Contagion in Financial Networks

Prasanna Gai[*] and Sujit Kapadia[†]

Abstract

This paper develops a general model of contagion in financial networks, identifying both its probability and potential impact. We explore how contagion risk is influenced by aggregate and idiosyncratic shocks, changes in network structure, and asset market liquidity. Our findings suggest that financial systems exhibit a *robust-yet-fragile* tendency: while the probability of contagion may be very low, the effects could be extremely widespread should problems occur. The resilience of the system to large shocks in the past is also unlikely to prove a reliable guide to future contagion.

Keywords: Contagion; Network models; Systemic risk; Liquidity risk; Random graphs.

JEL classification: D85; G21

1. Introduction

In modern financial systems, an intricate web of claims and obligations links the balance sheets of a wide variety of intermediaries, such as banks and hedge funds, into a network structure. The recent advent of sophisticated financial products, such as credit default swaps and collateralised debt obligations, has heightened the complexity of these balance sheet connections still further, making it extremely difficult for policymakers to assess the potential for contagion associated with the failure of an individual financial institution or from an aggregate shock to the system as a whole.[1]

The interdependent nature of financial balance sheets also creates an environment for feedback elements to generate amplified responses to any shock to the

---

[*]Bank of England, Threadneedle Street, London EC2R 8AH, U.K. Corresponding author. *E-mail address:* prasanna.gai@bankofengland.co.uk

[†]Bank of England, Threadneedle Street, London EC2R 8AH, U.K.

[1]See Rajan (2005) for a policymaker's view of the recent trends in financial development.

financial system. As Cifuentes et.al (2005) stress, the knock-on effect of an initial default of a financial institution on asset prices can trigger further rounds of default as other financial entities are forced to write down the value of their assets. Contagion due to the direct interlinkages of claims and obligations may, thus, be reinforced by contagion on the asset side – particularly when the market for key financial system assets is illiquid.

Existing theoretical work on contagion relies on simple network structures. Allen and Gale (2000) obtain results using a four-bank network, while Freixas et.al (2000) model a system with money-centre banks where banks on the periphery are linked to the centre, but not to each other. These models focus on direct interlinkages through the interbank market and suggest that contagion is less likely when the structure of claims is "complete" and balance sheets are highly connected, i.e. when every bank has symmetric exposures to each other. By contrast, when the structure of the financial system is "tiered" or asymmetric, with banks only having exposures to a few counterparties, the system is more fragile.

The generality of insights based on small networks with rigid structures to real-world contagion is clearly open to debate. Moreover, the literature largely fails to distinguish the probability of contagious default from its potential impact. The increasing emphasis by policymakers on quantifying the likelihood and potential cost of contagion risk in their financial stability work points to a need for an improved understanding of how financial intermediaries fit together as networks.[2]

This paper develops a general model of contagion in financial systems that explicitly accounts for the nature and scale of aggregate and idiosyncratic shocks, the complexity of network structure, and allows asset prices to interact with balance

---

[2] See Bank of England (2006) for some preliminary attempts to quantify and calibrate the main channels of loss/disturbance to the UK financial system.

sheets. The interactions between financial intermediaries following shocks make for non-linear system dynamics, and our model provides a framework for isolating the *probability* and *spread* of contagion when claims and obligations are interlinked. We find that financial systems exhibit a *robust-yet-fragile* tendency. While greater connectivity reduces the likelihood of widespread default, the impact on the financial system, should problems occur, could be on a significantly larger scale than hitherto. The model also highlights how *a priori* indistinguishable shocks can have very different consequences for the financial system. The resilience of the network to large shocks in the past is no guide to future contagion, particularly if shocks hit the network at particular pressure points associated with underlying structural vulnerabilities.

The intuition underpinning these results is straightforward. In a more connected system, the counterparty losses of a failing institution can be more widely dispersed to, and absorbed by, other entities. So increased connectivity and risk sharing may lower the probability of contagion. But *conditional* on the failure of one institution triggering contagious defaults, a higher number of financial linkages also increases the potential for contagion to spread more widely. In particular, greater connectivity increases the chances that institutions which survive the effects of the initial default will be exposed to more than one defaulting counterparty after the first round of contagion, thus making them vulnerable to a second-round default. The impact of any crisis that does occur could, therefore, be larger.

Our model draws on the mathematics of complex systems.[3] This literature describes the behaviour of connected groups of nodes in a network and predicts

---

[3] See Strogatz (2001) and Newman (2003) for authoritative and accessible surveys. Giesecke and Weber (2004) study contagion using the 'voter' model of interacting particle systems. But their network is constrained to a lattice structure and balance sheets are not defined.

the size of a susceptible cluster, i.e. the number of vulnerable nodes reached via the transmission of shocks along the links of the network. The approach relies on specifying all possible patterns of future transmission. Newman et.al (2001) and Watts (2002) show how probability generating function techniques can identify the number of a randomly selected node's first neighbours, second neighbours, and so on. Recursive equations are constructed to consider all possible outcomes and obtain the total number of nodes that the original node is connected to – directly and indirectly. *Phase transitions*, which mark the threshold(s) for extensive contagious outbreaks can then be identified.

In what follows, we build a realistic financial system and use these techniques to model the spread and probability of default, analytically and numerically. Unlike the generic, undirected graph model of Watts (2002), our model provides an explicit characterisation of balance sheets, making clear the direction of claims and obligations linking financial institutions. It also includes asset price interactions with balance sheets, allowing the effects of asset-side contagion to be clearly de-lineated. We illustrate the robust-yet-fragile tendency of financial systems and analyse how contagion risk changes with capital buffers, the degree of connectivity, and the liquidity of the market for failed banking assets.[4]

The structure of the paper is as follows. Section 2 describes the structure of the financial network, the transmission process for contagion, and analytical results characterising a default cascade. Section 3 uses numerical simulations to study the effects of failures of individual institutions and articulate the likelihood and extent of contagion. Section 4 considers the impact of liquidity effects on system

---

[4] In a recent paper, Alentorn et.al (2007) also use random graph techniques to model contagion in financial systems. Their results are strictly numerical and do not articulate probability or impact. The absence of any analytical underpinning also means that the intuition driving their findings is unclear.

stability. Section 5 discusses points of contact with the empirical literature on interbank contagion being pursued by central banks. A final section concludes.

2. The Model

2.1. Network Structure

Consider a financial network in which n financial intermediaries, 'banks' for short, are linked together by their claims on each other. In the language of graph theory, each bank represents a *node* on the graph and the interbank exposures of bank i define the *links* with other banks. These links are *directed*, reflecting the fact that interbank exposures comprise assets as well as liabilities. Figure 1 shows an example of a directed financial network in which there are five banks.

A crucial property of graphs such as those in Figure 1 is their *degree distribution*. In a directed graph, each node has two degrees, an *in-degree*, the number of links that point into the node, and an *out-degree*, which is the number pointing out. Incoming links to a node or bank reflect the interbank assets/exposures of that bank, i.e monies owed to the bank by a counterparty. Outgoing links from a bank, by contrast, correspond to its interbank liabilities. In what follows, the joint distribution of in- and out-degree governs the potential for the spread of shocks through the network. A feature of our analysis is that this joint degree distribution is entirely arbitrary, though a specific form is assumed in our numerical simulations in section 3.

Suppose that the total assets of each bank are normalised to unity and that these consist of interbank assets, A_i^{IB}, and illiquid external assets, such as mort-

gages, A_i^M. Then

$$A_i^{IB} + A_i^M = 1 \quad \forall \, i. \tag{1}$$

We assume that the total interbank asset position of every bank is evenly distributed over each of its incoming links and is independent of the number of links the bank has. If banks have no incoming links, $A_i^{IB} = 0$.

The normalised liabilities of bank i are composed of interbank liabilities, L_i^{IB}, as well as exogenously given customer deposits, D_i. The condition for bank i to be solvent is therefore

$$(1 - \phi) A_i^{IB} + q A_i^M - L_i^{IB} - D_i > 0, \tag{2}$$

where ϕ is the fraction of defaulted banks with obligations to bank i and q is the resale price of the illiquid asset. The value of q may be less than one in the event of asset sales by banks in default, but equals one if there are no 'fire sales'. We make a zero recovery assumption, namely that when a linked bank defaults, bank i loses all of its interbank assets held against that bank.[5] The solvency condition can also be expressed as

$$\phi < \frac{K_i - (1 - q) A_i^M}{A_i^{IB}}, \text{ for } A_i^{IB} \neq 0, \tag{3}$$

where $K_i = A_i^{IB} + A_i^M - L_i^{IB} - D_i$ is the bank's capital buffer, i.e the difference between the book value of its assets and liabilities.

To model the dynamics of contagion, we suppose that all banks in the network are initially solvent and that the network is perturbed at time $t = 1$ by the initial default of a single bank. Let j_i denote the number of incoming links for bank i.

---

[5] This assumption is likely to realistic in the midst of a crisis: in the immediate aftermath of a default, the recovery rate and the timing of recovery will be highly uncertain and banks' funders are likely to assume the worst-case scenario. Nevertheless, in our numerical simulations, we show that our results are robust to relaxing this assumption.

Since linked banks each lose a fraction $1/j_i$ of their interbank assets when a single counterparty defaults, it is clear that the only way default can spread is if there is a neighbouring bank for which

$$\frac{K_i - (1-q)\, A_i^M}{A_i^{IB}} < \frac{1}{j_i}. \tag{4}$$

We define banks that are exposed in this sense to the default of a single neighbour as *vulnerable* and other banks as *safe*. The vulnerability of a bank clearly depends on its in-degree, j. Specifically, a bank with in-degree j is vulnerable with probability

$$v_j = P\left[\frac{K_i - (1-q)\, A_i^M}{A_i^{IB}} < \frac{1}{j}\right] \quad \forall\, j \geq 1. \tag{5}$$

Further, the probability of a bank having in-degree j, out-degree k and being vulnerable is $v_j \cdot p_{jk}$, where p_{jk} is the joint degree distribution of in- and out-degree.

The model structure described by equations (1) to (5) captures several features of interest in systemic risk analysis. First, the nature and scale of adverse aggregate or macroeconomic events can be interpreted as a negative shock to the stock of illiquid assets, A_i^M, or equivalently, to the capital buffer, K_i. Second, idiosyncratic shocks (e.g fraud shocks) can be modelled by assuming the exogenous default of a bank. Third, the structural characteristics of the financial system are described by the distribution of interbank linkages, p_{jk}. And finally, liquidity effects associated with the potential knock-on effects of default on asset prices are captured by allowing q to vary. To keep matters simple, we initially fix $q = 1$, returning later to endogenise it.

2.2. Generating Functions and the Transmission of Shocks

In sufficiently large networks, for contagion to spread beyond the first neighbours of the initially defaulting bank, those neighbours must themselves have outgoing links (i.e. liabilities) to other vulnerable banks.[6] We therefore define the generating function for the joint degree distribution of a vulnerable bank as

$$\mathcal{G}(x,y) = \sum_{j,k} v_j \cdot p_{jk} \cdot x^j \cdot y^k. \tag{6}$$

The generating function contains all the same information that is contained in the degree distribution, p_{jk}, and the vulnerability distribution, v_j, but in a form that allows us to work with sums of independent draws from different probability distributions. Specifically, for our purposes, it generates all of the moments of the degree distribution of only those banks that are vulnerable. Appendix A provides a detailed description of the key properties of generating functions used in this paper.

Since every interbank asset of a bank is an interbank liability of another, every outgoing link for one node is an incoming link for another node. This means that the average in-degree in the network, $\frac{1}{n}\sum_i j_i = \sum_{j,k} j p_{jk}$, must equal the average out-degree, $\frac{1}{n}\sum_i k_i = \sum_{j,k} k p_{jk}$. We refer to this quantity as the *average degree* and denote it by

$$z = \sum_{j,k} j p_{jk} = \sum_{j,k} k p_{jk}. \tag{7}$$

---

[6] If the number of nodes, n, is sufficiently large, banks are highly unlikely to be exposed to more than one failed bank after the first round of contagion, meaning that safe banks will never fail in the second round. This assumption clearly breaks down either when n is small or when contagion spreads more widely. However, the logic of this section still holds in both cases: in the former, the exact solutions derived for large n will only approximate reality (this is confirmed by the numerical results in section 3); in the latter, the exact solutions will apply but the scale of contagion will be affected, as discussed further in section 2.4.

Using equation (7), the average degree of a vulnerable bank to other vulnerable banks, z_v, can be obtained from the generating function as

$$z_v = \frac{\partial \mathcal{G}}{\partial x} \Big|_{x,y=1} = \frac{\partial \mathcal{G}}{\partial y} \Big|_{x,y=1} = v_j \cdot z. \tag{8}$$

From $\mathcal{G}(x,y)$, we can define a single-argument generating function, $G_0(y)$, for the number of links leaving a randomly chosen vulnerable bank. This is given by

$$\begin{aligned} G_0(y) &= \mathcal{G}(1,y) \\ &= \sum_{j,k} v_j \cdot p_{jk} \cdot y^k. \end{aligned} \tag{9}$$

Note that

$$\begin{aligned} \mathcal{G}(1,1) &= G_0(1) \\ &= \sum_{j,k} v_j \cdot p_{jk} \end{aligned} \tag{10}$$

so that $G_0(1)$ yields the fraction of banks that are vulnerable.

We can also define a second single-argument generating function, $G_1(y)$, for the number of links leaving a bank reached by following a randomly chosen incoming link. Because we are interested in the propagation of shocks from one bank to another, we require the degree distribution, q_{jk}, of a vulnerable bank that is a random neighbour of our initially chosen bank. The larger the in-degree of this second bank, the more likely it is to be a neighbour and lie at the end of a randomly chosen outward link.[7] So the probability of choosing it is proportional to jp_{jk} and

---

[7] See Feld (1991) and Newman (2003) for a detailed analysis of this point.

the corresponding generating function is

$$G_1(y) = \sum_{j,k} v_j \cdot q_{jk} \cdot y^k = \frac{\sum\limits_{j,k} v_j \cdot j \cdot p_{jk} \cdot y^k}{\sum\limits_{j,k} j \cdot p_{jk}}. \tag{11}$$

We now describe the distribution of the *cluster* of vulnerable banks that can be reached by following a randomly chosen directed link, following an initial default. The size and distribution of the vulnerable cluster characterises the spread of default across the financial network. As Figure 2 illustrates, the pattern of transmission can take many different forms. We can follow a randomly chosen directed link and find a single bank at its end with no further outgoing connections emanating from it. This bank may be safe (s) or vulnerable (v). Or we may find a vulnerable bank with one, two, or more links emanating from it to further clusters.

Let $H_1(y)$ be the generating function for the probability of reaching an outgoing cluster of given size by following an outgoing link after an initial default. As shown in Figure 2, the total probability of all possible forms can be represented self-consistently as the sum of probabilities of having only a single node, having a single node connected to one other component, two other components, and so on. More formally,

$$H_1(y) = \Pr[\textit{reach safe bank}] + \Pr[\textit{arrive at vulnerable bank}]$$
$$+ \Pr[\textit{arrive at vulnerable bank with clusters}].$$

If a generating function generates the probability distribution of some property, k, of an object, then the sum of that property over n independent such objects is distributed according to the n^{th} power of the generating function (see Appendix). We can use this property to describe the distribution of cluster sizes as follows:

$$H_1(y) = \Pr[\textit{reach safe bank}] + y \sum_{j,k} v_j \cdot q_{jk} \cdot [H_1(y)]^k, \tag{12}$$

where the leading factor of y accounts for the one vertex at the end of the initial edge. Using equation (11) and noting that $G_1(1)$ represents the probability that a random neighbour of a vulnerable bank is vulnerable, allows us to write equation (12) in recursive form as

$$H_1(y) = 1 - G_1(1) + yG_1(H_1(y)). \tag{13}$$

It remains to establish the distribution of outgoing vulnerable cluster sizes to which a randomly chosen bank belongs. There are two possibilities that can arise. First, a randomly chosen bank may be safe, in which case the size of the vulnerable cluster is zero. Second, it may have in-degree j and out-degree k, and be vulnerable, the probability of which is $v_j \cdot p_{jk}$. In this second case, each link leads to a vulnerable cluster whose size is drawn from the distribution generated by $H_1(y)$. So the size of the vulnerable cluster to which a randomly chosen bank belongs is generated by

$$\begin{aligned} H_0(y) &= \Pr[bank\ safe] + y\sum_{j,k} v_j \cdot p_{jk} \cdot [H_1(y)]^k \\ &= 1 - G_0(1) + yG_0[H_1(y)]. \end{aligned} \tag{14}$$

And, in principle, we can calculate the complete distribution of cluster sizes by solving equation (13) for $H_1(y)$ and substituting the result into equation (14).

2.3. Phase Transitions

Although it is not usually possible to find a closed-form expression for the complete distribution of cluster sizes in a network, we can obtain closed form expressions for the moments of its distribution from equations (13) and (14). In

particular, the average vulnerable cluster size, $\mathcal{S}$, is given by

$$\mathcal{S} = H_0'(1). \tag{15}$$

Noting that $H_1(1) = 1$, it follows from equation (14) that

$$\begin{aligned} H_0'(1) &= G_0[H_1(1)] + G_0'[H_1(1)] H_1'(1) \\ &= G_0(1) + G_0'(1) H_1'(1). \end{aligned} \tag{16}$$

And we know from equation (13) that

$$H_1'(1) = \frac{G_1(1)}{1 - G_1'(1)}. \tag{17}$$

So substituting equation (17) into (16) yields

$$\mathcal{S} = G_0(1) + \frac{G_0'(1) G_1(1)}{1 - G_1'(1)} \tag{18}$$

From equation (18), it is apparent that the points which mark the *phase transitions* at which the average vulnerable cluster size diverges are given by

$$G_1'(1) = 1, \tag{19}$$

or, equivalently, by

$$\sum_{j,k} j \cdot k \cdot v_j \cdot p_{jk} = z \tag{20}$$

where we have used equations (7) and (11).

The term $G_1'(1)$ is the average out-degree of a vulnerable first neighbour, counting only those links that end up at another vulnerable bank. If this quantity is less than one, all vulnerable clusters are small and contagion dies out quickly since the number of vulnerable banks reached declines. But if $G_1'(1)$ is greater than one, a "giant" vulnerable cluster – a vulnerable cluster whose size scales linearly

with the size of the whole network – exists and occupies a finite fraction of the network. In this case, system-wide contagion is possible: with positive probability, a random initial default at one bank can lead to the spread of default across the entire vulnerable portion of the financial network.

As z increases, the $\sum_{j,k} j \cdot k \cdot p_{jk}$ term in equation (20) increases monotonically but v_j falls. So equations (19) and (20) will either have two solutions or none at all. In the first case, there are two phase transitions and a continuous window of (intermediate) values of z for which contagion is possible. For values of z that lie outside the window and below the lower phase transition, the $\sum_{j,k} j \cdot k \cdot p_{jk}$ term is too small and the network is insufficiently connected for contagion to spread (consider what would happen in a network with no links); for values of z outside the window and above the upper phase transition, the v_j term is too small and contagion cannot spread because there are too many safe banks.

2.4.   The Probability and Spread of Contagion

From a system stability perspective, we are primarily interested in contagion within the giant vulnerable cluster. This only emerges for intermediate values of z, and only when the initially defaulting bank is either in the giant vulnerable cluster or directly adjacent to it. The likelihood of contagion is, therefore, directly linked to the size of the vulnerable cluster within the window.[8] Intuitively, near both the lower and upper phase transitions, the probability of contagion must be close to zero since the size of the vulnerable cluster is either curtailed by limited connectivity or by the presence of a high fraction of safe banks. The probability of contagion is thus non-monotonic in z: initially, the risk-spreading effects stemming

---

[8] Note that this is not given by (18) since this equation is derived on the assumption that there are no cycles connecting subclusters. This will not hold in the giant vulnerable cluster.

from a more connected system will increase the size of the vulnerable cluster and the probability of contagion; eventually, however, risk-sharing effects that serve to reduce the number of vulnerable banks dominate, and the probability of contagion falls.[9]

Near the lower phase transition, the *conditional* spread of contagion (i.e. conditional on contagion breaking out) corresponds to the size of the giant vulnerable cluster. But, for higher values of z, once contagion has spread through the entire vulnerable cluster, the assumption that banks are adjacent to no more than one failed bank breaks down. So 'safe' banks may be susceptible to default and contagion can spread well beyond the vulnerable cluster. Therefore, the fraction of banks affected by episodes of contagion will be greater than the probability of contagion breaking out, with the difference being magnified as z increases. Indeed, near the upper phase transition, the system exhibits a robust-yet-fragile tendency, with episodes of contagion occurring rarely, but spreading widely when they do take place.

From equation (20), the size of the contagion window is larger if, for a given j, v_j is larger. Greater levels of vulnerability also increase the size of the giant vulnerable cluster and, hence, the probability of contagion within the range of intermediate z values. Therefore, it is clear from equation (5) that an adverse shock which erodes capital buffers will both increase the probability of contagion and extend the range of z for which contagious outbreaks are possible.

---

[9] In the special case of a uniform (Poisson) random graph in which each possible link is present with independent probability p, an analytical solution for the size of the giant vulnerable cluster can be obtained using techniques discussed in Watts (2002) and Newman (2003). Since this does not account for the possibility of contagion being triggered by nodes directly adjacent to the vulnerable cluster, it does not represent an analytical solution for the probability of contagion. However, it highlights that the size of the giant vulnerable cluster, and hence the probability of contagion, is non-monotonic in z.

3.   Numerical Simulations

3.1.   Methodology

To illustrate our results, we calibrate the model and simulate it numerically. Although the findings generalise to random graphs with arbitrary degree distributions, we assume a (Poisson) random graph in which each possible link in the graph is present with independent probability p. The average degree, z, is allowed to vary in each simulation. And although our model applies to networks of fully heterogeneous financial intermediaries, we take the capital buffers and asset positions on banks' balance sheets to be identical.

As a benchmark, we consider a network of 1000 banks. Clearly, the number of financial intermediaries in a system depends on how the system is defined and what counts as a financial intermediary. But several countries have banking networks of this size, and a figure of 1,000 intermediaries also seems reasonable if we are considering a global financial system involving investment banks, hedge funds, and other players.

The initial assets of each bank are chosen so that they comprise 80% external assets and 20% interbank assets – the 20% share of interbank assets is broadly consistent with the figures for developed countries reported by Upper (2006). Banks' capital buffers are set at 4%, a figure calibrated from data contained in the 2005 published accounts of a range of large, international financial institutions.[10] Since each bank's interbank assets are evenly distributed over its incoming links, interbank liabilities are determined endogenously within the network structure. And the liability side of the balance sheet is 'topped up' by customer deposits until the

---

[10] Further details are available on request from the authors.

total liability position equals the total asset position.

In the experiments that follow, we draw 1,000 realisations of the network for each value of z. In each of these draws, we shock one bank at random, wiping out all of its external assets – this type of idiosyncratic shock may be interpreted as a fraud shock. The failed bank defaults on all of its interbank liabilities. As a result, neighbouring banks may also default if their capital buffer is insufficient to cover their loss on interbank assets. Any neighbouring banks which fail are also assumed to default on all of their interbank liabilities, and the iterative process continues until no new banks are pushed into default.

Since we are only interested in the likelihood and conditional spread of system-wide contagion, we wish to exclude very small outbreaks of default outside the giant vulnerable cluster from our analysis. So when calculating the probability and conditional spread of contagion, we only count episodes in which over 5% of banks default. As well as being analytically consistent on the basis of numerical simulations, a 5% failure rate seems a suitable lower bound for defining a systemic financial crisis.

3.2. Results

Figure 3 summarises the benchmark case. In this and all subsequent diagrams, the scale of contagion measures the fraction of banks which default, *conditional* on contagion over the 5% threshold breaking out.

The benchmark simulation confirms the results and intuition of sections 2.3 and 2.4. Contagion only occurs within a certain window of z. Within this range, the probability of contagion is non-monotonic in connectivity, peaking at approximately 0.8 when z is slightly less than 4. Near the lower phase transition, the

conditional spread of contagion is approximately the same as the frequency of contagion – in this region, contagion breaks out when any bank in the giant vulnerable cluster is shocked and spreads to the entire cluster, but not beyond.

For higher values of z, however, a large proportion of banks in the network fail when contagion breaks out. Of particular interest are the points near the upper phase transition: when $z > 8$, contagion never occurs more than five times in 1,000 draws; but in each case where it does break out, every bank in the network fails. This highlights that *a priori* indistinguishable shocks to the network have vastly different consequences for contagion.

In Figure 4, we compare our benchmark results with the limiting case, since our analytical results only strictly apply in the limit as $n \to \infty$. Watts (2002) notes that numerical results in random graph models approximate analytical solutions in the vicinity of $n = 10,000$. Figure 4 demonstrates that a smaller number of nodes in the benchmark simulation does not fundamentally affect the results: the contagion window is widened slightly, but the qualitative results of the analytical model remain intact.

Figure 5 considers the effects of varying banks' capital buffers. As expected, an erosion of capital buffers both widens the contagion window and increases the probability of contagion for fixed values of z.[11] For small values of z, the extent of contagion is also slightly greater when capital buffers are lower but, in all cases, the scale of contagion reaches one for sufficiently high values of z. When the capital buffer is increased to 5%, however, this occurs well after the peak probability of contagion. This neatly illustrates how increased connectivity can simultaneously reduce the probability of contagion but increase its spread

---

[11] Reduced capital buffers may also increase the likelihood of an initial default. Therefore, they may contribute to an increased probability of contagion from this perspective as well.

conditional on it breaking it out.

Figure 6 illustrates how changes in the average degree and capital buffers jointly affect the *expected* number of defaults in the system. Since this diagram does not isolate the probability of contagion from its potential spread, rare but high-impact events appear in the benign (flat) region as the expected number of defaults in these cases is low. Figure 6 serves to highlight another non-linear feature of the system: when capital buffers are eroded to critical levels, the level of contagion risk can increase extremely rapidly.

Finally, in Figure 7, we relax the zero recovery assumption. Instead, we assume that when a bank fails, its default in the interbank market equals its asset shortfall (i.e. its outstanding loss after its capital buffer is absorbed) plus half of any remaining interbank liabilities, where the additional amount is interpreted as reflecting bankruptcy costs that are lost outside the system.[12] As we might expect, this reduces the likelihood of contagion because fewer banks are vulnerable when the recovery rate can be positive. But it is also evident that relaxing the zero recovery assumption does not fundamentally affect our broad results.

3.3.   Interpretation and Discussion

Contagious crises occur infrequently in developed countries, suggesting that financial systems are located near, or above, the upper phase transition of our model. The findings of Soramaki et al. (2006), who report average degrees in financial systems of 15, are consistent with this. Given that banks' capital buffers are generally set to withstand 99.9% of credit risk shocks, it is not inconceivable

---

[12] Since interbank assets make up 20% of each bank's total asset position, interbank liabilities must, on average, make up 20% of total liabilities. Therefore, for the average bank, the maximum bankruptcy cost under this assumption is 10% of total assets / liabilities, which accords with the empirical estimates of bankruptcy costs in the banking sector reported by James (1991).

that a one-in-a-thousand event might be needed to trigger contagion.

Our framework implies that financial systems exhibit a robust-yet-fragile tendency. Although the likelihood of contagion may be very low, its potential impact could be extremely widespread. Moreover, even if contagion from idiosyncratic shocks never occurs when banks have relatively high capital buffers, Figure 5 highlights that if an adverse aggregate shock, such as a macroeconomic downturn, erodes capital buffers, the system could be susceptible to contagion risk.

A priori indistinguishable shocks also have vastly different consequences in our model. Although the system may be robust to most shocks of a given size, if it is hit by a similarly sized shock at a particular pressure point, possibly reflecting a structural weakness, the ensuing financial instability could be significant. It cautions against assuming that the past resilience of the financial system to large shocks will continue to apply to future shocks of a similar magnitude. Related work by Albert et.al (2000) considers this issue further.

4. Liquidity Risk

We now incorporate liquidity effects into our analysis. When a bank fails, financial markets may have a limited capacity to absorb the illiquid external assets which are sold. As a result, the asset price may be depressed. Following Schnabel and Shin (2004) and Cifuentes et.al (2005), suppose that the price of the illiquid asset, q, is given by

$$q = e^{-\alpha x}, \tag{21}$$

where $x > 0$ is the fraction of system (illiquid) assets which have been sold onto the market (if assets are not being sold onto the market, $q = 1$). We calibrate α so

that the asset price falls by 10% when one-tenth of system assets have been sold.

We integrate this pricing equation into our numerical simulations. Specifically, when a bank defaults, all of its external assets are sold onto the market, reducing the asset price according to equation (21). We assume that when the asset price falls, the external assets of all other banks are marked-to-market to reflect the new asset price. From equation (5), it is clear that this will reduce banks' capital buffers and has the potential to make some banks vulnerable, possibly tipping them into default.

The incorporation of liquidity risk introduces a second potential source of contagion into the model from the asset-side of banks' balance sheets. Note, however, that liquidity risk only materialises upon default. Realistically, asset prices are likely to be depressed by asset sales before any bank defaults. So accounting only for the post-default impact likely understates the true effects of liquidity risk.

Figure 8 illustrates the effects of incorporating liquidity risk into the model. As we might expect, liquidity effects magnify the scale of contagion when it breaks out. The contagion window also widens.

As shown, liquidity effects do not drastically alter the main results of our model. But this should not be taken to mean that liquidity effects are unimportant. In part, the limited effect of liquidity risk reflects the already high spread of contagion embedded in the benchmark scenario. But if a fraction of banks were assumed to be totally immune to counterparty credit risk (i.e. they would survive even if all their counterparties defaulted), then liquidity risk would probably be much more significant in amplifying the scale of contagion. And, to the extent that liquidity risk materialises before any bank defaults, it can be viewed as having the potential to erode capital buffers and increase the likelihood of an initial default.

5.   Relationship to the Empirical Literature

The existing empirical literature has largely tended to use actual or estimated data on interbank lending to simulate the effects of the failure of an individual bank on financial stability.[13]   The evidence of contagion risk from idiosyncratic shocks is mixed. Furfine (2003) and Wells (2004) report relatively limited scope for contagion in the U.S. and U.K. banking systems. By contrast, Upper and Worms (2004) and Van Lelyveld and Liedorp (2006) suggest that contagion risk may be somewhat higher in Germany and the Netherlands. Upper (2006) provides a comprehensive survey of these country studies and highlights their shortcomings.

Contagion due to aggregate shocks is examined by Elsinger et.al (2006) who combine a matrix model of interbank lending in the Austrian banking system with models of credit risk. They take draws from a distribution of risk factors and compute the effects on banks' solvency, calculating the probability and the severity of contagion. Their findings echo the results reported in our paper. While contagious failures are relatively rare, if contagion does occur, it affects a large part of the banking system. The worst case of contagious default affects nearly 40% of the Austrian banking system as measured by the share of failed bank assets in total assets.

Counterfactual simulations have also been used to assess how changes in the structure of interbank loan markets affect the risk of contagion. But these results do not show a clear relationship. Degryse and Nguyen (2005) and Mistrulli (2005) consider how contagion risk has evolved in Belgium and Italy as their banking structures have shifted away from a comparatively complete graph structure to-

---

[13] A parallel literature explores contagion risk in payment systems – see, for example, Angelini et al. (1996).

wards one with multiple money-centre banks. Their findings suggest that whilst this shift appears to have reduced contagion risk in Belgium, the possibility of contagion risk in Italy appears to have increased.

As noted by Upper (2006), existing empirical studies are plagued by data problems and the extent to which reported interbank exposures reflect true linkages is unclear: generally, interbank exposures are only reported on a particular day once a quarter and exclude a range of items, including intraday and secured exposures. As such, they underestimate the true scale of financial connectivity. And studies attempting to analyse the effects of changes in network structure on contagion risk are constrained by short time series for the relevant data series.

6. Conclusion

In this paper, we develop a general model of contagion in financial networks that speaks to the growing financial integration of recent times. Our model applies broadly to systems of agents linked together by their financial claims on each other, including through interbank markets and payment systems. While greater connectivity may reduce the probability of contagion, it could also increase its severity should problems occur. Adverse aggregate shocks and liquidity risk also amplify the likelihood and scale of contagion.

Our results suggest that financial systems may exhibit a robust-yet-fragile tendency. They also highlight how *a priori* indistinguishable shocks can have vastly different consequences – financial market participants and policymakers would be unwise to draw too much comfort from the resilience of the financial system to recent large shocks.

The approach provides a framework for modelling contagion risk when true

linkages are unknown. Added realism to the model can be incorporated by, for example, using real balance sheets for each bank and calibrating the joint degree distribution to match observed data. Extending the model in this direction could help guide the empirical modelling of contagion risk and is left for future work.

Acknowledgments

References

Albert, R, Jeong, H and A-L Barabasi (2000), Attack and error tolerance of complex networks, *Nature*, 406, 378-82.

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

Alentorn, A, Nier, E, Yang, J and T Yorulmazer, Network models and financial stability, *Journal of Economic Dynamics and Control*, forthcoming.

Angelini, P, Maresca, G and D Russo (1996), Systemic risk in the netting system, *Journal of Banking and Finance*, 20, 853-868.

Bank of England (2006), *Financial Stability Report*, 20.

Cifuentes, R, Ferrucci, G and H S Shin (2005), Liquidity risk and contagion, *Journal of the European Economic Association*, 3, 556-566.

Degryse, H and G Nguyen (2005), Interbank exposures: an empirical examination of systemic risk in the Belgian banking system, *Working Paper #43*, National Bank of Belgium.

Elsinger, H, Lehar, A and M Summer (2006), Risk assessment for banking systems, *Management Science*, forthcoming.

Feld, S (1991), Why your friends have more friends than you do, *American Journal of Sociology*, 96, 1464-77.

Freixas, X, Parigi, B and J.C Rochet (2000), Systemic risk, interbank relations and liquidity provision by the central bank, *Journal of Money, Credit and Banking*, 32, 611-38.

Furfine, C (2003), Quantifying the risk of contagion, *Journal of Money, Credit and Banking*, 35, 111-28.

Giesecke, K and S Weber (2004), Cyclical correlations, credit contagion and portfolio losses, *Journal of Banking and Finance*, 28, 3009-3036.

James, C (1991), The losses realized in bank failures, *Journal of Finance*, 46, 1223-1242.

Mistrulli, P (2005), Interbank lending patterns and financial contagion, *mimeo*, Banca d'Italia.

Newman, M (2003), Random graphs as models of networks, in *Handbook of Graphs and Networks*, Bornholdt, S and H G Schuster (eds.), Wiley-VCH, Berlin.

Newman, M, Strogatz, S and D Watts (2001), Random graphs with arbitrary degree distributions and their applications, *Physical Review E*, 64, 026118.

Rajan, R (2005), Has financial development made the world riskier?, paper presented at the *Federal Reserve Bank of Kansas City Economic Symposium on the Greenspan Era: Lessons for the Future*, Jackson Hole, 25-27 August.

Schnabel, I and H S Shin (2004), Liquidity and contagion: the crisis of 1763, *Journal of the European Economic Association*, 2, 929-68.

Soramaki, K, Bech, M, Arnold, J, Glass, R and W Beyeler (2006), The topology of interbank payment flows, *Federal Reserve Bank of New York Staff Report #243*, March.

Strogatz, S (2001), Exploring complex networks, *Nature*, 410, 268-276.

Upper, C (2006), Contagion due to interbank credit exposures: what do we know, why do we know it, and what should we know? Assessing the danger of contagion with counterfactual simulations, *mimeo*, Bank for International Settlements, Basel.

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

Van Lelyveld, I and F Liedorp (2006), Interbank contagion in the Dutch banking sector: a sensitivity analysis, *International Journal of Central Banking*, 2, 99-134.

Watts, D (2002), A simple model of global cascades on random networks, *Proceedings of the National Academy of Sciences*, 99, 5766-5771.

Wells, S (2004), Financial interlinkages in the United Kingdom's interbank market and the risk of contagion, *Bank of England Working Paper #230*.

Appendix: Generating Functions

Let Y be a discrete random variable taking values in $\{0, 1, 2, ...\}$ and let $p_r = P[Y = r]$ for $r = 0, 1, 2...$

Then the *(probability) generating function* of the random variable Y of the distribution, p_r $(r = 0, 1, 2, ...)$, is

$$G(x) = E(x^Y) = \sum_{r=0}^{\infty} x^r P[Y = r] = \sum_{r=0}^{\infty} p_r x^r.$$

Theorem 1 The distribution of Y is uniquely determined by the generating function, $G(x)$.

Proof Since $G(x)$ is convergent for $|x| < 1$, we can differentiate it term by term in $|x| < 1$. Therefore

$$G'(x) = p_1 + 2p_2 x + 3p_3 x^2 + ...$$

and so $G'(0) = p_1$. Repeated differentiation gives

$$G^{(i)}(x) = \sum_{r=i}^{\infty} \frac{r!}{(r-i)!} p_r x^{r-i}$$

and so $G^{(i)}(0) = i! p_i$. Therefore, we can recover $p_0, p_1, p_2...$ from the generating function.

Theorem 2

$$E[Y] = \lim_{x \to 1} G'(x)$$

and, provided that x is continuous at $x = 1$,

$$E[Y] = G'(1).$$

Proof

$$G'(x) = \sum_{r=1}^{\infty} r p_r x^{r-1}$$

Therefore, for $x \in (0, 1)$, $G'(x)$ is a non-decreasing function of x, bounded above

by

$$E[Y] = \sum_{r=1}^{\infty} r p_r.$$

Choose $\varepsilon > 0$ and N large enough that $\sum_{r=1}^{N} r p_r \geq E[Y] - \varepsilon$. Then

$$\lim_{x \to 1} \sum_{r=1}^{\infty} r p_r x^{r-1} \geq \lim_{x \to 1} \sum_{r=1}^{N} r p_r x^{r-1}$$

$$= \sum_{r=1}^{N} r p_r \geq E[Y] - \varepsilon$$

Since this is true for all $\varepsilon > 0$,

$$\lim_{x \to 1} G'(x) = E[Y].$$

Provided that x is continuous at $x = 1$, the second result follows immediately.

Theorem 3

$$E[Y(Y-1)] = \lim_{x \to 1} G''(x)$$

and, provided that x is continuous at $x = 1$,

$$E\left[Y\left(Y-1\right)\right] = G''\left(1\right)$$

Proof

$$G''\left(x\right) = \sum_{r=2}^{\infty} r\left(r-1\right) p_r x^{r-2}$$

and the remainder of the proof is the same as the proof of theorem 2.

Theorem 4 If $Y_1, Y_2, ..., Y_n$ are independent random variables with generating functions $G_1\left(x\right), G_2\left(x\right), ..., G_n\left(x\right)$, then the generating function of $Y_1 + Y_2 + ... + Y_n$ is $G_1(x) \cdot G_2(x) \cdot ... \cdot G_n\left(x\right)$.

Proof

$$E\left[x^{Y_1+Y_2+...+Y_n}\right] = E\left[x^{Y_1} \cdot x^{Y_2} \cdot ... \cdot x^{Y_n}\right] \tag{22}$$

Since $Y_1, Y_2, ..., Y_n$ are independent random variables, the standard result from probability theory that functions of independent random variables are also independent implies that $x^{Y_1}, x^{Y_2}, ..., x^{Y_n}$ are independent. Therefore, using the properties of expectation, we can rewrite (22) as

$$\begin{aligned} E\left[x^{Y_1+Y_2+...+Y_n}\right] &= E\left[x^{Y_1}\right] \cdot E\left[x^{Y_2}\right] \cdot ... \cdot E\left[x^{Y_n}\right] \\ &= G_1(x) \cdot G_2(x) \cdot ... \cdot G_n\left(x\right). \end{aligned}$$
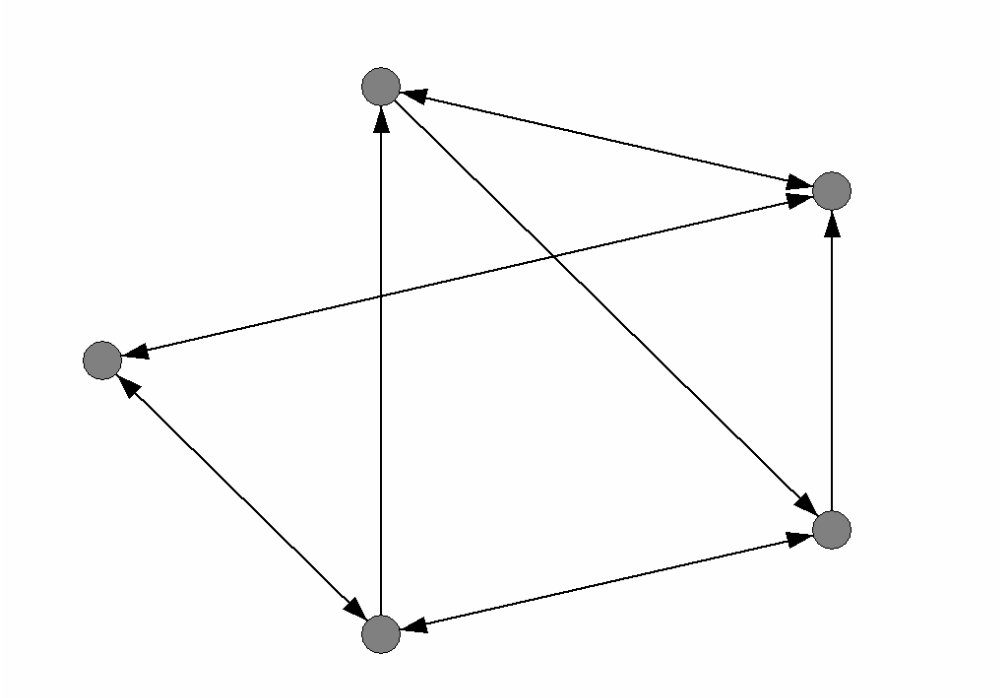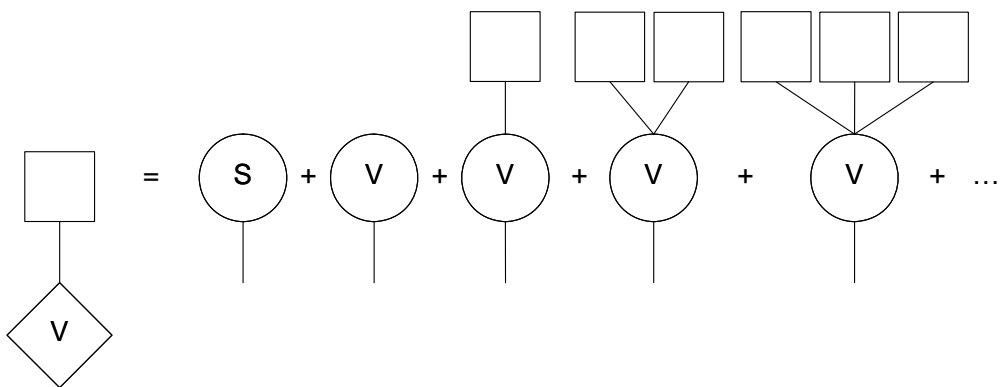
Figure 1: A Directed Network with Five Nodes

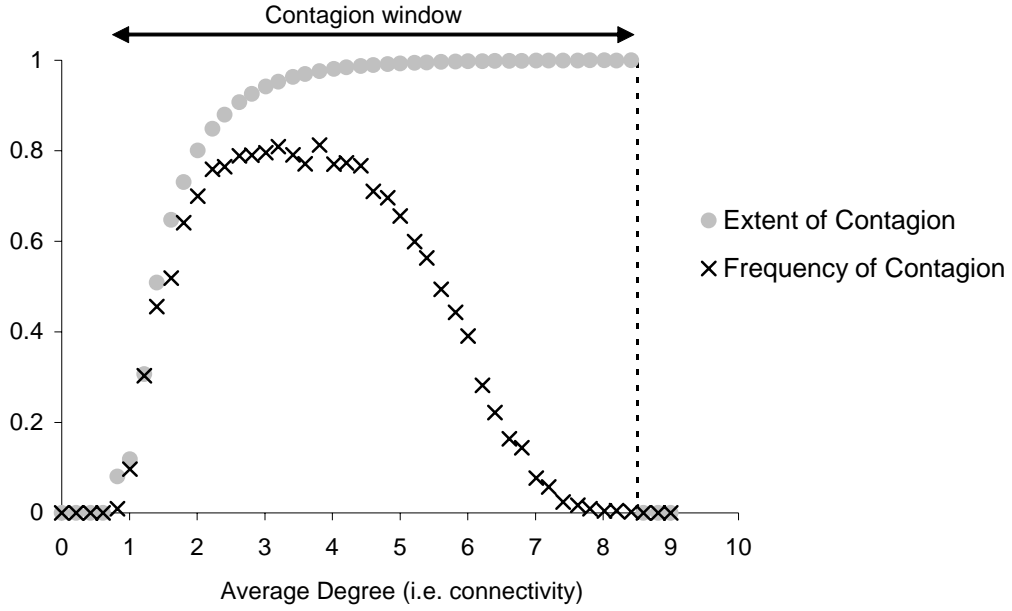Figure 2: Transmission of Contagion implied by Equation (12)
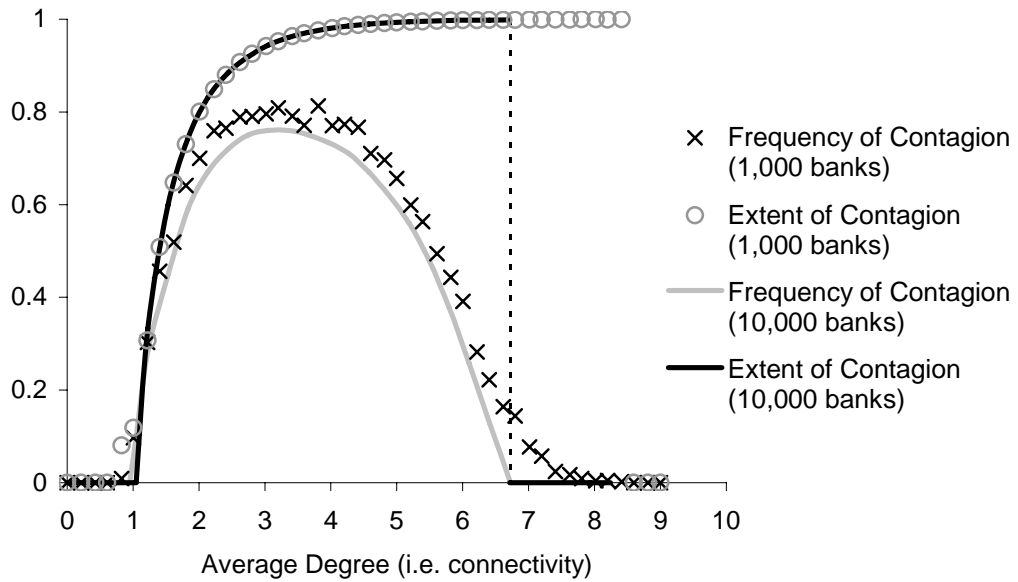
Figure 3: The Benchmark Case

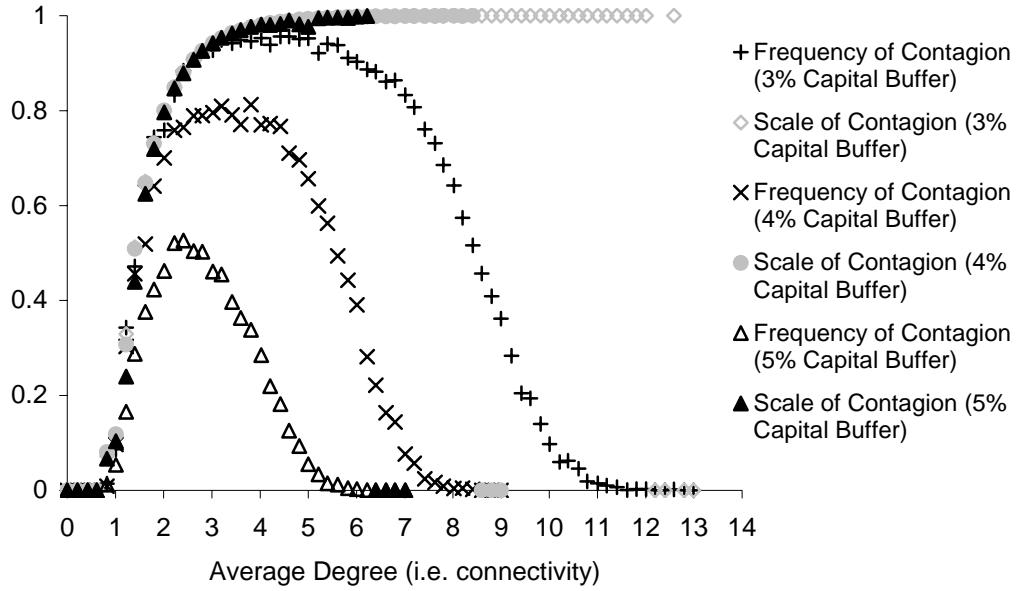Figure 4: Benchmark and Analytical Solutions Compared

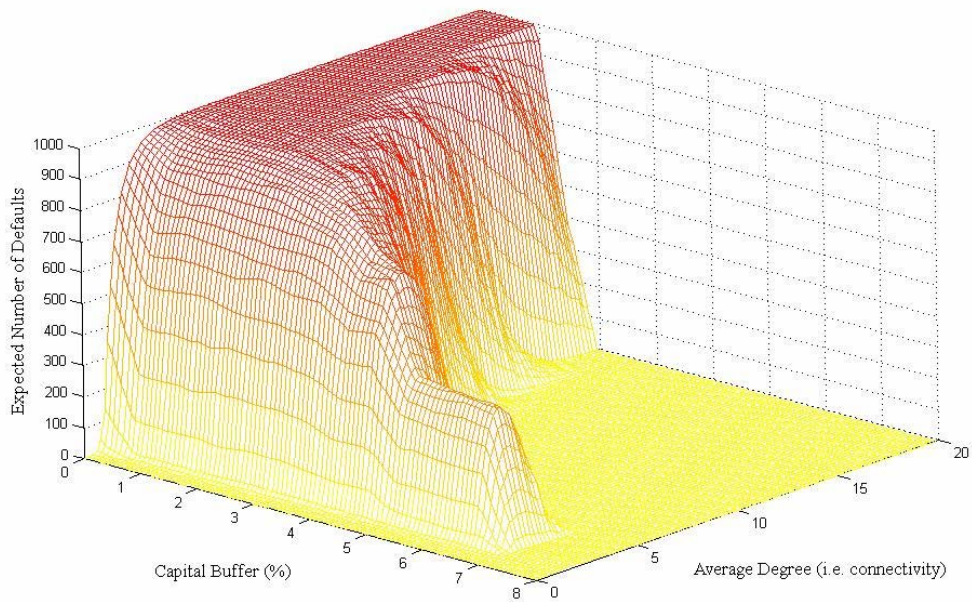Figure 5: Varying the Capital Buffer

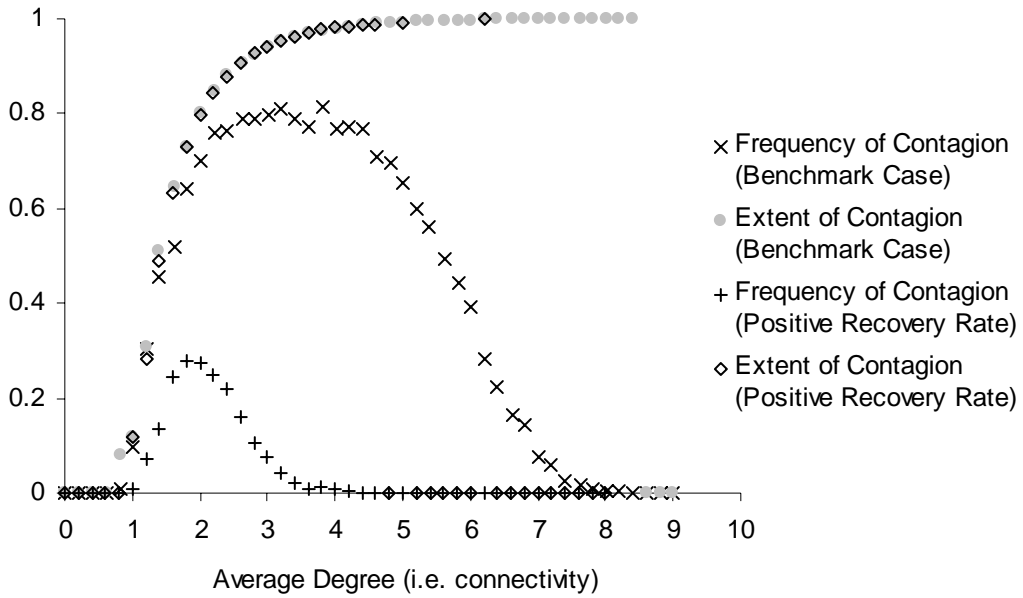Figure 6: Connectivity, Capital Buffers and the Expected Number of Defaults
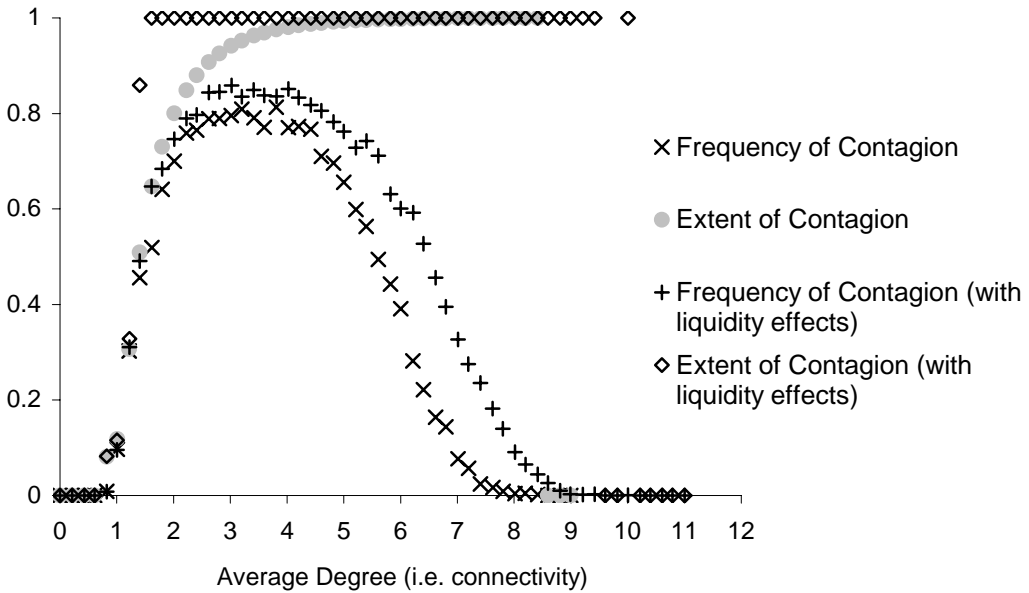
Figure 7: Relaxing the Zero Recovery Rate Assumption

Figure 8: Liquidity Effects and Contagion