Chapter 1
Systemic risk in banking networks without Monte Carlo simulation

James P. Gleeson, T. R. Hurd, Sergey Melnik, and Adam Hackett

Abstract An analytical approach to calculating the expected size of contagion events in models of banking networks is presented. The method is applicable to networks with arbitrary degree distributions, permits cascades to be initiated by the default of one or more banks, and includes liquidity risk effects. Theoretical results are validated by comparison with Monte Carlo simulations, and may be used to assess the stability of a given banking network topology.

1.1 Introduction

The study of contagion in financial systems is currently very topical. "Contagion" refers to the spread of defaults through a system of financial institutions, with each successive default causing increasing pressure on the remaining components of the system. The term "systemic risk" refers to the contagion-induced threat to the financial system as a whole, due to the default of one (or more) of its component institutions, and it has become a familiar term since the failure of Lehman Brothers and the rescue of AIG in the autumn of 2008.

J. P. Gleeson
MACSI, Department of Mathematics & Statistics, University of Limerick, Ireland, e-mail: james.gleeson@ul.ie

T. R. Hurd
Department of Mathematics and Statistics, McMaster University, Canada e-mail: hurdt@mcmaster.ca

S. Melnik
MACSI, Department of Mathematics & Statistics, University of Limerick, Ireland, e-mail: sergey.melnik@ul.ie

A. Hackett
MACSI, Department of Mathematics & Statistics, University of Limerick, Ireland, e-mail: adam.hackett@ul.ie

Interbank (IB) networks constitute financial systems that range in size from dozens to thousands of institutions (Boss *et al.*, 2004; Upper and Worms, 2004; Wells, 2002). An IB network may be modelled as a (directed) graph; the *nodes* or *vertices* of the network are individual banks, while the *links* or *edges* of the network are the loans from one bank to another. Such systems are vulnerable to contagion effects, and the importance of studying these complex networks has been highlighted by Andrew Haldane, Executive director of Financial Stability at the Bank of England in his speech (Haldane, 2009), in which he posed the following challenge: 'Can network structure be altered to improve network robustness? Answering that question is a mighty task for the current generation of policymakers'.

The study of complex networks has advanced rapidly in the last decade or so, with large-scale empirical datasets becoming readily available for a variety of social, technological, and biological networks (see Newman, 2010, 2003; May *et al.*, 2008, for reviews). By virtue of their size and complexity, such networks are amenable to statistical descriptions of their characteristics. The *degree distribution* p_k of a network, for example, gives the probability that a randomly-chosen node of the network has degree k, i.e., that it is connected by k edges to neighbours in the network. While classical random graph models of networks (Erdös and Rényi, 1959) have Poisson degree distributions, many empirical networks have been found to possess "fat-tailed" or "scale-free" degree distributions, where the probability of finding nodes of degree k decays as a power law in k ($p_k \propto k^{-\beta}$) for large k, in contrast to the exponential decay with k of the Poisson distribution (Newman, 2003).

This structural (topological) aspect of real-world networks has important implications for dynamical systems which run on the nodes of the network graph, see Barrat *et al.* (2008) for a review. For example, the rate of disease spread on networks depends crucially on whether or not they have fat-tailed degree distributions. As a consequence, there is considerable interest in the effect of network structure on a range of dynamics. Cascade-type dynamics occur whenever the switching of a node into a certain state increases the probability of its neighbours making the same switch. Examples include cascading failures in power-grid infrastructure (Motter and Lai, 2002), congestion failure in communications networks (Moreno *et al.*, 2003), the spread of fads on social networks (Watts, 2002), and bootstrap percolation problems (Baxter *et al.*, 2010), among others (Lorenz *et al.*, 2009). Building on earlier work on the random field Ising model of statistical physics (Dhar *et al.*, 1997), the expected size of cascades has recently been determined analytically for a range of cascade dynamics and (undirected) network topologies (Gleeson, 2008b; Gleeson and Cahalane, 2007). Our goal in thus paper is to extend and develop these methods for application to default contagion on (directed) interbank networks.

Although the importance of network topologies has been recognized for many years in the finance and economics literature (e.g., Allen and Gale, 2000), it is only with the publication of empirical datasets for large-scale interbank networks (Boss *et al.*, 2004; Upper and Worms, 2004; Wells, 2002) that theoretical models have moved beyond small networks and simple topologies. In this paper we focus on models for contagion on interbank networks exemplified by those of Gai and Kapadia (2010) ("GK" for short) and of Nier *et al.* (2007) ("NYYA" for short), which have attracted significant recent attention (May and Arinaminpathy, 2010; Haldane and May, 2011). We develop an analytical approach to calculating the expected size of contagion events in networks of a prescribed topology. The calculation is "semi-"analytical because it requires the iteration of a nonlinear map to its fixed point, but it nevertheless offers significantly faster calculation than Monte Carlo simulation. This reduces the computational burden of interbank network simulations, hence making network theory more useful for practical applications. Moreover, the analytical approach gives insights into the mechanisms of contagion transmission in a given network topology, and enables formulas relating critical parameter values to be derived.

Our work extends the seminal paper of May and Arinaminpathy (2010) by moving beyond their assumption that every bank in the network is identical (i.e., that all banks have the same numbers of debtors and creditors). As shown by May and Arinaminpathy, this "mean-field" assumption gives reasonably accurate results for Erdös-Rényi random networks, which have independent Poisson distributions for in- and out-degrees. This means that each bank in such a network is similar to the "average" bank. However, real-world banking networks often have fat-tailed degree distributions (Boss *et al.*, 2004), meaning that there is a significant probability of finding a bank with in-degree (or out-degree) very different to the mean degree. To analyze contagion on such networks we need to move beyond the mean-field assumption. Moreover, unlike May and Arinaminpathy, our formalism allows us to consider how the extent of the contagion is affected by the size of the bank which initiates the cascade, and so to inform the question of which banks are 'too big to fail'.

The remainder of this paper is structured as follows. In Section 1.2 we review the models of GK and NYYA. Sections 1.3 and 1.4 develop a general theoretical framework for analyzing such models, while in Section 1.5 we compare the results of our analytical approach with full Monte-Carlo simulations, and discuss conclusions in Section 1.6. Three appendices give details of several results that are not crucial to the main flow of the paper.

1.2 Models of contagion in banking networks

Fig. 1.1 Skeleton structure of the network locality of bank i. Bank i is in the $(j,k) = (3,2)$ class, since it has 3 debtors and 2 creditors in the interbank (IB) network.

We consider simplified models of banking networks, as introduced by GK and NYYA. As noted in May and Arinaminpathy (2010), such "deliberately oversimplified" mathematical models are caricatures of real banking networks, but may nevertheless lead to useful insights. These model networks can be considered as generated in two steps. First, a "skeleton" structure of N nodes (representing banks) and directed edges (to represent the interbank positions) is created. This structure should be a realization from the ensemble of all possible directed networks which are consistent with the joint probability p_{jk} (the probability that a randomly chosen node has j in-edges and k out-edges). We choose the following convention for the direction of edges: an arrow on an edge representing an interbank position ("loan" for short) points from the debtor bank to the creditor bank, see Figure 1.1. This convention ensures that shocks due to defaults on loans travel in the direction of the arrows on the edges. Thus p_{jk} is the probability that a randomly-chosen bank in the system has j debtors (or, more strictly, that it has j asset loans, since multiple links are possible) and k creditors (strictly speaking, k liability loans).

In the second step, each node (bank) of the skeleton structure is endowed with a balance sheet and the edges between banks are weighted with loan magnitudes. This process is performed in such as way as to ensure the banking system so represented is fully in equilibrium (i.e., assets exceed liabilities for each bank) in the absence of exogenous shocks. Once the banking networks are generated, the cascade dynamics can be implemented to examine the effects of various types of shocks. In Monte Carlo implementations, each step of the process (skeleton structure/balance sheets/dynamics) is repeated many times to simulate the ensemble of possible systems. The most common output from such simulations is the expected fraction of defaulted banks in steady-state (i.e., when all cascades have run their course) for the prescribed p_{jk} network topology.

We stress that this two-step procedure is only one of many possible alternatives for generating an ensemble of random networks. However, it is easily explained and reproducible by other researchers, and proves amenable to analysis. As a "deliberately oversimplified" model of the true complexities of banking networks, it is not suitable for calibration to real network data in its current form, but may nevertheless provide a starting point for improving our understanding of the interplay between network topology and default contagion cascades.

1.2.1 Generating model networks

We first discuss the creation of the skeleton structure for N banks (or nodes) consistent with a prescribed p_{jk} distribution. It is usually assumed that N is large (indeed theoretical results are proven only in the $N \to \infty$ limit), but in practice values of N as low as 25 have been successfully examined (see Results section). In each realization, N pairs of (j,k) variables are drawn from the p_{jk} distribution. For each pair (j,k), a node is created with j in-edge stubs and k out-edge stubs. Then a randomly-chosen out-stub is connected to a randomly-chosen in-stub to create a directed edge of the network. This process is continued until all stubs are connected. Note it is possible under this process for multiple edges to exist between a given pair of nodes, or for a node to be linked to itself, but both these likelihoods become negligibly small (proportional to $1/N$) as $N \to \infty$. Note also that interbank positions are not netted, so directed edges may exist in both directions between any two nodes of the banking network.

The second step of the network generation process, the creation of balance sheets

Fig. 1.2 Schematic balance sheet of banks in the $(j,k) = (3,2)$ class.

for each bank node, can vary considerably from model to model. In both the GK and NYYA models, the balance sheet quantities of a node depend on its in-degree (number of debtors) j and out-degree (number of creditors) k; we collectively refer to all banks with j debtors and k creditors as the "(j,k)-class". The total assets a_{jk} of a (j,k)-class bank are the sum of its external assets e_{jk} (such as property assets), and its interbank assets, i.e., the sum of its j loans to other banks, see Fig. 1.2. The liabilities side of the balance sheet is composed of the interbank liabilities (sum of the k loans taken from other banks) and customer deposits. The amount by which the total assets exceed the total liabilities is termed the *net worth* of the bank, and is denoted c_{jk} for banks in the (j,k) class. Within both the GK and NYYA models the net worth c_{jk} is assumed (in the initial, shock-free, state) to be proportional to the total assets a_{jk} of the bank:

$$c_{jk} = \gamma a_{jk}, \tag{1.1}$$

where the constant of proportionality γ is termed the "percentage net worth" or "capital reserve fraction". Note that shareholders' funds and subordinated debt are not considered here as useful to the loss absorption capacity; thus only three categories (interbank, customer deposits, and capital) appear on the liabilities side of the balance sheets.

An important difference between the GK and NYYA models is in how they assign values to loans, see Fig. 1.3. Recall the number of loans is determined by the

Fig. 1.3 Loan sizes in each of the models for a bank in the (j_i,k_i) class. In the GK model, all asset loans are of size $0.2/j_i$; liability loans are determined endogenously (by the random linking of in-stubs to out-stubs described in Section 1.2.1). In the NYYA model, every loan in the network is of equal size w.

number of directed edges in the skeleton structure of the first step, but there remains considerable freedom in allocating the weight to each edge. In the GK model (Fig. 1.3(a)), each bank is assumed to have precisely 20% of its assets as interbank assets, and all in-edges to a (j,k)-class node (i.e. all asset loans of a (j,k) bank) are assigned equal weight $0.2/j$ (in units where the total assets of every bank equals unity):

$$a_{jk} = 1, \quad e_{jk} = 0.8 \quad \text{for all } (j,k) \text{ classes.} \tag{1.2}$$

This case represents a maximum-diversity lending strategy, where banks give loans of equal size to all their debtors (Gai and Kapadia, 2010).

	GK	NYYA
total assets of a (j,k)-class bank	$a_{jk} = 1$	$a_{jk} = \tilde{e} + w\max(j,k)$
net worth of a (j,k)-class bank	$c_{jk} = \gamma a_{jk}$	$c_{jk} = \gamma a_{jk}$
size of asset loans of (j,k)-class bank	$\frac{0.2}{j}$	w
external assets of (j,k)-class bank	$e_{jk} = 0.8$	$e_{jk} = \tilde{e} + w\max(0,k-j)$

Table 1.1 Summary of main balance sheet quantities within the GK and NYYA models (see Gai and Kapadia (2010) and Nier *et al.* (2007) for details).

In the model of NYYA, on the other hand, the same weight w is assigned to all directed edges in the network (Fig. 1.3(b)). A (j,k)-class node therefore has interbank assets of jw, and interbank liabilities of kw. To ensure all banks are initially solvent, NYYA describe a process for distributing a pool of external assets over the banks (see Nier *et al.* (2007) for details). As a consequence, the resulting total assets and external assets may respectively be written as

$$a_{jk} = w\max(j,k) + \tilde{e}, \quad e_{jk} = a_{jk} - jw \quad \text{for all } (j,k) \text{ classes,} \tag{1.3}$$

where $\tilde{e}$ is related to the pool of external assets. The balance sheet quantities and their definitions within the two models considered are summarized in Table 1.1.

1.2.2 Contagion mechanisms

Having generated the banking system via the network skeleton structure and balance sheet allocations, the dynamics of cascading defaults can then be investigated. Recall that the banks' balance sheet have been set up so that the system is initially in equilibrium, i.e., total assets for each bank equals its total liabilities plus its net worth. The effect of an exogenous shock is simulated, typically by setting to zero the external assets of one (or more) banks. The shocked bank(s) may be chosen randomly from all banks in the simulation, or a specific (j,k)-class may be targeted—the latter case allows us to investigate the impact of the size of the initially shocked bank upon the final cascade size (see Results section). The initial exogenous shock is intended to model, for example, a sudden decrease in the market value of the external assets held by the bank. The decrease may lead to a situation where the total liabilities of the bank now exceed the total assets: in this case, the bank is deemed to be in default As a consequence, the bank will be unable to repay its creditors the full values of their loans; the loans from these creditors to the defaulted bank are termed "distressed". The creditors (in network terminology, the out-neighbors of the original "seed" bank) experience a shock to their balance sheets at the next timestep due to writing-down the value of the distressed loans. If at any time the total of the shocks received by a bank (i.e. the total losses to date on its loan portfolio) exceeds the net worth of the bank, then its liabilities exceed its assets, and it is deemed to be in default. The defaulted bank then passes shocks to its creditors in the system, and so the cascade or contagion may spread through the banking network. Timesteps are modelled as being discrete, with possibly many banks defaulting simultaneously in each timestep, and with the shocks transmitted to their creditors taking effect in the following timestep.

The mechanism of shock transmission is treated differently by GK and by NYYA, and this is an important distinction between the models.

1.2.2.1 Shock transmission in the GK model

In the GK model, defaulted banks do not repay any portion of their outstanding interbank debts because the timescale for any recovery on these defaulted loans is assumed to exceed the timescale of the contagion spread in the system. Consequently, all creditors of a bank which defaulted in timestep n receive, at timestep $n+1$, a shock of magnitude equal to the total size of their loan to the defaulted bank. If multiple banks defaulted at timestep n, then a bank which is a creditor of several of these will receive multiple shocks at timestep $n+1$. Specifically, if the creditor bank is in the (j,k) class, then it receives a total shock of size $0.2\mu/j$, where μ is the number of its asset loans which defaulted at timestep n (since each loan is of size $0.2/j$, see Table 1.1). This process of shock transmission continues until there are no new defaults, at which point the cascade terminates.

1.2.2.2 Shock transmission in the NYYA model

The NYYA model allows for the possibility of non-zero recovery on defaulted loans. Suppose the total shock received by a (j,k)-class bank from all its defaulted debtors is of size σ, and this shock is sufficient to make the bank default, i.e., $\sigma > c_{jk}$. The amount $\sigma - c_{jk}$ by which total liabilities now exceed total assets for the bank is distributed evenly among the k creditors of the bank, with the proviso that no creditor can lose more than the size w of its original loan (recall every loan in the NYYA system is the same size w, see Table 1.1). Thus the shock transmitted to each creditor of the defaulted bank is

$$\min\left(\frac{\sigma - c_{jk}}{k}, w\right). \tag{1.4}$$

As in the GK model, shocks transmitted from banks which default at timestep n will affect the creditor banks at timestep $n+1$, and a cascade of banks failures may ensue. This cascade mechanism bears some resemblance to the “fictitious default” cascade used by Eisenberg and Noe (2001) (“EN” for short) to determine the clearing payment vector in a system with defaults, see Appendix A. However, the NYYA cascades are not identical to the EN cascades. When a bank defaults in the NYYA model, it transmits a once-off shock to each of its creditors, but then plays no further role in the dynamics of the system. In particular, any shocks received by this bank subsequent to its default do not affect its creditors. In contrast, the EN clearing algorithm effectively requires defaulted banks to transmit newly-received shocks to their creditors. Although the EN algorithm is not the main focus of this paper, we present in Section 1.5 (see Figures 1.5(a) and 1.6(a)) numerical results for the fraction of defaults in EN cascades. The results are qualitatively similar, though not identical, to those obtained using the NYYA contagion dynamics, the difference being most notable in cases where a large fraction of the network is in default.

1.2.3 Liquidity risk

In both the GK and NYYA dynamics, it is possible to include liquidity risk effects in a simple fashion. Suppose that at timestep n, a fraction ρ^n of all banks in the system have already defaulted. It is plausible that the market value of external assets (e.g., property) will be adversely affected by the weakened banking system. A bank needing to liquidate its external assets may, for example, find it difficult to realise the full value in a “fire sale” scenario. To model the effects of this system-wide effect, we assume that at timestep n the external assets of a (j,k)-class bank are marked-to-market as

$$e_{jk}\exp\left(-\alpha\rho^n\right). \tag{1.5}$$

The liquidity risk parameter α measures the influence of the system contagion upon asset prices; note when $\alpha = 0$ the external asset values are constant over time, but for $\alpha > 0$ the asset values decrease with increasing contagion levels. This effect is included in the dynamics of the GK and NYYA models by subtracting the quantity $e_{jk}\left[1-\exp(-\alpha\rho^n)\right]$ from the net worth c_{jk} of the (j,k)-class banks. Thus, for example, banks default in the NYYA model if the incoming shock s is bigger than $c_{jk} - e_{jk}\left[1-\exp(-\alpha\rho^n)\right]$ (the fire-sale adjusted net worth), and the shock transmission equation (1.4) is generalized to

$$\min\left(\frac{\sigma - c_{jk} + e_{jk}\left[1-\exp(-\alpha\rho^n)\right]}{k}, w\right), \tag{1.6}$$

for $\alpha \geq 0$. A similar modification applies in the GK model. Interestingly, if α is sufficiently large, the liquidity risk effect can lead to banks defaulting even if they receive no shocks from debtors, because their net worth is obliterated by the fall in market value of their external assets. Consequences of this are explored in the Results section.

1.2.4 Monte Carlo simulations

The steps needed to study the models using Monte Carlo simulation are now clear. In each realization a skeleton structure for a network of N nodes with joint in- and out-degree distribution p_{jk} is first constructed. Then balance sheets are assigned to each node, consistent with the specific model chosen (see Table 1.1). The cascade of defaults initiated by an exogenous shock to one (or more) banks proceeds on a timestep-by-timestep basis, following the dynamics of either the zero recovery (GK) or non-zero recovery (NYYA) prescription for shock transmission. When no further defaults occur, the fraction of defaulted banks (the "cascade size") is recorded, and then another realization may begin. When a sufficiently large number of realizations are recorded, the average cascade size (and potentially further statistics, i.e., the variance, of the cascade size) may be calculated in a reproducible (up to statistical scatter) manner. Monte Carlo simulations of this type were implemented in GK and NYYA; our focus in the remainder of this paper is on analytical approaches to predicting the average size of cascades, and so avoiding the need for computationally expensive numerical simulations.

1.3 Theory

In this section we derive analytical equations which allow us to calculate the expected fraction of defaults in a banking network with a given topology (defined by p_{jk}). Like related approaches for cascades on undirected networks (Gleeson and Cahalane, 2007; Gleeson, 2008b), the method is only approximate for finite-sized networks because it assumes the $N \to \infty$ limit of infinite system size. However, in practice we find it nevertheless gives reasonably accurate results for networks as small as $N = 25$ banks (see Section 1.5).

1.3.1 Thresholds for default

We begin by defining the threshold level M^n_{jk} as the maximum number m of distressed loans that can be sustained by a (j,k)-class bank at timestep n without the bank defaulting at timestep $n+1$. If a (j,k)-class bank has m defaulted debtors, with $m > M^n_{jk}$, then it will default in the subsequent timestep, otherwise it will remain solvent. As we show below, the GK model is easily expressed in terms of thresholds, but thresholds can be defined for the NYYA model only under an approximating assumption.

In the GK model a bank in the (j,k) class has total assets of unity ($a_{jk} = 1$), net worth of $c_{jk} = \gamma a_{jk} = \gamma$, and each distressed loan carries a shock of $0.2/j$. In the absence of a liquidity risk (fire sale) factor, the (j,k) bank would then default if the sum of the shocks it receives from its m defaulted debtors exceeds its net worth, i.e., if $0.2m/j > \gamma$, giving $M^n_{jk} = \lfloor 5jc_{jk} \rfloor$, where $\lfloor\cdot\rfloor$ is the floor function (returning the greatest integer less than or equal to its argument). Liquidity risk may also be included in models of this type by appropriately reducing the effective net worth, and we can write the threshold levels in their most general form as

$$M_{jk}^{n} = \min\left\{ j, \max\left\{ \lfloor 5jc_{jk} - 5je_{jk}\left(1 - e^{-\alpha\rho^{n}}\right) \rfloor, -1 \right\} \right\}. \tag{1.7}$$

Here e_{jk} is the value of external assets for (j,k)-class banks, α is the liquidity risk parameter introduced in Section 1.2 and we constrain M_{jk}^{n} to be between -1 and j. Note that this expression for M_{jk}^{n} is constant over time n if $\alpha = 0$, and is decreasing in time if α is positive and ρ^{n} is increasing.

In the NYYA model the size of the write-down shock on a newly-distressed loan depends on how large the shock received by the debtor bank was compared to its net worth. This means that there will, in general, be a distribution of shocks of various sizes in the system, and this distribution will change in time. Denoting the distribution of shock sizes by $S^{n}(\sigma)$—so that at timestep n a randomly-chosen distressed loan (i.e. an out-edge of a defaulted bank node) carries a shock of size σ with probability $S^{n}(\sigma)$—we would require m-fold convolutions of $S^{n}(\sigma)$ to correctly describe the shock received by a bank with m distressed asset loans (as the sum of m independent draws of shock values from $S^{n}(\sigma)$). It is clearly desirable to find a simple parametrization of $S^{n}(\sigma)$ to make the model computationally tractable, even at the loss of some accuracy. With this in mind, we approximate the true value of the shock received by a bank with m distressed loans at timestep n by ms^{n}, where s^{n} is the average shock on all distressed loans in the system at that timestep. Effectively we are replacing the true distribution $S(\sigma)$ of shock sizes by a delta function distribution: $S^{n}(\sigma) \mapsto \delta(\sigma - s^{n})$, where s^{n} is the average shock $s^{n} = \int \sigma S^{n}(\sigma) d\sigma$; in other words, every distressed loan at timestep n is assumed to have equal recovery value $w - s^{n}$. This approximation turns out to work rather well because in cases where many debtors are in default, the total shock received by a creditor is well approximated by m times the average shock. However we will also show examples (in the Results section) where the approximation of the shock distribution $S^{n}(\sigma)$ by a delta function leads to less accurate results.

Using this approximation, the NYYA threshold levels are:

$$M_{jk}^{n} = \min\left\{ j, \max\left\{ \lfloor \frac{1}{s^{n}} \left[c_{jk} - e_{jk}\left(1 - e^{-\alpha\rho^{n}}\right) \right] \rfloor, -1 \right\} \right\}. \tag{1.8}$$

The time dependence of the thresholds in this case is due to both liquidity risk ($\alpha > 0$), and to the time-varying nature of the (mean) shock size s^{n}. In Appendix B we derive an iteration equation for s^{n}, consistent with the general model (1.12)–(1.13) below and based on the approximation of the true shock size distribution by a delta function.

1.3.2 General theory

We consider (j,k)-class banks, of which there are approximately Np_{jk} in any given network realization (for sufficiently large N). Each bank in the (j,k) class has j debtors, each of which may be either solvent or in default at a specific time. Given that a bank is in the (j,k) class, we define $u^n_{jk}(m)$ as the probability that the bank (i) is solvent at timestep n and (ii) has m distressed asset loans (due to the default of the corresponding debtors in earlier cascades). According to its definition, the sum of $u^n_{jk}(m)$ over all m gives the fraction of (j,k)-class banks which are solvent at timestep n:

$$\sum_{m=0}^{j} u^n_{jk}(m) = 1 - \rho^n_{jk}, \tag{1.9}$$

where ρ^n_{jk} denotes the fraction of (j,k)-class banks which are in default at timestep n. In a slight abuse of mathematical terminology we will refer to $u^n_{jk}(m)$ as a "distribution", but note from (1.9) that the sum of $u^n_{jk}(m)$ over all m is not unity.

We consider how the states of the banks change from timestep n to timestep $n+1$, and update the $u^n_{jk}(m)$ distribution accordingly. The update occurs in two stages: first a "node update" stage, where the states of the banks are updated, followed by an "edge update", where the $u^n_{jk}(m)$ distribution is updated to give $u^{n+1}_{jk}(m)$. In the node update stage, banks in the (j,k) class default if their number of distressed loans m at timestep n exceeds their threshold M^n_{jk} (see Section 1.3.1). Thus the newly defaulting fraction of (j,k)-class banks is made up of those who were previously solvent but now have m values above threshold. These newly defaulted banks increase the total default fraction of the (j,k) class by the amount:

$$\rho^{n+1}_{jk} - \rho^n_{jk} = \sum_{m=M^n_{jk}+1}^{j} u^n_{jk}(m). \tag{1.10}$$

Each newly defaulted (j,k)-class bank is a debtor of k other banks in the system and correspondingly triggers k newly-distressed loans: this is the edge update stage between timestep n and timestep $n+1$. The number of newly-distressed loans in the network due to defaults in the (j,k) class of banks is approximately $Np_{jk}k(\rho^{n+1}_{jk} - \rho^n_{jk})$ (since there are Np_{jk} such banks, each newly-defaulted with probability $\rho^{n+1}_{jk} - \rho^n_{jk}$, and each with k creditors). Summing over all classes gives

$$N\sum_{j,k} kp_{jk}\left(\rho^{n+1}_{jk} - \rho^n_{jk}\right) \tag{1.11}$$

as the number of newly-distressed loans in the system. The total number of loans which were not distressed at timestep n is similarly calculated as $N\sum_{j,k} kp_{jk}\left(1-\rho^n_{jk}\right)$. So the probability that a previously-undistressed loan will be distressed at timestep $n+1$ is given by

$$f^{n+1} = \frac{\sum_{j,k} kp_{jk}\left(\rho^{n+1}_{jk} - \rho^n_{jk}\right)}{\sum_{j,k} kp_{jk}\left(1-\rho^n_{jk}\right)} = \frac{\sum_{j,k} kp_{jk}\sum_{m=M^n_{jk}+1}^{j} u^n_{jk}(m)}{\sum_{j,k} kp_{jk}\sum_{m=0}^{j} u^n_{jk}(m)}. \tag{1.12}$$

Consider a (j,k)-class bank which remains solvent and has exactly m distressed asset loans at timestep $n+1$. This bank was also solvent at timestep n and had some number $\ell \leq \min(m, M^n_{jk})$ of distressed asset loans at timestep n. Amongst the remaining $j-\ell$ asset loans of this bank, exactly $m-\ell$ of the loans must have become newly distressed due to the debtor bank having defaulted in the first stage of the update: this happens independently to each of the $j-\ell$ loans with probability f^{n+1}. If we introduce the convenient notation $B^k_i(p)$ for the binomial probability $\binom{k}{i} p^i (1-p)^{k-i}$, the probability that a (j,k)-class bank remains solvent and has exactly m distressed asset loans at timestep $n+1$ can be written as

$$u^{n+1}_{jk}(m) = \sum_{\ell=0}^{\min(m,M^n_{jk})} B^{j-\ell}_{m-\ell}\left(f^{n+1}\right) u^n_{jk}(\ell). \tag{1.13}$$

Equations (1.12) and (1.13) together define the updating of the state variables $u_{jk}(m)$ and f in terms of the $u_{jk}(m)$ distribution at timestep n. Given the initial condition—for instance, if a randomly-chosen fraction ρ^0 of all banks are initially subject to default-causing shocks, this is $u^0_{jk}(m) = \left(1-\rho^0\right) B^j_m\left(\rho^0\right)$—it is straightforward to iterate the system given by (1.12) and (1.13) forward through the discrete timesteps until it converges to a steady state. The total fraction of defaulted banks in the system at timestep n is given by summing (1.9) over all (j,k) classes:

$$\rho^n = 1 - \sum_{j,k} p_{jk} \sum_{m=0}^{j} u^n_{jk}(m), \tag{1.14}$$

and the steady-state value of this quantity (as $n \to \infty$) is reported for various cases in Section 1.5 below.

In Section 1.4 we prove that a certain class of models, including GK, admits an exact reduction of the system described here to just two state variables. In the GK model, and for the case where a fraction ρ^0 of the banks are chosen at random to be the seed defaults, the fraction of bank defaults ρ^n and the fraction of edge defaults g^n are given by the recurrence

$$\rho^{n+1} = \rho^0 + \left(1-\rho^0\right) \sum_{j,k} p_{jk} \sum_{m=M^n_{jk}+1}^{j} B^j_m\left(g^n\right) \tag{1.15}$$

$$g^{n+1} = \rho^0 + \left(1-\rho^0\right) \sum_{j,k} \frac{k}{z} p_{jk} \sum_{m=M^n_{jk}+1}^{j} B^j_m\left(g^n\right), \tag{1.16}$$

with the initial condition $g^0 = \rho^0$.

For the NYYA model, we use the mean-shock-size approximation discussed in Section 1.3.1, so the thresholds M^n_{jk} are given by equation (1.8). Then the iteration equation for s^n (see Appendix B), along with equations (1.12) and (1.13), gives us a system of equations for $u^{n+1}_{jk}(m)$, f^{n+1}, and s^{n+1} in terms of the values of these quantities at the previous timstep. Results for both models are compared with Monte Carlo simulations in Section 1.5.

1.4 Simplified theory

In this section we show that the iteration of the system defined by equations (1.12) and (1.13) in order to obtain the expected fraction of defaulted banks (as given by equation (1.14)) may be dramatically simplified in certain cases. A sufficient condition for this simplified theory to exactly match the full theory of equations (1.12) and (1.13) is:

Condition 1: For every (j,k) class with $p_{jk}>0$, the threshold level M^n_{jk} is a non-increasing function of n.

This condition holds if the threshold levels for each (j,k) class are constant, or decreasing with time, as in the GK model. For the NYYA model, cases where the shock size decreases over time may have thresholds M^n_{jk} which increase with n, and so this model does not satisfy Condition 1.

1.4.1 Simplified theory for GK

Focussing now on the GK model, whose thresholds (1.7) satisfy Condition 1, we claim that at timestep n, the distribution for the number m of distressed loans of solvent banks is a binomial distribution, at least for m values below the threshold:

$$u^n_{jk}(m)=\left(1-\rho^0_{jk}\right)B^j_m(g^n) \qquad \text{for } m\le M^n_{jk}, \tag{1.17}$$

and the fraction of distressed edges is

$$g^n=\sum_{j,k}\frac{k}{z}p_{jk}\rho^n_{jk}. \tag{1.18}$$

Here ρ^0_{jk} is the initially defaulted fraction of (j,k)-class banks and ρ^n_{jk} is the defaulted fraction of (j,k)-class banks at timestep n. For the case $m>M^n_{jk}$, the values $u^n_{jk}(m)$ are slightly more complicated in form: they are given by the update equation (1.13) for level n, with the right-hand side given using (1.17) at the level $n-1$. As we show below, the result (1.17) is sufficient to determine the expected fraction of defaulted banks at any timestep n.

To prove our claim, we use an induction argument, showing that if the sub-threshold distribution at timestep n is assumed to take the form (1.17), (1.18) then the distribution at timestep $n+1$ (as given by equation (1.13) of the full theory) is also of the form (1.17), (1.18). Substituting for $u^n_{jk}(\ell)$ in (1.13) using (1.17) yields

$$u_{jk}^{n+1}(m) = \left(1-\rho_{jk}^{0}\right) \sum_{\ell=0}^{\min\left(m, M_{jk}^{n}\right)} B_{m-\ell}^{j-\ell}\left(f^{n+1}\right) B_{\ell}^{j}\left(g^{n}\right). \tag{1.19}$$

To satisfy (1.17) at timestep $n+1$ we need only consider values of m between 0 and M_{jk}^{n+1}, and by Condition 1 we have $M_{jk}^{n+1} \leq M_{jk}^{n}$, so that $0 \leq m \leq M_{jk}^{n+1} \leq M_{jk}^{n}$, and thus the upper limit on the summation in (1.19) is $\min\left(m, M_{jk}^{n}\right) = m$. The sum in (1.19) is therefore a convolution sum of two binomial distributions, which is itself a binomial distribution:

$$u_{jk}^{n+1}(m) = \left(1-\rho_{jk}^{0}\right) B_{m}^{j}\left(g^{n+1}\right) \qquad \text{for } m \leq M_{jk}^{n+1}, \tag{1.20}$$

Here g^{n+1} is given by $g^{n+1} = g^{n} + (1-g^{n}) f^{n+1}$. One can now use (1.12) and (1.18) to verify that

$$g^{n+1} = \sum_{j,k} \frac{k}{z} p_{jk} \rho_{jk}^{n+1}. \tag{1.21}$$

By assuming the form (1.17), (1.18) at timestep n we have shown the full theory yields the corresponding result (1.20), (1.21) at timestep $n+1$. The induction proof is completed by verifying that the initial condition is given by

$$u_{jk}^{0}(m) = \left(1-\rho_{jk}^{0}\right) B_{m}^{j}\left(g^{0}\right) \qquad \text{for } m = 0 \text{ to } j, \tag{1.22}$$

$$g^{0} = \sum_{j,k} \frac{k}{z} p_{jk} \rho_{jk}^{0} \tag{1.23}$$

306 which is of the form (1.17), (1.18).

Using the binomial distribution for u_{jk}^{n} in (1.9) and (1.10) gives the update equations for ρ^{n+1} and g^{n+1} in terms of the parameter g^{n} only:

$$\begin{aligned}
\rho^{n+1} = \sum_{j,k} p_{jk} \rho_{jk}^{n+1} &= 1 - \sum_{j,k} p_{jk} \left(1-\rho_{jk}^{0}\right) \sum_{m=0}^{M_{jk}^{n}} B_{m}^{j}\left(g^{n}\right) \\
&= 1 - \sum_{j,k} p_{jk} \left(1-\rho_{jk}^{0}\right) \left(1 - \sum_{m=M_{jk}^{n}+1}^{j} B_{m}^{j}\left(g^{n}\right)\right) \\
&= \rho^{0} + \sum_{j,k} p_{jk} \left(1-\rho_{jk}^{0}\right) \sum_{m=M_{jk}^{n}+1}^{j} B_{m}^{j}\left(g^{n}\right),
\end{aligned} \tag{1.24}$$

and

$$g^{n+1} = \sum_{j,k} \frac{k}{z} p_{jk} \rho_{jk}^{n+1} = \sum_{j,k} \frac{k}{z} p_{jk} \left[\rho_{jk}^0 + \left(1 - \rho_{jk}^0\right) \sum_{m=M_{jk}^n+1}^{j} B_m^j(g^n) \right]$$
$$= \rho^0 + \sum_{j,k} \frac{k}{z} p_{jk} \left(1 - \rho_{jk}^0\right) \sum_{m=M_{jk}^n+1}^{j} B_m^j(g^n), \qquad (1.25)$$

where $\rho^0 = \sum_{j,k} p_{jk}\rho_{jk}^0$ is the overall fraction of initially defaulted banks. In the case where a fraction ρ^0 of the banks are chosen at random to be the seed defaults we have $\rho_{jk}^0 = \rho^0$ for all (j,k) classes, and equations (1.24) and (1.25) reduce to equations (1.15) and (1.16).

The expected size of global cascades in a given GK-model network has essentially been reduced to solving the single equation (1.16), since ρ^{n+1} can be immediately determined by substituting g^n into (1.15). Equation (1.16) is of the form $g^{n+1} = J(g^n)$, and the function $J(\cdot)$ is non-decreasing on $[0,1]$. It follows that $g^{n+1} \geq g^n$ for all n, and iteration of the map leads to the solution g^∞ of the fixed-point equation $g^\infty = J(g^\infty)$. The corresponding steady-state fraction of defaulted banks is determined by substituting g^∞ for g^n in (1.15).

Equations of this sort, giving the expected size of cascades on directed networks, have been previously derived in various contexts (Gleeson, 2008a; Amini *et al.*, 2010). In Gleeson (2008a), the main focus is on percolation-type phenomena (see also the undirected networks case Gleeson (2008b)), while Amini *et al.* (2010) consider more complicated dynamics but take the limit $\rho^0 \to 0$. The general case (1.24), (1.25) where initial default fractions can be different for each (j,k) class has not, to our knowledge, been considered previously, even in Monte Carlo simulations.

In the limit $\rho^0 \to 0^+$, the scalar map $g^{n+1} = J(g^n)$ has a fixed point at $g^n = 0$, but it is an unstable fixed point if $J'(0) > 1$, where J' is the derivative of the function J. Thus the condition for an infinitesimally small seed fraction to grow to a large-scale cascade may, using (1.16), be written as

$$J'(0) = \sum_{j,k} \frac{jk}{z} p_{jk} \Theta\left[\tfrac{0.2}{j} - c_{jk}\right] > 1, \qquad (1.26)$$

where the GK threshold (1.7) for $m = 1$ and $\rho^0 = 0$ has been used, and Θ is the Heaviside step function ($\Theta(x) = 1$ for $x > 0$; $\Theta(x) = 0$ for $x \leq 0$). This “cascade condition” has been derived in a rather different fashion by GK; they extend Watts’ (2002) percolation theory approach from his work on undirected networks to the case of directed networks considered here. In Gleeson and Cahalane (2007); Gleeson (2008b), the generalization of this result to cases where ρ^0 is finite but small has been given for cascades on undirected networks. Similar “higher-order cascade conditions” may similarly be derived for this directed-network case, but are beyond the scope of the present paper.

1.4.2 Frequency of contagion events

The simplified equations (1.15), (1.16), and indeed the more general method of Section 1.3, allow the specification of a fraction ρ^0 (or ρ^0_{jk} in the case of targetted (j,k) classes) of initially defaulted bank nodes. This fraction need not be small, and this feature allows us to investigate features of systemic risk due to simultaneous failure of more than one bank (see Results section). However, most work to date has focussed exclusively on the case where a single initially defaulted bank leads to a cascade of defaults through the network. Because our theory assumes an infinitely large network, some special attention must be paid to the case of a single "seed" default in the GK model. As we show in Appendix C, in this model the locality of the seed node determines whether, in a given realization, a cascade will reach global size, or remain restricted to a small neighborhood of the seed. The distribution of cascade sizes observed in single-seed GK simulations is thus typically bimodal: only a certain fraction (termed the *frequency*) of cascades reach a network-spanning size, the remainder remain small (typically only a few nodes). The average *extent* (i.e. size) of the global cascades is given by equations (1.15), (1.16), whereas the frequency of cascades which escape the neighborhood of the seed may be expressed in terms of the size of connected components for a suitable percolation problem, see Appendix C and the Results section. The NYYA model does not exhibit this sensitivity to the details of the neighborhood of the seed node(s), so its distribution of cascade sizes is quite narrowly centered on the mean cascade size given by theory; the same comment applies to the GK model with multiple seed nodes.

1.5 Results

1.5.1 GK model

Fig. 1.4 Theory and Monte Carlo simulation results for GK model on Erdös-Rényi networks with $N = 10^4$ nodes and mean degree z. The percentage net worth is set to $\gamma = 3.5\%$ for all cases. Cascades which exceed 0.5% of the network are considered as "global" cascades; the "extent" of contagion is the average size of these global cascades, while the "frequency" is the fraction of all cascades that become global cascades. In (b), the effects of non-zero liquidity risk are clearly seen for lower z values, and cause the appearance of a discontinuous transition which is not present in the $\alpha = 0$ case of (a). Monte Carlo numerical results are averages over 5000 realizations.

Figure 1.4(a) compares results of Monte Carlo simulations of the GK model (symbols) with the results of the simplified theory of equations (1.15) and (1.16). As in Fig. 1 of the GK paper, we show the extent and frequency (see Appendix C) of contagion resulting from a single seed default in Erdös-Rényi directed random graphs with $N = 10^4$ nodes. The mean degree z of the network is varied to investigate the effects of connectivity levels upon the contagion spread. In such networks the in- and out-degree of a node (i.e., the number of debtors and creditors of a bank) are independent, and the joint distribution p_{jk} is a product of Poisson distributions:

$$p_{jk} = \frac{z^j}{j!} e^{-z} \frac{z^k}{k!} e^{-z}. \tag{1.27}$$

The formula for the contagion window derived in Gai and Kapadia (2010) (which is the same as our equation (1.26)) predicts that cascades occur for z values between 1 and 7.477, but our theory also accurately predicts the expected magnitude of these events. Moreover, as shown in Fig. 1.4(b), our theory also accurately incorporates the effects of the liquidity risk model (1.5), capturing the discontinuous transition in cascade size which appears above $z = 1$ for the case $\alpha = 0.1$.

1.5.2 NYYA benchmark case

Fig. 1.5 Expected steady-state default fraction in Erdös-Rényi random graphs with mean degree $z = 5$. Monte Carlo numerical simulation results are averages over 5000 realizations. In the networks with $N = 25$ nodes, cascades are initiated by the default of a single randomly-chosen node; in the larger networks with $N = 250$, 10 randomly-chosen nodes are defaulted to begin the cascade; theory uses $\rho^0 = 1/25$.

Figure 1.5(a) examines the benchmark case of NYYA; note our Monte Carlo simulation results match those presented in Chart 1 of Nier *et al.* (2007). The fraction of defaults (extent of contagion) is here plotted as a function of the percentage net worth parameter γ, as defined in equation (1.1). The network structure is again Erdös-Rényi, with p_{jk} given by (1.27), and mean degree $z = 5$. We also show Monte Carlo results for the default fraction resulting from the clearing vector algorithm of Eisenberg and Noe (see Appendix A). This algorithm gives results which are qualitatively similar in behavior (though not identical) to those generated by the NYYA shock transmission dynamics described in equation (1.6). As in the NYYA paper, our Monte Carlo simulations use $N = 25$ nodes (banks) in each realization, and cascades are initiated by a single randomly-chosen bank being defaulted by an exogenous shock. Despite this relatively small value of N, we find very good agreement between the theoretical prediction (which assumes the $N \to \infty$ limit) from equations (1.12) and (1.13), and the Monte Carlo simulation results. The theory also enables us to examine the case where multiple banks are defaulted to begin the cascade. We demonstrate this by also showing numerical results for a larger Erdös-Rényi network of $N = 250$ nodes, with the same mean degree $z = 5$. In order to match the seed fraction of defaults, cascades in the larger networks are initiated by simultaneously shocking 10 randomly-chosen banks (each shock being calibrated to wipe out the external assets of the bank), so $\rho^0 = 1/25 = 0.04$. The numerical results for this case are almost indistinguishable from the $N = 25$ case, and both cases match very well to the theory curve.

In Figure 1.5(b) we increase the liquidity risk parameter from $\alpha = 0$ (as in Figure 1.5(a)) to $\alpha = 0.05$ and $\alpha = 0.1$. For clarity, the results of the Eisenberg-Noe dynamics are not shown here, but as in Figure 1.5(a), they are qualitatively similar to the simulation results using the NYYA shock transmission dynamics. The theory predicts a discontinuous transition in ρ at γ values between 2% and 3% for the $\alpha = 0.05$ and $\alpha = 0.1$ cases, but this is not well reproduced in Monte Carlo simulations with $N = 25$ nodes and $\rho^0 = 1/N$ (triangles). However, this is due to finite-N effects (i.e., due to having a finite-sized network whereas theory assumes the $N \to \infty$ limit), as can be seen by the much closer agreement between the theory and the $N = 250$ (with 10 seed defaults) case (filled circles) for $\alpha = 0.05$.

A more serious discrepancy between theory and numerics can be seen in the γ range 4% to 5%. Here the theory underpredicts the cascade size, and the difference is unaffected by increasing the size of the network. Detailed analysis of this case reveals that the root of the discrepancy is in fact the simplifying assumption made for the shock size distribution $S^n(\sigma)$ in the NYYA case (see Section 1.3.1). By replacing all shocks with the mean shock size we are underestimating (at timestep $n > 1$) the residual effects of the large shock which propagated from the first defaulted node(s) at timestep $n = 1$. Indeed, if we modify the Monte Carlo simulations to artificially replace all shocks at each timestep by their mean, we find excellent agreement between theory and numerics over all γ values. We conclude that the simplifying assumption $S^n(\sigma) \to \delta(\sigma - s^n)$ of the shock size distribution may lead to some errors, and further work on approximating $S^n(\sigma)$ by analytically tractable distributions is desirable. Despite this caveat, overall the theory works very well on the Erdös-Rényi random graphs studied by NYYA.

1.5.3 Networks with fat-tailed degree distributions

As noted in May and Arinaminpathy (2010), empirical data on banking networks indicates that their in- and out-degree distributions are fat-tailed, and so it is important that theoretical approaches not be restricted to Erdös-Rényi networks. Accordingly, for Figure 1.6 we generate a network with joint in- and out-degree distribution given by

$$p_{jk} = C\delta_{jk}k^{-1.7} \quad \text{for } k = 5, 10, 15, \ldots, 50. \tag{1.28}$$

Here C is a normalization constant (so that $\sum_{j,k} p_{jk} = 1$), and the exponent 1.7 has been chosen to be similar to that found for the in-degree distribution in the empirical data set of Boss *et al.* (2004). The Kronecker delta δ_{jk} appears in (1.28) to give our networks very strong correlations between in- and out-degrees: in contrast to the independent j and k distributions of (1.27), here we set the in- and out-degree of every node to be equal (i.e., each bank has equal numbers of debtors and creditors).

We also consider for the first time the effect on the contagion of the size of the initially defaulting bank. If the single bank to be defaulted by the initial exogenous shock is chosen randomly from a specific (j,k) class, denoted (j',k'), then the initial values of ρ^n_{jk} are

$$\rho^0_{jk} = \begin{cases} \frac{1}{Np_{j'k'}} & \text{for } (j,k) = (j',k') \\ 0 & \text{for all other } (j,k) \text{ classes.} \end{cases} \tag{1.29}$$

The corresponding initial conditions for $u_{jk}(m)$ are:

$$u^0_{jk}(m) = \begin{cases} \left(1 - \frac{1}{Np_{j'k'}}\right) B^{j'}_m(g^0) & \text{for } (j,k) = (j',k') \\ B^j_m(g^0) & \text{for all other } (j,k) \text{ classes,} \end{cases} \tag{1.30}$$

where

$$g^0 = \sum_{j,k} \frac{k}{z} p_{jk} \rho^0_{jk} = \frac{k'}{Nz} \tag{1.31}$$

is the fraction of loans (edges) in the network which are initially distressed (i.e. have their debtor bank in default). We use $N = 200$ banks and ignore liquidity effects: $\alpha = 0$. All other parameters are as in the benchmark case of NYYA (Nier *et al.*, 2007).

Fig. 1.6 Comparison of theory and Monte Carlo numerical simulations for banking networks with joint in- and out- degree distribution (1.28), with $N = 200$ nodes. Cascades are initiated by targeting a single node of a specific (j,k) degree class. Monte Carlo simulation results are averages over 5000 realizations. The dashed lines in (b) mark the critical γ values given by (1.33).

Figure 1.6(a) shows the theoretical and numerical results for the case where one of the largest banks in the network (i.e., with $j' = k' = 50$) is targeted initially. Note that the theory accurately matches to the NYYA Monte Carlo simulation results; also note that the Eisenberg-Noe clearing vector case is (at low γ values) somewhat further removed from the NYYA dynamics than in previous figures.

Figure 1.6(b) compares the results of Fig. 1.6(a) to the case where the targeted bank is from the class with $j' = k' = 30$, i.e., a mid-sized bank in this network. Theory and numerics again match well, and over most of the γ range the smaller target bank leads to smaller cascade sizes. Interestingly however, near $\gamma = 2\%$ is a range where the smaller target bank actually generates a larger cascade than the bigger target bank—this phenomenon is clearly visible in both numerical and theoretical results. To explain it, we consider the threshold levels at timestep $n = 0$ (and with $\alpha = 0$). The initially-targeted bank was subject to an exogenous shock that wiped out its external assets and each of its out-edges (liability loans) carries a residual shock (cf. (1.4)) of magnitude

$$s^0 = \min\left(\frac{e_{j'k'} - c_{j'k'}}{k'}, w\right), \tag{1.32}$$

where (j', k') denotes the class of the targeted bank. If a single such shock is to cause further defaults, say of a (j,k)-class node, then the threshold M^0_{jk} must be zero (cf. equation (1.8)). This requires $c_{jk} < s^0$ (note $\alpha = 0$ here), or, using (1.1),

$$\gamma < \frac{s^0}{a_{jk}}. \tag{1.33}$$

The largest critical value s^0/a_{jk} for γ occurs for the lowest $j = k$ value (because of the dependence of a_{jk} on degree, see Table 1.1) and this γ value for each case is marked by the vertical dashed lines in Figure 1.6(b)—note in each case it matches the location of the sudden change in the contagion size. Essentially, this is the level of γ below which a single shock of magnitude s^0 can cause further defaults (moreover, our argument indicates that these further defaults will be among the smallest banks in the system). The shock magnitude s^0 given by (1.32) (see Table 1.1 for details of e_{jk} and c_{jk}) is a non-increasing function of k', and in the crucial γ range the value of s^0 is less for $k' = 50$ than for $k' = 30$. This is reflected in the respective critical values for γ, and allows the $k' = 30$ case to cause larger cascades than the $k' = 50$ case, at least while these cascades are relatively small.

1.6 Discussion

In this paper we have introduced an analytical method for calculating the expected size of contagion cascades in the banking network models of Gai and Kapadia (2010) and Nier *et al.* (2007). Our method may be applied to cases with:

- an arbitrary joint distribution p_{jk} of in- and out-degrees (i.e., numbers of debtors and creditors) for banks in the network. This includes fat-tailed distributions; see equation (1.28) and Fig. 1.6;
- arbitrary initial conditions for the cascade, including the targeting of one or more banks of a specified size (see Fig. 1.6);
- liquidity risk effects (see Figs. 1.4 and 1.5).

In the general case, the theory gives a set of discrete-time update equations (equations (1.12), (1.13), and (1.42)) for a vector of unknowns $\mathbf{g}^n$, which is composed of the state variables f^n, $u^n_{jk}(m)$, and s^n. The update equations may be written in the form $\mathbf{g}^{n+1} = \mathbf{H}(\mathbf{g}^n)$ and this vector mapping is iterated to steady-state to find the fixed point solution $\mathbf{g}^\infty = \mathbf{H}(\mathbf{g}^\infty)$, hence giving the expected fraction of defaults ρ^∞, see Figs. 1.5 and 1.6 for examples. Under certain conditions it proves possible to simplify the equations to be iterated: as shown in Section 1.4, this reduces the vector $\mathbf{g}^n$ to a scalar g^n, with iteration map $g^{n+1} = J(g^n)$. The GK model is of this type, and the simplified equations (1.15) and (1.16) were used to generate the theoretical results in Fig. 1.4. In all cases we find very good agreement between Monte Carlo simulations and theory, even on relatively small ($N = 25$) networks.

We expect it will prove possible to improve and extend these results in several ways. As noted in Section 1.5.2, the approximation of the shock size distribution in the NYYA model leads to some loss of accuracy, and this merits further attention. It is also desirable to develop analytical methods for calculating the frequency of cascades caused by single seeds in the GK model (see Appendix C). Even in its current form, however, the theory presented here is ideally suited to the study of some policy questions. For example, suppose the models are modified so that the capital reserve fraction γ is not the same for all banks in the system, instead depending on the size of the bank (i.e. $\gamma \mapsto \gamma_{jk}$). This requires only a slight modification of the existing equations. The question is then: how should γ_{jk} depend on the (j,k) class in order to optimally reduce the risk of contagion-induced systemic failure? Other possible extensions, such as allowing for the existence of subgroups of banks with different levels of interbank assets or with interbank loans/liabilities drawn from a prescribed distribution, are required to begin modelling the important non-homogeneities that are seen in the real banking system, and these will be the subject of future work.

For these and similar questions, it is likely that a general cascade condition (or "instability criterion"), analogous to equation (1.26) for the GK model, will prove very useful. Cascade conditions for dynamics with vector mappings have been derived for undirected networks (see Gleeson (2008b) and references therein), so we believe that similar methods may be applied to the directed networks analyzed here.

Finally, it is hoped that the methods introduced here will prove extendable beyond the stylized models of Gai and Kapadia (2010) and Nier *et al.* (2007), and in particular that related methods will be applicable to datasets from real-world banking networks. Ideally, such datasets would include information on bank sizes, connections, and the sizes of loans (Bastos *et al.*, 2011). Modelling the distribution of loan sizes within a semi-analytical framework will be challenging, but the understanding gained of how network topology affects systemic risk on toy models will no doubt prove important to finding the solution.

Acknowledgements

We acknowledge the work of undergraduate students Niamh Delaney and Arno Mayrhofer on an early version of the simulation codes used in this paper. Discussions with the participants at the Workshop on Financial Networks and Risk Assessment, hosted by MITACS at the Fields Institute, Toronto in May 2010 (particularly Rama Cont and Andreea Minca) are also gratefully acknowledged, as are the comments of Sébastien Lleo and Mark Davis. This work was funded by awards from Science Foundation Ireland (06/IN.1/I366, 06/MI/005 and 11/PI/1026), from an INSPIRE: IRCSET-Marie Curie International Mobility Fellowship in Science Engineering and Technology, and from the Natural Sciences and Engineering Research Council of Canada.

Appendix A: Generalized Eisenberg-Noe clearing vector cascades

This Appendix provides a summary of the financial cascade framework of Eisenberg and Noe (2001), placed in a slightly more general context. Extending their work somewhat (Eisenberg and Noe (2001) combine the quantities Y_i and D_i into a single quantity $e_i = \mathrm{Y}_i - \mathrm{D}_i$), we identify the following stylized elements of a financial system consisting of N "banks" (which may include non-regulated leveraged institutions such as hedge funds). The assets A_i of bank i at a specific time consists of *external assets* Y_i (typically a portfolio of loans to external debtors) plus *internal assets* Z_i (typically in the form of interbank overnight loans). The liabilities of the bank includes *external debts* D_i (largely in the form of bank deposits, but also including long term debt) and *internal debt* X_i. The bank's *equity* is defined by $\mathrm{E}_i = \mathrm{Y}_i + \mathrm{Z}_i - \mathrm{D}_i - \mathrm{X}_i$ and is constrained to be non-negative.

The amounts $\mathrm{Y}, \mathrm{Z}, \mathrm{D}, \mathrm{X}$ refer to the notional value, or face value, of the loans, and are used to determine the relative claims by creditors in the event a debtor defaults. Internal debt and assets refer to contracts between the N banks in the system. Banks and institutions that are not part of the system are deemed to be part of the exterior, and their exposures are included as part of the external debts and assets. Let $\bar{L}_{ij}$ denote the notional exposure of bank j to bank i, that is to say, the amount i owes j. Note the constraints that hold for all i

$$\mathrm{Z}_i = \sum_j \bar{L}_{ji}, \quad \mathrm{X}_i = \sum_j \bar{L}_{ij}, \quad \sum_i \mathrm{Z}_i = \sum_i \mathrm{X}_i, \quad \bar{L}_{ii} = 0,$$

and that the matrix of exposures $\bar{L}$ is not symmetric.

A.1 Default cascades

A healthy bank manages its books to maintain mark-to-market values with sufficient "economic capital" to provide an "equity buffer" against shocks to its balance sheet. This means that the bank maintains its asset-to-equity ratio A_i/e_i above a fixed threshold Λ_i (a typical value imposed by regulators might be 12.5).

Following a bank-specific catastrophic event, such as the discovery of a major fraud, or a system wide event, the assets of some banks may suddenly contract by more than the equity buffer. Assets are then insufficient to cover the debts, and these banks are deemed insolvent. The assets of an insolvent bank must be quickly liquidated, and any proceeds go to pay off that bank's creditors, in order of seniority. We now discuss three simple settlement mechanisms for how an insolvent bank i is removed from the system.

- Version A, the original mechanism of Eisenberg and Noe (2001), supposes that external debt is always senior to internal debt. We define fractions $\pi_{ij} = \bar{L}_{ij}/\mathrm{X}_i$. If p_i denotes the amount available to pay i's internal debt, this amount is split amongst creditor banks in proportion to π_{ij}, that is bank j receives $\pi_{ij}p_i$. Given $\mathbf{p} = [p_1,\ldots,p_N]$, the clearing conditions are $p_i = 0$ if $\mathrm{Y}_i - \mathrm{D}_i + \sum_j \pi_{ji}p_j < 0$ and $p_i = \min(\mathrm{Y}_i - \mathrm{D}_i + \sum_j \pi_{ji}p_j, \mathrm{X}_i)$ if $\mathrm{Y}_i - \mathrm{D}_i + \sum_j \pi_{ji}p_j \geq 0$. We can write this as

$$p_i = F_i^{(A)}(\mathbf{p}) := \min(\mathrm{X}_i, \max(\mathrm{Y}_i + \sum_j \pi_{ji}p_j - \mathrm{D}_i, 0)), \quad i = 1,\ldots,N \tag{1.34}$$

- Version B supposes that external and internal debt have equal seniority. We define fractions $\tilde{\pi}_{ij} = \bar{L}_{ij}/(\mathrm{D}_i + \mathrm{X}_i)$. If $\tilde{p}_i$ denotes the amount available to pay i's total debt, creditor bank j receives $\tilde{\pi}_{ji}\tilde{p}_i$ while the external creditors receive $\mathrm{D}_i\tilde{p}_i/(\mathrm{D}_i + \mathrm{X}_i)$. The clearing conditions are:

$$\tilde{p}_i = F_i^{(B)}(\tilde{\mathbf{p}}) := \min(\mathrm{D}_i + \mathrm{X}_i, \mathrm{Y}_i + \sum_j \tilde{\pi}_{ji}\tilde{p}_j), \quad i = 1,\ldots,N.$$

- Most simply, Version C supposes as in the GK model that the recovery from any insolvent bank is zero. That means the amount p_i available to pay i's internal debt is simply

$$p_i = F_i^{(C)}(\mathbf{p}) := \mathrm{X}_i \Theta(\mathrm{Y}_i - \mathrm{D}_i + \sum_j \pi_{ji}p_j)$$

where Θ denotes the Heaviside function.

Under each of these settlement mechanisms, any solution $\mathbf{p} = (p_1,\ldots,p_N) \in \mathbb{R}_+^N$ of the clearing conditions is called a "clearing vector". In the subsequent discussion we consider only version A. The existence result extends easily to versions B and C by considering fixed points of the monotonic mappings $F^{(B)}, F^{(C)} : \mathbb{R}_+^N \to \mathbb{R}_+^N$.

Proposition 1. *Consider a financial system with* $\mathrm{Y} = [\mathrm{Y}_1,\ldots,\mathrm{Y}_N], \mathrm{D} = [\mathrm{D}_1,\ldots,\mathrm{D}_N]$ *and matrix* $\bar{L} = (\bar{L}_{ij})_{i,j=1\ldots,N}$. *Then the mapping* $F^{(A)} : \mathbb{R}_+^N \to \mathbb{R}_+^N$ *defined by* (1.34) *has at least one clearing vector or fixed point* $\mathbf{p}^*$. *If in addition the system is "regular" (a natural economic constraint on the system), the clearing vector is unique.*

Proof: Existence is a straightforward application of the Tarski Fixed Point Theorem to the mapping F acting on the complete lattice

$$[\mathbf{0}, \bar{X}] := \{\mathbf{x} = [x_1,\ldots,x_N] \in \mathbb{R}_+^N : 0 \leq x_i \leq \bar{X}_i, i = 1,\ldots,N\}.$$

One simply verifies the easy monotonicity results that for any vectors $\mathbf{0} \leq \mathbf{p} \leq \mathbf{p}' \leq \mathrm{X}$ one has

$$\mathbf{0} \leq F^{(A)}(\mathbf{0}) \leq F^{(A)}(\mathbf{p}) \leq F^{(A)}(\mathbf{p}') \leq F^{(A)}(\mathrm{X}) \leq \mathrm{X}$$

(where $\mathbf{a} \leq \mathbf{b}$ for vectors means $a_i \leq b_i$ for all $i = 1,\ldots,N$). For the definition of "regular" and the uniqueness result, please see Eisenberg and Noe (2001).

A.2 Clearing Algorithm

Cascades of defaults arise when primary defaults trigger further losses to the remaining banks. The above proposition proves the existence of a unique "equilibrium" clearing vector that characterizes the end result of any such cascade. The following algorithm for version A of the settlement mechanism resolves the cascade to the fixed point $\mathbf{p}^*$ in at most $2N$ iterations by constructing an increasing sequence of defaulted banks $A^k \cup B^k, k = 0,1,\ldots$. Analogous (but simpler) algorithms are available for settlement mechanisms B and C.

1. **Step 0** Determine the primary defaults by writing a disjoint union $\{1,\ldots,N\} = A^0 \cup B^0 \cup C^0$ where

$$\begin{aligned} A^0 &= \{i | \mathrm{Y}_i + \mathrm{Z}_i - \mathrm{D}_i < 0\} \\ B^0 &= \{i | \mathrm{Y}_i + \mathrm{Z}_i - \mathrm{D}_i - \mathrm{X}_i < 0\} \setminus A^0 \\ C^0 &= \{1,\ldots,N\} \setminus (A^0 \cup B^0). \end{aligned}$$

2. **Step** $k,\ k = 1,2,\ldots$ Solve the $|B^{k-1}|$ dimensional system of equations:

$$p_i = \mathrm{Y}_i - \mathrm{D}_i + \sum_{j \in C^{k-1}} \pi_{ji} \mathrm{X}_j + \sum_{j \in B^{k-1}} \pi_{ji} p_j,\ i \in B^{k-1}$$

and define result to be $\mathbf{p}^{k*}$. Define a new decomposition

$$\begin{aligned} A^k &= A^{k-1} \cup \{i \in B^{k-1} | p_i^{k*} \le 0\} \\ B^k &= (B^{k-1} \setminus A^k) \cup \{i \in C^{k-1} | \mathrm{Y}_i - \mathrm{D}_i + \sum_{j \in C^{k-1}} \pi_{ji} \mathrm{X}_j + \sum_{j \in B^{k-1}} \pi_{ji} p_j^{k*} \le \mathrm{X}_i\} \\ C^k &= \{1,\ldots,N\} \setminus (A^k \cup B^k) \end{aligned}$$

and correspondingly

$$p_i^k = \begin{cases} 0 & i \in A^k \\ \mathrm{Y}_i + \sum_{j \in C^k} \pi_{ji} \mathrm{X}_j + \sum_{j \in B^k} \pi_{ji} p_j^{k*} - \mathrm{D}_i & i \in B^k \\ \mathrm{X}_i & i \in C^k. \end{cases} \tag{1.35}$$

If $A^k = A^{k-1}$ and $B^k = B^{k-1}$, then halt the algorithm and set $A^* = A^k, B^* = B^k, \mathbf{p}^* = \mathbf{p}^{k*}$. Otherwise perform step $k+1$.

Appendix B: Updating of average shock strength for NYYA model

Assuming a delta function distribution approximating $S^n(\sigma)$ as in Section 1.3.1, we need to count the number of loans (edges in the directed network) which link defaulted banks to solvent banks. In the notation of Section 1.3.2, the number of such "d-s" (for "defaulted-to-solvent") edges in the network at timestep n is

$$N\sum_{j,k} p_{jk} \sum_{m=0}^{j} m\, u_{jk}^{n}(m), \tag{1.36}$$

since each solvent bank with m defaulted debtors contributes m d-s edges to the total. We assume that all these d-s edges at timestep n carry an equal shock s^n.

Now consider the situation at timestep $n+1$. Some of the d-s edges from timestep n are still d-s edges, although others will have become d-d ("defaulted-to-defaulted") edges. We count the number of d-s edges which remained as d-s from timestep n to timestep $n+1$ as

$$A^{\text{old}} = N\sum_{j,k} p_{jk} \sum_{m=0}^{M_{jk}^{n}} m\, u_{jk}^{n}(m). \tag{1.37}$$

Note the upper limit of M_{jk}^{n} for the sum over m (cf. equation (1.36)); this arises because the creditor banks in question remain solvent at timestep $n+1$.

The other mechanism generating d-s edges at timestep $n+1$ is the default of the debtor end of a timestep-n s-s (solvent-to-solvent) edge. Similar to (1.36), we can count the number of s-s edges at timestep n as

$$N\sum_{j,k} p_{jk} \sum_{m=0}^{j} (j-m) u_{jk}^{n}(m), \tag{1.38}$$

since each (solvent) (j,k)-class bank with m defaulted debtors must also have $j-m$ solvent debtors. Each of the s-s edges at timestep n becomes an d-s edge at timestep $n+1$ if (i) the debtor bank defaults during the timestep, and (ii) the creditor bank remains solvent to at least timestep $n+1$. Noting that (i) occurs with probability f^{n+1} (see equation (1.12) of the main text), and that (ii) requires $m \le M_{jk}^{n}$, we obtain the number of new d-s edges at timestep $n+1$ as

$$A^{\text{new}} = f^{n+1} N\sum_{j,k} p_{jk} \sum_{m=0}^{M_{jk}^{n}} (j-m) u_{jk}^{n}(m). \tag{1.39}$$

The total number of d-s edges at timestep $n+1$ is then $A^{\text{old}} + A^{\text{new}}$, while the cumulative total of the shock sizes transmitted by these edges is

$$s^n A^{\text{old}} + \widetilde{s} A^{\text{new}}, \tag{1.40}$$

where $\widetilde{s}$ is the average shock on each newly-distressed loan (using (1.6) of the main text):

$$\widetilde{s} = \frac{\sum_{j,k} k p_{jk} \sum_{m=M_{jk}^n+1}^{j} u_{jk}^n(m) \min\left(\frac{m s^n - c_{jk} + e_{jk}[1-\exp(-\alpha\rho^n)]}{k}, w\right)}{\sum_{j,k} k p_{jk} \sum_{m=M_{jk}^n+1}^{j} u_{jk}^n(m)}. \tag{1.41}$$

Thus, under the simplifying assumption on the shock size distribution ($S^n(\sigma) \mapsto \delta(\sigma - s^n)$) , we model the shocks on d-s edges at timestep $n+1$ to each be of equal size s^{n+1}, where

$$s^{n+1} = \frac{s^n A^{\text{old}} + \widetilde{s} A^{\text{new}}}{A^{\text{old}} + A^{\text{new}}}, \tag{1.42}$$

with A^{old}, A^{new}, and $\widetilde{s}$ given in terms of u_{jk}^n by equations (1.37), (1.39), and (1.41), respectively. This gives an update equation for s^n in terms of known quantities from timestep n.

Appendix C: Frequency of cascades for single-seed initiation in GK model

In this Appendix we consider the frequency of cascades in the GK model when initiated by a single seed node. Mathematically, our theory applies to the limiting case $N \to \infty$ of a sequence of networks of size N, with $\lfloor \rho^0 N \rfloor$ seed nodes. In Monte Carlo simulations of real banking networks, the size N of the system is fixed, and the case of a single seed corresponds to a fraction $\rho^0 = 1/N$ of initial defaults. The mechanism of cascade initiation in the infinite-N network may be understood as follows. As in Watts (2002), we call bank nodes *vulnerable* if they default due to a single defaulting loan. When the cascade condition (1.26) is satisfied, a giant connected cluster of vulnerable nodes exists in the network. The fractional size of this vulnerable cluster is denoted S_v, and it may be calculated by solving a site percolation problem for the directed network (see Meyers, Newman, and Pourbohloul, 2006) in a similar fashion to the calculation for undirected networks in Watts (2002):

$$S_v = \sum_{jk} p_{jk} \left[1 - (1-\phi)^j\right] \Theta\left[\frac{0.2}{j} - c_{jk}\right], \tag{1.43}$$

where ϕ is the non-zero solution of the equation

$$\phi = \sum_{jk} \frac{k}{z} p_{jk} \left[1 - (1-\phi)^j\right] \Theta\left[\frac{0.2}{j} - c_{jk}\right]. \tag{1.44}$$

Here, as in Watts (2002), the Θ term plays the role of a degree-dependent site occupation probability: sites (nodes) are deemed occupied if they are vulnerable in the sense defined above, and this happens if the shock due to a single defaulting loan ($0.2/j$) exceeds their net worth c_{jk}. In Figure 1.7 we directly calculate the size of the largest vulnerable cluster in a single realization of an Erdös-Rényi network with

$N = 10^4$ nodes and mean degree z (cf. Figure 1.4) and show that it closely matches to the analytical result (1.43).

The *extended vulnerable cluster* (Watts, 2002), which takes up a fraction S_e of the network, consists of nodes which are debtors of at least one bank in the vulnerable cluster. If a seed node is part of the extended vulnerable cluster, it immediately causes the default of its creditor in the vulnerable cluster, which in turn leads to default of other nodes in the vulnerable cluster, and so on until the entire vulnerable cluster is in default. Nodes outside the vulnerable cluster (i.e. banks which can withstand the default of a single asset loan) may also be defaulted later on in this cascade as the percentage of defaulted banks increases; the result is a global cascade of expected size ρ^∞, given by equation (1.15). On the other hand, if no seed node is part of the extended vulnerable cluster, then no further defaults will occur and the cascade immediately terminates. Thus, if only a single seed node is used in each realization, we expect cascades of size ρ^∞ to occur in a fraction S_e of realizations (corresponding to cases where the seed node lies in the extended vulnerable cluster), and no cascades to occur in the remaining fraction $1-S_e$ of realizations. The size S_e of the extended vulnerable cluster thus determines the frequency of global cascades among the set of single-seed realizations. The size of S_e was calculated analytically in Gleeson (2008b) for the undirected networks case, but the corresponding derivation for directed networks is non-trivial. Instead, we directly calculate the size of the largest extended vulnerable cluster in the network, and show in Figure 1.7 that it corresponds very closely to the frequency of global cascades in the large ensemble of Monte Carlo simulations of Figure 1.4 in the main text.

Fig. 1.7 Sizes of vulnerable cluster (S_v) and of extended vulnerable cluster (S_e) as calculated directly from (for each value of mean degree z) a single Erdös-Rényi network with $N = 10^4$ nodes. The vulnerable cluster size is compared with the analytical result of equation (1.43), while the extended vulnerable cluster is shown to closely match the frequency of global cascades in the single-seed GK model (cf. Figure 1.4).

As argued in Gleeson (2008b), the frequency of cascades increases with the number $\lfloor \rho^0 N \rfloor$ of seed nodes used as

$$1-(1-S_e)^{\lfloor \rho^0 N \rfloor}, \tag{1.45}$$

which reduces to S_e for the single-seed case ($\rho^0 = 1/N$) and to 1 for the case where ρ^0 remains a finite fraction as $N \to \infty$. The frequency of cascades (of size ρ^∞) in the GK model initiated by a single default is thus S_e, whereas if multiple seeds (say, 10 initial defaults among 1000 banks) are used we find that almost all cascades are of size ρ^∞.

References

Allen, F. and Gale, D. 2000 Financial contagion, *Journal of Political Economy*, **108**, 1–33.

Amini, H., Cont, R., and Minca, A. 2010 Resilience to contagion in financial networks, working paper.

Barrat, A., Vespignani, A., and Barthélemy, M. 2008 *Dynamical Processes on Complex Networks*, Cambridge University Press.

Bastos, E. , Cont, R., and Moussa, A. 2010 The Brazilian banking system: network structure and systemic risk analysis, working paper.

Baxter, G. J. , Dorogovtsev, S. N., Goltsev, A. V., and Mendes, J. F. F. 2010 Bootstrap percolation on complex networks, *Phys. Rev. E*, **82**, 011103.

Boss, M., Elsinger, H., Thurner, S. & Summer, M. 2004 Network topology of the interbank market. *Quantitative Finance* **4**, 677–684.

Dhar, D., Shukla, P., and Sethna, J. P. 1997 Zero-temperature hysteresis in the random-field Ising model on a Bethe lattice, *J. Phys. A: Math. Gen.* **30**, 5259–5267.

Eisenberg, L. and Noe, T. H. 2001 Systemic risk in financial systems,*Management Science,* **47**, (2) 236–249.

Elsinger, H., Lehar, A. & Summer, M. 2004 Analyzing systemic risk in the European banking system: a portfolio approach, Vienna University working paper.

Elsinger, H., Lehar, A. & Summer, M. 2006 Using market information for banking system risk assessment, *International Journal of Central Banking*, **2**, (1) 137–165.

Erdös, P. and Rényi, A. 1959 On random graphs, *Publicationes Mathematicae*, **6**, 290–297.

Gai, P. and Kapadia, S. 2010 Contagion in financial networks, *Proc. R. Soc. A*, **466** (2120) 2401–2423.

Gleeson, J. P. 2008a Mean size of avalanches on directed random networks with arbitrary degree distributions, *Phys. Rev. E*, **77**, 057101.

Gleeson, J. P. 2008b Cascades on correlated and modular random networks, *Phys. Rev. E*, **77**, 046117.

Gleeson, J. P. and Cahalane, D. J. 2007 Seed size strongly affects cascades on random networks, *Phys. Rev. E*, **75**, 056103.

Haldane, A. G. 2009 Rethinking the financial network, online: http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf.

Haldane, A. G. and May, R. M. 2011 Systemic risk in banking ecosystems, *Nature*, **469** (7330) 351–355.

Lorenz, J., Battiston, S., and Schweitzer, F. 2009 Systemic risk in a unifying framework for cascading processes on networks, *Eur. Phys. J. B*, **71** (4) 441–460.

May, R. M. and Arinaminpathy, N. 2010 Systemic risk: the dynamics of model banking systems, *J. R. Soc. Interface,* **7**, (46) 823–838.

May, R. M., Levin, S. A., and Sugihara, G. 2008 Ecology for bankers, *Nature*, **451**, 893–895.

Meyers, L. A., Newman, M. E. J., and Pourbohloul, B. 2006 Predicting epidemics on directed contact networks, *J. Theor. Biol.*, **240**, 400–418.

Moreno, Y., Paster-Satorras, R., Vázquez, A., and Vespignani, A. 2003 Critical load and congestion instabilities in scale-free networks, *Europhys. Lett.*, **62**, 292–298.

Motter, A. E. and Lai, Y.-C. 2002 Cascade-based attacks on complex networks, *Phys. Rev. E*, **66**, 065102(R).

Newman, M. E. J. 2003 The structure and function of complex networks, *SIAM Review*, **45** (2), 167–256.

Newman, M. E. J. 2010 *Networks: An Introduction*, Oxford University Press.

Nier, E., Yang, J., Yorulmazer, T., and Alentorn, A. 2007 Network models and financial stability, *J. Economic Dynamics & Control*, **31**, 2033–2060.

Upper, C. & Worms, A. 2004 Estimating bilateral exposures in the German interbank market: is there a danger of contagion? *European Economic Review* **48**, 827–849.

Watts, D. J. 2002 A simple model of global cascades on random networks, *Proc. Nat. Acad. Sci.*, **99**, (9) 5766–5771.

Wells, S. 2002 UK interbank exposures: systemic risk implications. *Bank of England Financial Stability Review* December, 175–182.