Risk in the Croatian Banking System*. Retrieved May 3, 2011, from http://hnbnetra.hnb.hr/publikac/istrazivanja/w-020.pdf.

Lelyveld, I. V., & Liedorp, F. (2004). Interbank Contagion in the Dutch Banking Sector : A Sensitivity Analysis �. *Central Banking*, *31*(0), 99-133.

Lublóy, A. (2005). Domino effect in the Hungarian interbank market. *Hungarian Economic Review*, *52*(4), 377–401. Retrieved March 21, 2011, from http://www.efmaefm.org/efma2005/papers/3-agnes_paper.pdf.

Mistrulli, P. E. (2010). Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking & Finance*, *35*(5), 1114-1127. Elsevier B.V. doi: 10.1016/j.jbankfin.2010.09.018.

Müller, J. (2006). Interbank Credit Lines as a Channel of Contagion. *Journal of Financial Services Research*, *29*(1), 37-60. doi: 10.1007/s10693-005-5107-2.

Nier, E., Yang, J., Yorulmazer, T., & Alentorn, A. (2008). Working Paper No . 346 Network models and financial stability Working Paper No . 346. *Network*, (346).

Sheldon, G., & Maurer, M. (1998). Interbank Lending and Systemic Risk : An Empirical Analysis for Switzerland. *Journal of Economics*, *134*(1997), 685-704.

Stern, G. H., & Feldman, R. J. (2004). *Too Big To Fail*. Washington, DC: Brookings Institution Press.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable.* New York: Random House.

Thomas, L. C., Edelman, D. B., & Crook, J. N. (2002). *Credit Scoring and its Applications*. Philadelphia: SIAM.

Upper, C., & Worms, A. (2004). Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review*, *48*(4), 827–849. Elsevier. Retrieved March 21, 2011, from http://linkinghub.elsevier.com/retrieve/pii/S0014292104000145.

Wells, S. (2002). UK interbank exposures: systemic risk implications. *Financial Stability Review*, *13*(December), 175–182. Retrieved April 28, 2011, from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:UK+interbank+exposures+:+systemic+risk+implications#0.