ELSEVIER

Journal of Economic Behavior & Organization
Vol. 61 (2006) 525–542

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Systemic risk on the interbank market

Giulia Iori [a,*], Saqib Jafarey [a], Francisco G. Padilla [b]

[a] *Department of Economics, City University, Northampton Square, London EC1V 0HB, UK*
[b] *Department of Mathematics, Kings College, Strand, London WC2R 2LS, UK*

Received 3 September 2003; accepted 21 July 2004
Available online 18 July 2006

Abstract

We simulate interbank lending. Each bank faces fluctuations in liquid assets and stochastic investment opportunities that mature with delay, creating the risk of liquidity shortages. An interbank market lets participants pool this risk but also creates the potential for one bank's crisis to propagate through the system. We study banking systems with homogeneous banks, as well as systems in which banks are heterogeneous. With homogeneous banks, an interbank market unambiguously stabilizes the system. With heterogeneity, knock-on effects become possible, but the stabilizing role of interbank lending remains so that the interbank market can play an ambiguous role.

JEL classification: E5; G21

Keywords: Systemic risk; Contagion; Interbank lending

1. Introduction

In banking systems, there is a tendency for crises to spread from institution to institution. This tendency is referred to as 'systemic risk' and, at a large enough scale, leads to 'systemic failure'. Three types of sources for systemic failure have been noted in the literature. First, there can be a 'bank run', a self-fulfilling panic in which depositors and investors attempt to withdraw funds leading to a collapse of the system (see Diamond and Dybvig, 1983; Donaldson, 1992; de Bandt, 1999; Jacklin and Bhattacharya, 1988; Calomiris and Kahn, 1996; Cowen and Kroszner, 1989). Second, given that banks invest in similar types of assets, a large enough failure by one bank can lead to a fall in the price of its assets and affect the solvency of other banks that hold the same

* Corresponding author. Tel.: +44 20 7040 4575; fax: +44 20 7040 8580.
E-mail addresses: g.iori@city.ac.uk (G. Iori), s.s.jafarey@city.ac.uk (S. Jafarey), francisco.padilla@mth.kcl.ac.uk (F.G. Padilla).

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2004.07.018

asset (see, e.g., Radelet and Sachs, 1998, or Edison et al., 2000; Luangaram and Miller, 1998, and Allen and Gale, 2000). The third aspect of systemic risk arises from inter-locking exposures among financial institutions, which form a basis for mutual insurance, on one hand (see Allen and Gale) and on the other, create the potential for one institution's failure to have 'knock on' effects on the financial health of other institutions.

The focus of this paper is on the third source of systemic risk. In particular, we look at a specific form of inter-locking exposure, viz. interbank lending. Such lending is very short term, mainly overnight, and allows banks facing liquidity shortages to cover their obligations and reserve requirements by borrowing from surplus banks. Thus, interbank lending represents one form of safety net for individual banks. At the same time, by creating inter-locking exposure, it has the potential to create knock-on effects from the failure of debtor banks to the balance sheets and reserve holdings of lender banks. Thus, there is a tradeoff between mutual insurance and systemic risk on the overall stability of the system under interbank lending.

The empirical dimensions of knock-on effects have been studied in Humphrey (1986), Angelini et al. (1996), Furfine (2003) and Upper and Worms (2002). The first two papers study knock-on possibilities associated with intra-day netting arrangements between banks, Humphrey for the United States and Angelini et al. for the Italian netting system. Furfine studies the bilateral exchange of interbank credit in the federal funds market of the United States. Upper and Worms analyse the German interbank market. Each of the papers estimates the impact of failure in one bank upon the financial health of the system as a whole. A common theme of their results is that the degree of contagion depends on the size of losses imposed by failing debtor banks on creditor banks in the system, but the empirical focus of these papers is on the size of knock effects to the neglect of the mutual insurance role.

The full tradeoff between individual and system-wide risk is explicitly considered in Allen and Gale and, more recently, Kahn and Santos (2005), although neither paper focuses exclusively on the interbank market. Moreover, these two papers are mainly interested in studying strategic behaviour by banks and the resulting optimality of mutual insurance arrangements between them.

In this paper, by contrast, we are explicitly concerned with the potential for the interbank market to act as a *propagation* mechanism for liquidity crises. Given that such a role is by no means unambiguous, we believe that it is worth studying for its own sake and in abstraction from the reasons for the onset of crises or the microeconomic details of individual behaviour. While these are undeniably important issues, they are not the primary concern of this paper. Rather, the focus is on how the interplay between individual risk characteristics and the interbank market leads to accentuating (or dampening, for that matter) liquidity crises.

In this sense, our focus can be said to be on microstructure rather than on microeconomics. The strategies of the participants are imposed exogenously and the choice between them is not derived from first principles, but given these elements, the impact of different patterns of stochastic properties and bank interdependence is evaluated through numerical simulations. The crises that are generated in our model are not meant to be realistic in their origins, their duration, or their severity; what matters is how these crises vary *relative to each other* under different statistical scenarios. The point of reference for one crisis will be others produced by the same model using different risk characteristics. This approach is related to a small but growing literature that uses techniques from statistical mechanics to study the dynamics of economic and financial systems (see, for example, Bak et al., 1993; Lux and Marchesi, 1999; Iori, 2002; Cont and Bouchaud, 2000).

Our results suggest that interbank lending contributes to a lower incidence of bank failures through the mutual insurance role but, at the same time, does create the tendency for the system

to display avalanches (i.e., episodes of multiple bank collapse). The avalanches observed in our model indicate the presence of both direct and indirect forms of contagion. The direct effects arise because the failure of one bank has knock-on effects on its creditors, while the indirect effect suggests that bank failures tend to weaken the system and drive it to an unstable state in which it becomes susceptible to further simultaneous failures. Both possibilities are quite minimal when the interbank market consists of homogeneous banks, but their likelihood increases with heterogeneity. We also find that the holding of large liquid reserves, while unambiguously stabilising in the case of isolated banks, can lead to reducing the risk sharing role of the interbank market and can thus make banks more unstable when they are linked through such a market.

Section 2 describes the basic model under the assumption that banks are homogeneous. Section 3 discusses the results of simulating the model and then discusses extensions that introduce heterogeneity into the system. Section 4 concludes.

2. The basic model

The system operates in discrete time, which is denoted by $t = 0, 1, 2 \ldots$. At any time, t, there are a finite number of functioning banks, M_t. Each of these is labelled by k, $k = 1, 2, \ldots, M_t$.

The primary purpose of banks is to channel funds received from depositors towards productive investment. Investment opportunities come from entrepreneurs based in the non-financial sector. These opportunities are assumed to be stochastic and bank-specific in the sense that one bank cannot undertake the opportunity offered to another. It is assumed that resources are illiquid while tied up in investment. An investment made at time t matures at a future time $t + \tau$, although a return, ρ, is realised at each point, $t + 1, t + 2, \ldots t + \tau$. ρ is exogenous and, in all the cases discussed in this paper, risk-free.

Each bank receives stochastic shocks to its liquid reserves, arising ultimately from the deposit and withdrawal patterns of customers. We interpret these shocks broadly to encompass both cash deposits and withdrawals as well as electronic interbank transfers. Since liquidity fluctuations are unpredictable, a bank may find itself unable to meet payment obligations due to the illiquidity of its available assets.

If no interbank market is present, the mere inability to meet customer demands triggers off failure. If interbank lending is possible, an illiquid bank might seek funds not just to make payments but also to repay past creditors. If despite such efforts, a bank ends up with insufficient funds, we assume for simplicity that it too closes down. Credit linkages between banks are defined by a connectivity matrix, J_{ij}. J_{ij} is either one or zero; a value of one indicates that a credit linkage exists between banks i and j and zero indicates no relationship. J_{ij} are randomly chosen at the beginning of the simulation. c represents the probability that J_{ij} is one for any two banks. At one extreme, $c = 0$ represents the case of no interbank lending, while $c = 100$ (percent) represents a situation in which all banks can potentially borrow from and lend to each other.

We now describe the system in detail. We shall first describe the case where the interbank market does not exist at all ($c = 0$) and then discuss the case where such a market exists ($c > 0$).

In the initial period, $t = 0$, the economy starts with a finite number, M_0, of banks. The sequencing of events within each period is as follows. At the start of each period, each bank inherits an amount, L^k_{t-1}, of liquid assets from the previous period, which consists of

$$L^k_{t-1} = A^k_{t-1} + V^k_{t-1} - \sum_{s=1}^{\tau} I^k_{t-s}$$

where A^k denotes deposits held by the general public in bank k, V^k represents its equity, and I^k represents investment by bank k. At $t = 0$, the values of A^k_{-1}, V^k_{-1}, I^k_{-1}, I^k_{-2} and so on are chosen exogenously.

Given A^k_{t-1}, each bank pays an exogenous amount of interest, $r_a \geq 0$, at the start of period t and, at the same time, receives income, $\rho \sum_{s=1}^{\tau-1} I^k_{t-s}$ and $(1+\rho)I^k_{t-\tau}$ from investments made over the last τ periods. It then faces a stochastic pattern of withdrawals and new deposits from customers, resulting in a new value, A^k_t.

Instead of strictly controlling aggregate deposits, we assume that A^k_t are idiosyncratic random variables whose exact specification will be varied through the paper and will be described as relevant in the following sections.

Since the amount of liquidity available to a bank can vary within a period, we shall use $\hat{L}_t$ and $\tilde{L}_t$ to represent intra-period liquidity at various points within each period and L_t to represent end-of-period liquidity. After paying interest, receiving investment income and facing deposit shocks, intra-period liquidity is

$$\hat{L}^k_t = L^k_{t-1} + (A^k_t - A^k_{t-1}) - r_a A^k_{t-1} + \rho \sum_{s=1}^{\tau} I^k_{t-s} + I^k_{t-\tau}. \tag{1}$$

In the absence of an interbank market, any bank with $\hat{L}^k_t < 0$ is unable to make good at least part of its current obligations to depositors. It then shuts down and is removed from the system.

If $\hat{L}^k_t \geq 0$, the bank survives. Such a bank can undertake dividend payments to shareholders along with fresh investment activity. Both dividend payments and investment require liquidity, and dividends are paid before investment takes place.

Dividends are paid out of "excess" returns. In practice, $\hat{E}^k_t = \hat{V}^k_t / A^k_t$ is calculated for each surviving bank. $\hat{V}_t$ is middle-of-period net worth and is itself defined by

$$\hat{V}^k_t = \hat{L}^k_t + \sum_{s=1}^{\tau-1} I^k_{t-s} - A^k_t.$$

All banks with $\hat{E}^k_t > \chi$, where χ is a target capital:deposit ratio, are chosen as candidates to pay dividends. The actual dividend paid, D^k_t, is given by

$$D^k_t = \max\left[0, \min\left[\rho \sum_{s=1}^{\tau} I^k_{t-\tau} - r_a A^k_{t-1}, \hat{L}^k_t - R^k_t, \hat{L}^k_t + \sum_{s=1}^{\tau-1} I^k_{t-s} - (1+\chi)A^k_t\right]\right],$$

where R^k_t is the minimum reserve kept by bank k and is calculated as $R^k_t = \beta A^k_t$. The above formulation restricts dividend payments to equal investment income net of interest payments (if positive), but not if that violates either reserve or capital:deposit targets. Linking dividends to the capital:deposit ratio also has the advantage that the capital of surviving banks grows in proportion to their average deposits.

After dividends have been paid, the bank is assumed to undertake investment on the basis of its available liquidity, on the one hand, and its stochastic investment opportunity, on the other. Liquid reserves at this point are defined as $\tilde{L}^k_t = \hat{L}^k_t - D^k_t$. The amount available to invest is therefore $\tilde{L}^k_t - \beta A^k_t$.

We assume that at any time, t, each bank receives a random investment opportunity, ω^k_t. This opportunity defines the maximum possible investment it can undertake. The actual investment

satisfies

$$I_t^k = \min[\max[0, \tilde{L}_t^k - \beta A_t^k], \omega_t^k].$$

The end-of-period liquidity holding of the bank then becomes

$$L_t^k = \tilde{L}_t^k - I_t^k.$$

The end-of-period net worth of the bank is defined by

$$V_t^k = L_t^k + \sum_{s=0}^{\tau-1} I_{t-s}^k - A_t^k.$$

This completes the description of the system without an interbank market. We now discuss the system with lending.

In this case, at the start of each period, the amount, L_{t-1}^k, of liquid holdings held by each bank consists of

$$L_{t-1}^k = A_{t-1}^k + B_{t-1}^k + V_{t-1}^k - \sum_{s=1}^{\tau} I_{t-s}^k,$$

where B_{t-1}^k denotes total borrowing by bank k from other banks at time $t-1$. Note that B^k can be negative or positive and satisfies $\sum_{k=1}^{M} B^k = 0$. Borrowing is assumed to consist of one-period loans that require repayment in full in the period after which they are undertaken. From here onwards, we shall economise on notation and describe things qualitatively for the most part.

As before, at the very outset of period t, each bank's liquid resources change due to receipt of income from investment, payment of interest to depositors and a stochastic shock to deposits. The difference is that at this early stage, if liquid holdings become negative the bank can issue negotiable debt certificates to cover any excess of payments over its liquid reserves. However, the certificates have to be redeemed at the end of the period through borrowing from other banks. If this is not done, the bank fails and its debt certificates become worthless.

Each bank's next priority is to repay creditors, if any, from the previous period. Such payments have to be made in cash. For the purpose of simulations, it is assumed that a bank either immediately pays all it owes, if it can, or pays nothing at this stage. Hence, only banks with cash holdings in excess of their debt obligations (which equal $(1 + r_b)B_{t-1}^j$ for bank j) make payments. Here, r_b is the interest rate on interbank borrowing. Repayments go directly to the creditors from the past period. The intra-period cash holdings of the relevant banks are updated accordingly.

At this point, two types of banks can be distinguished, those with positive cash and those with negative cash. Accordingly, they are classified as potential lenders and potential borrowers. Note that banks that are potential borrowers at this stage could themselves be owed money by debtors from the past period. Borrowing banks make demands for loans equaling their debt obligations (interest plus principal) minus their current cash (which if negative adds to their demand). In this formulation, borrowing on the interbank market is restricted to short-term solvency needs (i.e., to repay depositors and creditors) and does not cover 'long-term' investments.

Surplus banks are assumed to give priority to dividend payments and investment. The latter assumption implies that arbitrage between the real sector and the interbank market is imperfect. One justification for this would be the specialised nature of bank investments in the real sector. Banks are more likely to possess *privileged* information about their non-bank borrowers than

about other banks within the system; hence, they are more likely to earn economic rents from lending to non-bank entrepreneurs than to other banks.[1] As a result, the returns on 'real' assets dominate those on interbank loans.

After investment, any excess left over is made available to other banks. There is admittedly no mechanism in the model to guarantee that the total demand for interbank loans equals the total supply.[2] Hence, as far as the simulation is concerned, it is assumed that each borrowing bank contacts lending banks in a random order, subject to the condition that borrower and lender are linked in an interbank relationship.

Once a borrowing bank contacts a lending bank, an agreement is reached between the two about how much credit will be exchanged. This is the minimum of the two banks' respective demand and supply. The bank that is left with an unfulfilled trade stays on and contacts another bank on the opposite side of the market and tries reaching further agreements. A borrowing bank does not receive actual funds until it has lined up enough credit to ensure that it will not fail during the current period. Once a bank has obtained sufficient credit, funds are transferred and the cash positions of all banks involved is updated. This continues until either all loanable funds are exhausted or all demands for credit are satisfied.

The process now reiterates itself through the following steps: banks that had not repaid creditors in the first round but have now borrowed enough liquid resources to pay off past debt entirely, do so; these payments go to their creditors and money holdings are updated accordingly; potential borrowers and lenders are determined for the next round; lenders undertake further investment, assuming they had unfulfilled investment opportunities from the previous round. Finally, a fresh round of borrowing and lending takes place. In principle, the process repeats itself until reiteration produces no further exchange of credit.

All banks that are left with negative holdings of liquid resources or holdings that fall short of their remaining debt obligations are deemed to be in default. These banks are removed from the system. Their remaining assets (upon liquidation) are distributed to depositors. If, afterwards, there is still some value left, this goes to creditors from the previous period. Each creditor's net worth is directly reduced by the amount owed to it by the failing bank minus the amount recovered.

3. Simulations and results

We started each simulation with 400 banks. In all the simulations presented here, the return on investment, ρ, was set at 1 percent, the interest rate on interbank borrowing, r_{b}, was 0.5 percent and the interest rate paid to depositors was $r_{\mathrm{d}} = 0$. χ, the capital:deposit ratio, was 30 percent. Parameters that vary across the simulations are identified separately.

Each bank, labeled by an index $k = 1, 2, \ldots, M$, is characterized by its size s^k. This size represents the initial amount of customer deposits. In the homogeneous case it is assumed that all banks have the same size and $s^k = \bar{A}$. In the heterogenous case banks' sizes are sampled from the positive side of a Normal distribution.

[1] This is not say that banks have *more* information about their clients than about their counterparts, just that the latter information is likely to be shared between banks and thus not a significant source of rents.

[2] This arises from the assumed constancy of interest rates. In practise, interbank markets experience considerable fluctuations in interest rates that would tend to dampen shortages. Nonetheless the demand for funds from an otherwise failing bank is likely to be inelastic, on the one hand, and on the other, such a bank might not find creditors at any interest rate. This is what we attempt to capture in the present paper. The limitations that arise as a result are discussed in Section 4.

3.1. Homogeneous case

In the initial set of simulations, we assumed that the banks were identical in the stochastic sense (i.e., that $s^k = \bar{A}$, $\forall k$ and $\bar{\omega}^k = \bar{\omega}$, $\forall k$). Each bank holds initial deposits, $A_0^k = \bar{A}$, worth 1000 units and initial equity, V_0, equal to 0.3 times the initial deposit. At the initial period each bank is assumed to have already made investments. The maturity period is set at $\tau = 3$.

Within this framework, we consider two alternative models of how bank deposits fluctuate over time. In the first one, model A, we take the traditional view that fluctuations are caused by random but mutually uncorrelated withdrawals of depositors. Under the assumption of independence, fluctuations in deposits are proportional to the square root of their size

$$A_t^k = \left| \bar{A} + \sigma_A \sqrt{\bar{A}} \epsilon_t \right|, \tag{2}$$

$$\omega_t^k = |\bar{\omega} + \sigma_\omega \eta_t| . \tag{3}$$

An alternative, model B, is to assume that local correlations exist in the behaviour of each given bank's depositors. These may arise from the fact that depositors of a given bank can observe and mimic each other's withdrawal pattern. In this case, banks will experience fluctuation in their deposit proportional to their mean size

$$A_t^k = |\bar{A} + \bar{A}\sigma_A \epsilon_t|, \tag{4}$$

$$\omega_t^k = |\bar{\omega} + \sigma_\omega \eta_t|. \tag{5}$$

In both cases $\epsilon_t \sim N(0, 1)$, $\eta_t \sim N(0, 1)$ and $\bar{\omega} = \delta\bar{A}$, with $0 < \delta < 1$. Note that this herding behaviour only acts locally, within each bank, so that fluctuation of deposits are still uncorrelated across different banks.

Since the purpose of the exercise was to study the dynamics of a self-contained system with a given initial number of banks, we excluded the possibility that failing banks would be replaced by new entrants. Along with the homogeneity of banks this led to a system that, depending on the average size of investment opportunities, was either highly stable with practically no banks failing or a system in which all banks failed in the long run. For this reason, we use as our measure of relative stability, not the total long-run incidence of bank failures, but the number of failures over any given period of time.

Both model A and B generate qualitatively similar results, so we report simulations only for model B.

Fig. 1 compares bank failures with different degrees of linkage in the interbank market. Linkage is described by a percentage, c. A linkage of c percent implies that each bank can exchange credit with c percent of other banks in the system. The vertical axis shows the number of surviving banks, and the horizontal axis measures time (up to $t = 1000$). Other parameters were $\beta = 0.2$, $\sigma_\omega = \sigma_A = 0.5$. Increasing linkage adds stability in the sense that in any period, more banks survive when the degree of linkage is greater. Indeed, with 100 percent linkage, the system becomes entirely stable. This is consistent with Allen and Gale's result on the risk sharing role of interbank lending. This pattern was very robust to changing the parameters of the simulation.

A second question concerned the role of reserve requirements. Fig. 2 shows how different reserve ratios affect the incidence of surviving banks for the case of no interbank credit market. Of parameters other than β, only σ_A was changed to 0.25 while the others were set as in Fig. 1. As the reserve ratio, β, increases from 10 to 70 percent in increments of 20 percent, the incidence of bank failures clearly falls. This conforms to conventional wisdom regarding the role of reserves.

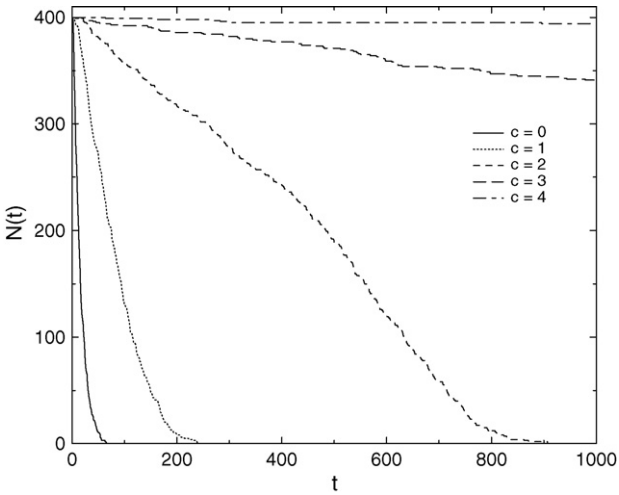

Fig. 1. Surviving homogeneous banks with $\sigma_A = 0.5$, $\sigma_\omega = 0.5$, $\beta = 20$ percent and different interbank linkages.

Fig. 3 depicts the same experiment as Fig. 2, except that interbank linkages are set at 1 percent. With an interbank market, the role of reserve requirements is less clear cut. In both cases, increasing β initially leads to an increase in the incidence of bank failures, but if β crosses a critical level, increasing it further results in fewer bank failures.

We explain these results as follows. Holding reserves keeps resources liquid and enables banks to meet unpredictable shifts in demand by depositors more effectively. Obviously, increasing reserves adds to the stability of individual banks. However, with interbank linkages, higher reserves also reduce the insurance provided by banks to each other. If reserves are set high enough relative to the degree of risk, the individual effect completely dominates and overall stability increases with reserves. Indeed, in this regime our simulations detected no activity on the interbank market since the need to keep high reserves mopped up all the excess liquidity at the disposal of surplus banks.

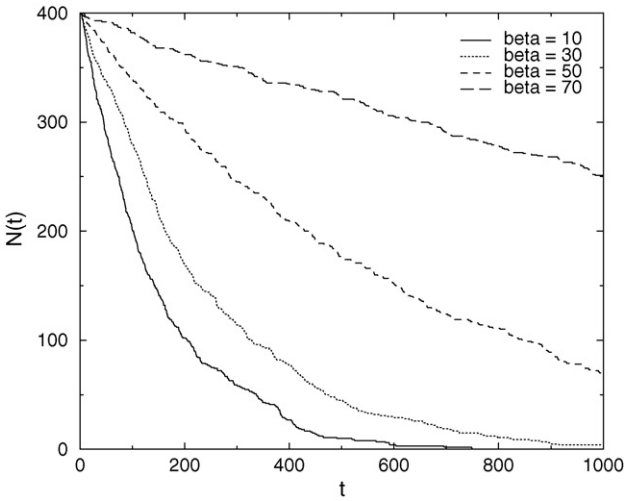

Fig. 2. Surviving homogeneous banks with $\sigma_A = 0.25$, $\sigma_\omega = 0.5$, different reserves and no interbank linkage.

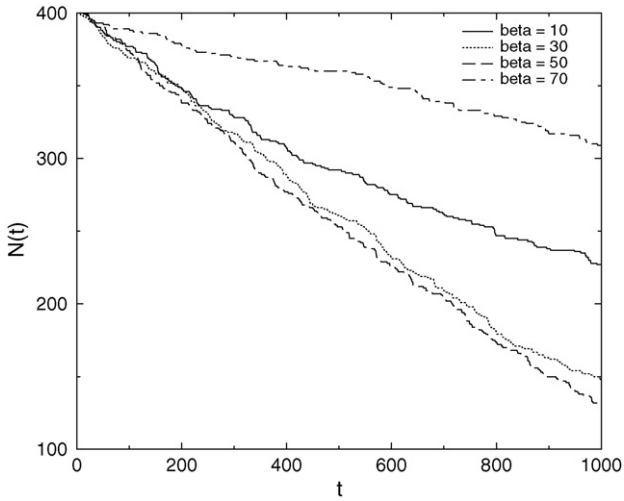

Fig. 3. Surviving homogeneous banks with $\sigma_A = 0.25$, $\sigma_\omega = 0.5$, different reserves and 1 percent linkage.

At lower values of β, reserves do not kill off interbank activity, and the individual stabilising effect is relatively small compared to the insurance effect. It is in this regime that increasing reserves leads to more frequent failures.

Changing the other parameters of the system had predictable effects. For example, increasing the volatility on deposits increased the rate of bank failures, while increasing the recovery rate from failed banks reduced the rate at which banks failed. However, the qualitative nature of the effects reported in Figs. 1–3 was not affected by varying these other parameters. These comparative static experiments will therefore not be discussed in further detail.

We now turn to the issue of contagious failures. Contagion would suggest that 'avalanches', periods in which many banks collapse together, occur. There appears to be no evidence of this in Fig. 1 where the time path of surviving banks declines more or less smoothly over time.

We also simulated a variant of the above model. Instead of letting the *aggregate* levels of deposits and investment opportunities average out on the basis of the law of large numbers, we controlled them deterministically. Aggregate investment opportunities were assumed to be a fixed constant over time. Aggregate deposits were fixed initially and subsequently reduced appropriately whenever a failing bank left behind obligations that exceeded the liquidation value of its remaining assets. Individual banks received stochastic shares of the aggregate deposits and investment opportunities.

In the variant, the basic results described by Figs. 1–3 were also observed, but in the counterparts to Figs. 1 and 3, whenever the interbank market was present, instead of a relatively uniform decline in the number of banks over time, there was an initial period of slow decline followed by a subsequent sharp decline. On a casual glance, this appeared to be a sign of contagious failure, but it turned out that the sharp decline was driven by the cessation of all activity on the interbank market rather than via knock-on effects.[3] When we tried a third variant in which aggregate quantities

[3] Essentially, as the number of surviving banks fell, surviving banks were able to get larger and larger *per capita* opportunities to invest, which eventually absorbed all their available deposits. This feature has been ruled out in the present model where *per capita* values of both deposit and investment opportunity are constant by construction.

remained deterministic but both the aggregate deposit and the investment opportunity shrank proportionately as banks failed, activity on the interbank market did not dry up, and the path of decline remained roughly uniform, as in the present version.

We report these results in order to emphasise that despite having explored several formulations of the stochastic risk facing individual banks, two properties were consistently observed: (i) increasing inter-bank linkages slows down the rate at which banks fail, and (ii) there are no significant 'knock-on' effects from the failure of individual banks.

One feature of the results shown in Figs. 1 and 3 was that no single bank became a significant debtor, and hence, the knock-on effects of any bank's failure did little damage to other banks. The homogeneity of banks appears to be a cause of this. The implication is that with homogeneous banks, the risk sharing role of the interbank market is dominant while the absence of significant amounts of borrowing or lending on an individual basis minimises the potential for knock-on effects within the system.

We therefore made a departure from the homogeneous case by introducing different types of banks.

3.2. Heterogeneous case

The first manner in which we pursued heterogeneity was by letting banks vary in size. Each bank's average deposit was chosen from a Gaussian distribution. Other characteristics, such as the average investment opportunity, were scaled in proportion to this parameter. We again consider two models of each bank's deposit fluctuations over time. In model A,

$$A_t^k = \left| s^k + \sigma_A \sqrt{s^k} \epsilon_t \right|, \tag{6}$$

$$\omega_t^k = |\bar{\omega} + \sigma_\omega \eta_t|. \tag{7}$$

Note that, with different sizes, in model A, the deposits of bigger banks will have a lower volatility:mean ratio than those of smaller banks. In model B,

$$A_t^k = |s^k + \sigma_A s^k \epsilon_t|, \tag{8}$$

$$\omega_t^k = |\bar{\omega} + \sigma_\omega \eta_t|. \tag{9}$$

In model B, all bank deposits have the same volatility:mean ratio. Hence, banks differ not in terms of the idiosyncratic risk arising from depositors' demands, but in terms of the size of their borrowing and lending requirements. Big banks are capable of generating both larger demands for and larger supplies of funds. In both cases $\bar{\omega} = \delta s^k$ and

$$s^k \sim |N(\mu_\sigma, \sigma^2)| = N^+(\bar{s}, \sigma_s^2).$$

Note that we have chosen the parameters so that the realized mean size remains fixed when we change the variance.

The second way in which we introduced heterogeneity was in the specification of investment opportunities. We call this variant model C. In this case, banks face identical distributions of deposits but differ by a scale factor in the distribution of investment opportunities,

$$A_t^k = |\bar{A} + \bar{A}\sigma_A \epsilon_t| \tag{10}$$

$$\omega_t^k = |i^k + \sigma_\omega \eta_t| \quad i_k \sim |N(\bar{i}, \sigma_i^2)|. \tag{11}$$

Before showing the results of the simulation we derive some qualitative results on the behaviour of the aggregate system under the assumption that banks are all connected by lending and borrowing to each other (i.e., $c = 1$). We need to compare quantities averaged over different formulations of heterogeneity (distribution of bank sizes and investment opportunities). In model A, the aggregate reserves, $A^{\mathrm{T}} = \sum_{k=1}^{M} A^k$, have a variance

$$\sigma_T^2 \sim \sum_{i=1}^{M} s^k.$$

This is a random variable that depends on the realization of the bank's size. Averaging over the realization of the s^k we obtain

$$E[\sigma_T^2] = \sum_{i=1}^{M} \bar{s} = M\bar{s}.$$

Also, the mean size of the total deposit,

$$E[A^{\mathrm{T}}] = M\bar{s}.$$

Thus, the volatility:mean ratio of the overall deposit, and hence the probability of default, on average decreases with the square root of the number of banks linked together but does not depend on the heterogeneity parameter σ_s. In model B, the aggregate variance is

$$\sigma_T^2 \sim \sum_{i=1}^{M} s_k^2.$$

and averaging over the realization of the s^k we obtain

$$E[\sigma_T^2] = \sum_{i=1}^{M} E[(s^k)^2] = \sum_{i=1}^{M} (\bar{s}^2 + \sigma_s^2) = M(\bar{s}^2 + \sigma_s^2).$$

Also in this case the volatility:mean ratio of the overall deposit decreases with the square root of M, and the probability of default is reduced, on average, when the interbank market is in place. Nonetheless, in this case, defaults become more likely to appear as the heterogeneity of the system increases.

We split the simulation analysis into two parts. First we investigate the role of connectivity and heterogeneity in the system before contagion effects could possibly take place, that is, up to the time the first event of default occurs (note that at this time, several defaults may occur simultaneously). To this end, we calculate the first time of default τ_1, averaged over 1000 different realizations of A^k and ω^k (sampled from the same distribution) as a function of c and σ_s. In this experiment $\sigma_\omega = 0.2$, $\sigma_A = 0.4$, and $\beta = 0.2$. The other parameters are as in the homogeneous case.

In all the following simulations, we set γ, the recovery rate, to zero.[4] Increasing connectivity improves the stability of the system (up to τ_1), while increasing heterogeneity makes the system more stable in model A (Fig. 4a) but more unstable in model B (Fig. 4b). In model A, bigger banks are inherently more stable, so heterogeneity seems to play a stabilising role because the introduction of larger banks contributes to the overall stability of the system. This is not true of model B in which all banks have the same mean:variance ratio.

[4] The connection between low recovery rates and the knock-on effects of bank failures has been considered in Furfine, whose simulation suggests that for contagion to be possible, recovery rates should not be too large.

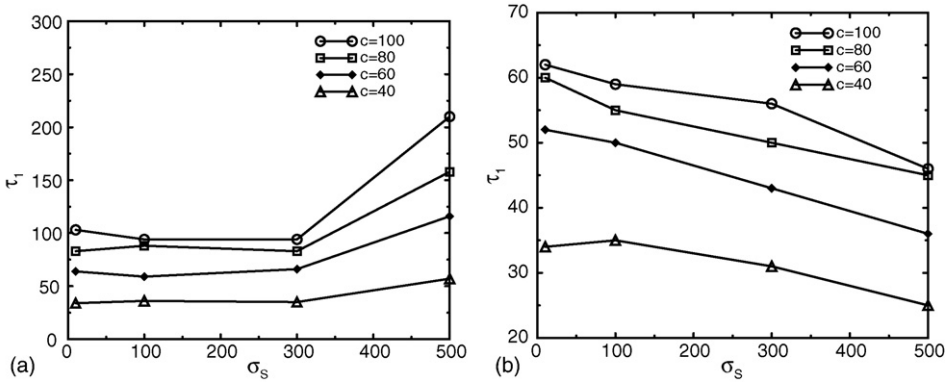

Fig. 4. First default time τ_1 vs. σ_s for model A (left) and for model B (right) at different level of connectividness, with $\sigma_A = 0.4$ and $\sigma_\omega = 0.2$.

To identify contagion effects more precisely, we analyzed a model with only two types of banks. Both faced identical distributions of deposits but differed in the distribution of investment opportunities: 'low' and 'high'. Due to low and relatively constant investment opportunities, the former set of banks tended to possess excess liquidity. In the absence of interbank lending, they never failed (see Fig. 5a: 'liquid' banks are type 0 and 'illiquid' ones are type 1).

However, with lending, they became exposed to the risk of default because of non-repayment by borrowing banks. Since non-repayment implies failure, the collapse of a borrowing bank was a direct trigger of the collapse of the lending one. Fig. 5b and c shows that as the degree of connectivity was increased across banks, the 'liquid' banks defaulted in increasing numbers. Nonetheless, the overall incidence of bank failures remained lower with interbank lending than without. This illustrates the trade-off between the insurance role of interbank credit and the possibility that it will result in 'knock-on' effects.

Having obtained these insights into the role of heterogeneous investment opportunities in the binary case, we went into the study of contagion in the case of continuously varying bank types. Contagion effects might play a more important role in model B and C than in model A because, as shown in Fig. 4, increasing heterogeneity tends to stabilise model A. For expositional purposes, in what follows, we only present the results for model C.

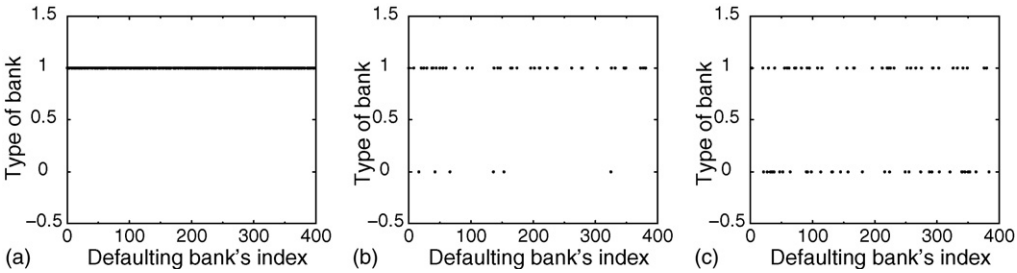

Fig. 5. Incidence of bank failures by bank type with (a) no interbank linkage, (b) 1 percent linkage, (c) 5 percent linkage. Banks are identified by an index from 1 to 400. 'Liquid' banks are type 0 and 'illiquid' ones are type 1. We add a point to the plot in correspondence of the banks that default. In the absence of interbank lending, banks of type 0 never fail. As the degree of connectivity increases across banks, the 'liquid' banks default in increasing numbers. Nonetheless, the overall incidence of bank failures remains lower with interbank lending than without.

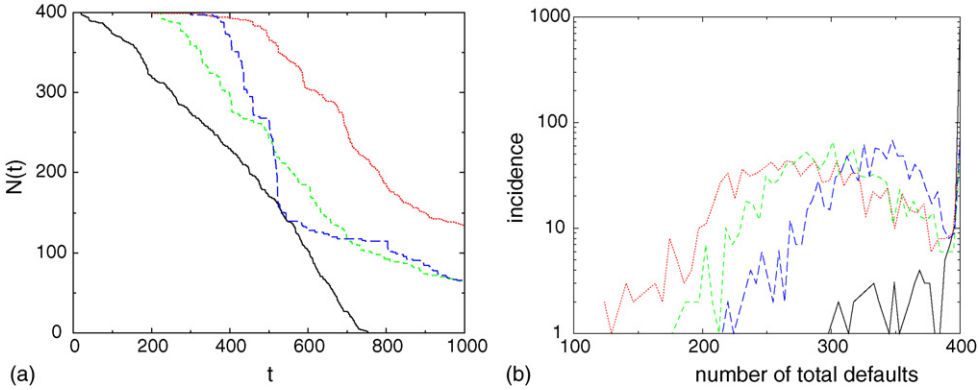

Fig. 6. Effect of increasing linkages on failure with heterogeneous investment opportunities and $\sigma_A = 0.4$, $\sigma_\omega = 0.2$, $\sigma_I = 0.5$. Connectivity varies from $c = 10$ percent (black full line), $c = 20$ percent (red dotted line), $c = 30$ percent (green dashed line), to $c = 60$ percent (blue long-dashed line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

In Fig. 6a with the variance of types $\sigma_I = 0.5$, $\sigma_A = 0.4$, and $\sigma_\omega = 0.2$, we show that increasing connectivity from 10 to 20 percent leads to less failures during the interval studied, but increasing connectivity to 30 and 60 percent leads to increasing rates of failures. This pattern was also observed for connectivities greater than 60 percent. The exact optimum (in terms of reducing the overall risk of failure) level of connectivity varies according to the degree of heterogeneity. With $\sigma_I = 0.1$, this optimum was close to $c = 50$ percent (this case is not shown). Hence, in contrast to the case of homogeneous banks, increasing the extent of the interbank market need not increase overall stability. Nonetheless as shown in Fig. 6a, the time of first default still increases with c in all cases. Thus, connectivity does seem to contribute to stabilising the system up to some point, but once defaults start they tend to snowball at high levels of connectivity. Note that the time path of surviving banks shows 'avalanches' (i.e., many banks collapsing over brief periods of time) at high levels of connectivity. This is consistent with the expectation that with high connectivity, knock on effects are taking place and causing higher rates of default once the first default has taken place. It is also interesting to note that the time of first default was found to decrease with σ_I, which is consistent with the intuition that heterogeneity has a destabilising effect on its own. The respective effects of connectivity and heterogeneity on time of first default are in line with those shown in Fig. 4 (model B).

To study further the effect of connectivity on overall rates of failure, we repeated the simulation (using the parameters in Fig. 6a) 1000 times. Fig. 6b shows the histogram for the incidence of total defaults for four different levels of connectivity, 10, 20, 30 and 60 percent, respectively. The diagram reinforces the relationship between connectivity and bank failure that was observed in the single simulation of Fig. 6; the average rate of failure goes down as c increases from 10–20, but then goes up beyond that. Interestingly, the variance in the number of failures becomes lower as c increases from 20–30 to 60 percent, implying that connectivity beyond the optimal point tends to lead invariably to more overall failures in the system.

In Fig. 7, we examine the issue of avalanches (i.e., the number of failures at one point of time). We compare the size of avalanches with increases for a single simulation (the one depicted in Fig. 6a). As connectivity increases, episodes of failure become more sparse (and as in Figs. 4 and 6a, the time before first default gets longer), but when they happen, the number of banks involved is

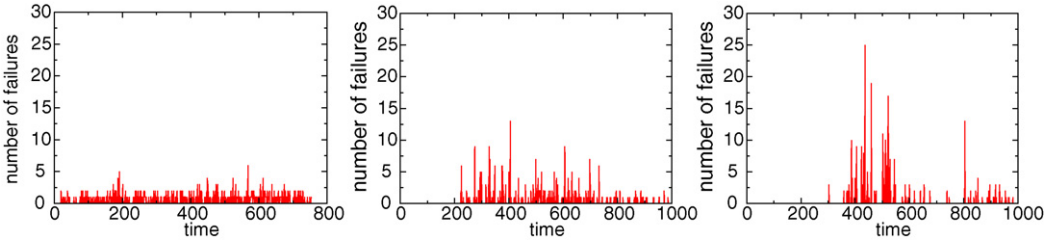

Fig. 7. Effects of increasing linkages on avalanches: $\sigma_A = 0.4$, $\sigma_\omega = 0.2$, $\sigma_I = 0.5$ and $c = 10$ percent (left), $c = 30$ percent (center) and $c = 60$ percent (right). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

greater. An interesting effect that can be observed is that as linkages increase, the system moves from a dynamics where only small avalanches are observed to one where avalanches of many different sizes become possible.

While the presence of avalanches in Fig. 7 gives heuristic support to the possibility of contagion, we now examine this issue more carefully. The structure of our model is such that, with different activities taking place at specific times within a period, it becomes impossible for a failing bank to drag down an otherwise healthy creditor within the same period. Thus, multiple contemporaneous failures are not necessarily evidence of intra-period contagion. This raises the question of why we observe a greater incidence of such failures as connectivity increases. It could be either because of delayed direct effects (i.e., banks that fail together in one period were direct creditors to a one or more debtor banks that generated significant losses due to failure in the previous period), but it could also be because of indirect forms of contagion (i.e., failure by one large debtor can drain the system of funds and make it fragile for a period of time before ongoing investment activity and the returns from it replenish these funds).

In any case, to go from the mere existence of avalanches to a more rigorous proof of contagion, we have to control for the effect that past failed loans have on a creditor bank's failure. In this regard, one issue that arises is that simply because a bank that fails had been a creditor to some other failing bank in an arbitrary past time period is not a convincing proof of contagion.

This is because even had the creditor received its full repayment it could possibly have overexposed itself to illiquid investments in the intervening time. We thus narrow the definition of direct contagion to focus on the number of banks that fail at some time given that they had lost funds due to defaults by their debtors in the period just before (i.e., had the loans that they had made the period before been repaid, these creditor banks would not have defaulted). In Fig. 8, we show the histogram of failures by such banks (as a percentage of all failing banks) for 1000 repetitions of the simulation using the parameters of Fig. 6a. The connectivity is set at 20, 30 and 60 percent. It is clear from the figure that direct contagion effects become more important as connectivity increases, since not only does overall failure increase (as shown in Fig. 6b) but the proportion of banks failing as a direct result of having lost funds due to their debtors' failures is going up even faster.

Nonetheless this effect, while clearly increasing with connectivity, only explains a small percentage of the overall failures. This implies that these simultaneous defaults can arise spontaneously as the consequence of the system reaching a critical state by its own intrinsic dynamics. One intuitive explanation is that of indirect contagion as explained in paragraph preceding Fig. 8 (i.e., as the system looses funds it becomes more fragile). Heuristic evidence for this can be seen in the center and right panel of Fig. 7, which suggest that instabilities builds up over time, leading to clusters of avalanches happening over brief periods of time.

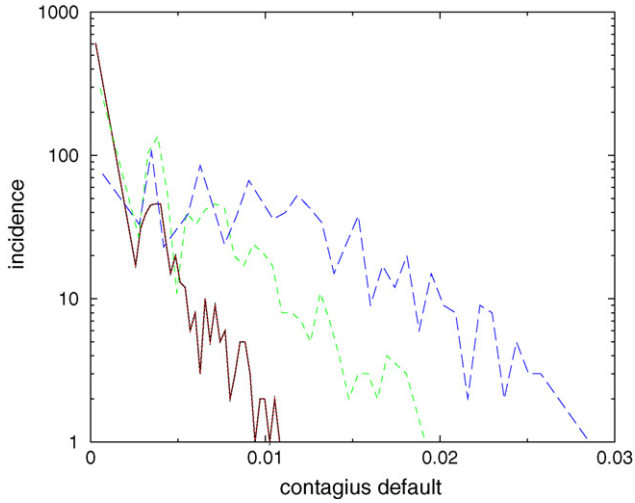

Fig. 8. Contagious default for $\sigma_A = 0.4$, $\sigma_\omega = 0.2$ and $\sigma_I = 0.5$. Connectivity varies from $c = 20$ percent (red full line), $c = 30$ percent (green dashed line), to $c = 60$ percent (blue long-dashed line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

This phenomenon resembles self-organized criticality in physical systems (Bak et al.). Critical states of the system are typically signaled by a power-law distribution of avalanches size (see Iori and Jafarey, 2001 for a more detailed discussion on this point).

Even though the main focus of this paper is on stability and contagion, it is interesting to study the effects of the interbank market on the level of investment activity. This is because interbank lending interacts with investment activity in an ambiguous way, it not only diverts funds from the latter, but by affecting the rate of bank failures, it also influences the overall level of funds available at any given time. In order to study the balance of these two effects, we measured (i) the level of investment carried out *per* investment opportunity, and (ii) the overall level of investment, against various levels of connectivity (10, 20 and 60 percent), with other parameters set as in Fig. 6a. The results are shown as histograms for 1000 repetitions in Fig. 9a and b respectively. In Fig.

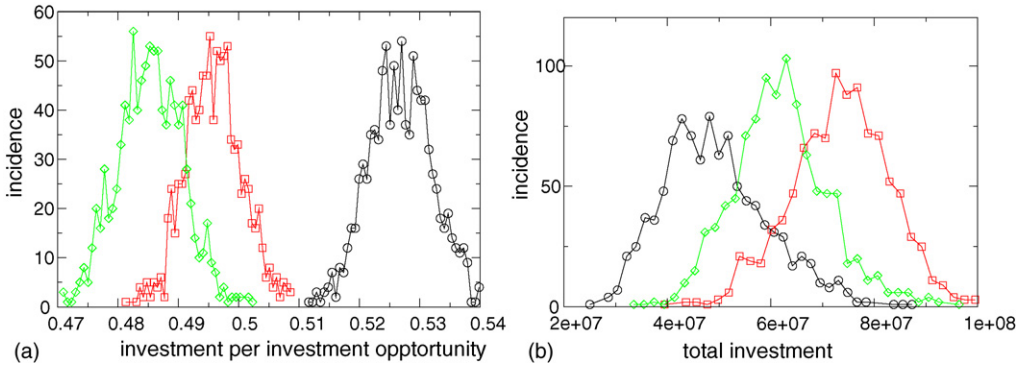

Fig. 9. (a) Investment per investment opportunity (left); (b) total investment (right). In both cases $\sigma_A = 0.4$, $\sigma_\omega = 0.2$, $\sigma_I = 0.5$ with $c = 10$ percent (black circles), $c = 20$ percent (red squares), $c = 60$ percent (green diamonds). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of the article.)

9a, only the level of investment per opportunity is observed, and this clearly falls as interbank lending increases, indicating the diversion of funds from investment to lending. The ambiguity comes out in Fig. 9b in which overall investment is recorded. In this case, the increase in systemic solvency in going from $c = 10$ to $c = 20$ percent, leads to higher overall levels of investment (despite the decrease in investment per opportunity), but since increasing connectivity too much lowers solvency, the level of investment is also less at $c = 60$ percent than at $c = 20$ percent. Thus, the stability provided by interbank lending results in an increase in overall funds available for investment, despite the diversionary effect identified in Fig. 9a. With too much connectedness, stability decreases, and so does overall investment.

4. Conclusions

This paper has studied the performance of the interbank market in its role as a safety net. The focus has been on how the statistical characteristics of a market's constituents and the nature of their interconnectedness affect the tradeoff between mutual insurance and potential for contagion.

We found that when banks are homogeneous, the insurance role of interbank lending prevails. In this situation, higher reserve requirements can lead to a higher incidence of bank failures. When banks are heterogeneous in average liquidity or average size, contagion effects may arise. We also found that such effects can be of both a direct (i.e., knock-on from a failing bank to its direct creditors) and an indirect (i.e., causing criticality in the system as a whole) nature. Interestingly, the direct effects, although systematically increasing with connectivity, remain quantitatively small, in analogy with similar findings from the empirical literature (see, e.g., Angelini et. al. for the Italian netting system). This suggests that criticality (rather than panics or direct knock-on effects) might be an important but neglected factor behind observed episodes of systemic failures. Obviously, this finding is specific to the model, but it offers an interesting further insight into the nature of conagion.

Despite the potential to create contagion, interbank lending always seems to stabilise the system in one particular sense: the elapsed time before the first failure is observed always to increase with connectivity. Our simulation results also indicate that heterogeneity alone can (as in models B and C) contribute to instability. One policy implication that suggests itself is that interbank lending relationships be confined to banks that share similar liquidity characteristics.

As for the effects of interbank lending on overall investment, we observe that while it does divert funds from the latter activity, it also indirectly contributes to it up to a point, that is up to the extent that it increases the solvency of the system by reducing the rate of default.

The insights obtained from our exercise should, of course, be taken as suggestive rather than as literal. We are conscious of the strategic limitations of the behavioural assumptions made in this paper. In reality, banks might alter their strategies in response to perceived systemic risk by altering the interest rates charged on interbank loans, the time structure of these loans, or even the degree of their participation in the market. Thus, the transmission mechanisms that generate contagion in our simulations are probably more rigid than in reality. But assuming that strategy changes take effect with a certain lag, our approach is suggestive of how an unanticipated shock to an individual bank could, under the identified circumstances, temporarily trigger a pattern of contagious failures, depending on the instantaneous statistical characteristics of the market and its participants.

Furthermore, our focus on the relationship between the statistical properties of an interbank market and its stability could help to identify circumstances under which individual banks are more likely to lend willingly to other banks as opposed to other circumstances under which they

are more likely to exercise caution. Thus, gaining insights for given bank strategies could be useful for understanding how such strategies would change in response to the overall statistical composition of a market.

Even disregarding the above caveats, the implication that interbank lending across heterogeneous banks can contribute to instability does not necessarily constitute an argument for government intervention. Before advocating such intervention, one would need to explain why banks did not endogenously sort themselves in the first place. Rochet and Tirole (1996) argue that a willingness by regulatory authorities to bail out failing institutions might reduce the incentive for banks to monitor and screen each other. The possibility of government bailouts might hinder endogenous separation.

The model of this paper can lend itself to some extensions regarding the design and stability of payments systems in general. For example, Carletti et al. (2002) have argued that bank mergers can reduce liquidity in the interbank market. Similarly, while in the present paper, in order to focus on the autonomous behaviour of the banking system, we have neglected any form of safety net other than that provided by interbank lending itself, in practise Central Banks provide external safety nets such as deposit insurance. The latter also often intervene especially when 'large' banks face crises. As a third extension, the bankruptcy procedure utilised in our model is highly stylised. Actual bankruptcy procedures vary across both countries and businesses (see Bliss, 2003). By and large, differences in bankruptcy laws will create differences in the size and nature of systemic effects. Each of these points is worthy of detailed analysis under substantial modifications of the present model, and this is left for future work.

Acknowledgements

For valuable comments and without implicating for any errors, we thank seminar participants at ICTP, Trieste, participants at the Society for Nonlinear Dynamics and Econometrics conference, Atlanta Federal Reserve, March 15–16 (2001), participants at the Seventh International Conference of the Society of Computational Economics, Yale University, New Haven, (USA), June 28–29, 2001, participants at Eleventh Symposium of the Society for Nonlinear Dynamics and Econometrics, Florence (Italy), March 13–15, 2003, participants at the Second Conference on Complex Behaviour in Economics: Modeling, Computing, and Mastering Complexity, Aix en Provence (France), May 8–11, 2003, for valuable comments. G.I. gratefully acknowledges financial support from the EPSRC (grant GR/R22629/01).

References

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

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

Bak, P., Chen, K., Scheinkman, J., Woodford, M., 1993. Aggregate fluctuations from independent sectoral shocks: self-organized criticality in a model of production and inventory dynamics. Ricerche Economichi 47, 3–30.

Bliss, R.R., 2003. Bankruptcy law and large complex financial organizations: a primer. Federal Reserve Bank of Chicago Economic Perspectives 48 (1Q), 48–58.

de Bandt, O., 1999. Competition among financial intermediaries and the risk of contagious failures. Ph.D. Dissertation, University of Chicago.

Calomiris, C.W., Kahn, C.M., 1996. The efficiency of self-regulated payment systems: Learning from the Suffolk system. Journal of Money, Credit and Banking 28, 766–797.

Carletti, E., Hartmann, P., Spagnolo, G., 2002. Bank mergers, competition and liquidity. Working Paper, European Central Bank Working Paper 292.

Cont, R., Bouchaud, J.P., 2000. Herd behaviour and aggregate fluctuation in financial markets. Macroeconomic Dynamics 4, 170–196.

Cowen, T., Kroszner, R., 1989. Scottish banking before 1845: a model for laissez faire?. Journal of Money, Credit and Banking 21, 221–231.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance and liquidity. Journal of Political Economy 91, 401–419.

Donaldson, R.G., 1992. Costly liquidation, interbank trade, bank runs and panics. Journal of Financial Intermediation 2, 59–82.

Edison, H.J., Luangaram, P., Miller, M., 2000. Asset bubbles, domino effects and "lifeboats": elements of the East Asian crisis. Economic Journal 110, 309–334.

Furfine, C., 2003. Interbank exposures: quantifying the risk of contagion. Journal of Money, Credit, and Banking 35 (1), 111–128.

Humphrey, D.B., 1986. Payments finality and risk of settlement failure. In: Saunders, A., White, L. (Eds.), Technology and the Regulation of Financial Markets: Securities, Futures and Banking. Lexington Books, Lexington, MA, pp. 97–120.

Iori, G., 2002. A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions. Journal of Economic Behaviour and Organization 49, 269–285.

Iori, G., Jafarey, S., 2001. Criticality in a model of banking crisis. Physica A 299, 205–212.

Jacklin, C.J., Bhattacharya, S., 1988. Distinguishing panics and information-based bank runs: welfare and policy implications. Journal of Political Economy 96, 568–592.

Kahn, C.M., Santos J., 2005. Endogenous Financial Fragility And Prudential Regulation, http://ssrn.com/abstract=680161.

Lux, T., Marchesi, M., 1999. Scaling and criticality in a stochastic multi-agent model of a financial market. Nature 397, 498–500.

Radelet, S., Sachs, J., 1998. The onset of the East Asian currency crisis. Working Paper 6680, NBER Working Paper No. W6680, http://ssrn.com/abstract=123429.

Rochet, J.-C., Tirole, J., 1996. Interbank lending and systemic risk. Journal of Money, Credit and Banking 28, 733–762.

Upper, C., Worms, A., 2002. Estimating Bilateral Exposures in the German Interbank Market: Is There a Danger of Contagion? Discussion. Deutsche Bundesbank, Discussion Paper No. 09/02, http://ssrn.com/abstract=304454.