BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers
No 234

Using counterfactual simulations to assess the danger of contagion in interbank markets

by Christian Upper

Monetary and Economic Department

August 2007

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The views expressed in them are those of their authors and not necessarily the views of the BIS.

Copies of publications are available from:

Bank for International Settlements
Press & Communications
CH-4002 Basel, Switzerland

E-mail: publications@bis.org

Fax: +41 61 280 9100 and +41 61 280 8100

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2007. All rights reserved. Limited extracts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)

ISSN 1682-7678 (online)

Abstract

Researchers at central banks increasingly turn to counterfactual simulations to estimate the danger of contagion owing to exposures in the interbank loan market. The present paper summarises the findings of such simulations, provides a critical assessment of the modelling assumptions on which they are based, and discusses their use in financial stability analysis. On the whole, such simulations suggest that contagious defaults are unlikely, but cannot be fully ruled out, at least in some countries. If contagion does take place, then it could lead to the breakdown of a substantial fraction of the banking system, thus imposing high costs to society. However, when interpreting these results, one has to bear in mind the potential bias caused by the very strong assumptions underlying the simulations. While robustness tests indicate that the models might be able to correctly predict whether or not contagion could be an issue and, possibly, also identify critical institutions, they are less suited for stress testing or for the analysis of policy options in crises, primarily due to their lack of behavioural foundations. Going forward, more work is needed on how to attach probabilities to the individual scenarios and on the microfoundations of the models.

JEL Classification Numbers: E58, G18, G21.

Keywords: Contagion, interbank lending, domino effects, systemic risk.

Contents

1. Introduction[1]

This paper reviews the use of counterfactual simulations to assess the danger of contagion arising from interbank lending. Since the late 1990s, such simulations have been performed by an increasing number of central banks and also feature in a recent handbook on stress testing published by the International Monetary Fund (Cihák (2007)). Typically, the simulations start from the assumption that a bank, or a group of banks, is not able to repay their borrowings in the interbank loan market and then compute the losses at the creditor banks. Contagious defaults occur if the losses on the exposures to the defaulting bank exceed the capital of a creditor. Since every default weakens the surviving bank, this could lead to a cascade of bank failures, resembling a chain of domino pieces.

Of course, defaults on interbank loans are only one mechanism through which the failure of one bank could have repercussions on other banks. Contagion could also take place if an institution does not meet its obligations in payments or in securities settlements systems.[2] Alternatively, contagion between banks might be the result of deposit withdrawals (bank runs).[3] The studies reviewed in this paper ignore these channels and consider direct contagion due to interbank lending only. This is not because I believe them not to be important, quite to the contrary. However, by focusing on one particular channel of contagion it is possible to compare a relatively homogenous set of papers and discuss their modelling assumptions in greater detail than would be the case with a broader focus. In addition, some, albeit not all, indirect channels of contagion are driven by the actions of agents attempting to protect themselves against the effects of direct contagion. For example depositors may run on a bank because they suspect it to be exposed to a failing institution and try to obtain payment before the bank fails due to this direct effect. In such cases, the simulations presented in this paper represent a benchmark against which agents assess their own actions.

Why has contagion due to interbank lending received so much attention that it merits, in my opinion, a survey paper? The substantial social costs associated with financial crises[4] make it imperative to prevent the spreading of financial distress from a single bank or a small subset of institutions to the financial system as a whole. For this reason, authorities have often chosen to bail out troubled banks rather than risking that their default might provoke the failure of other institutions. For example, almost three quarters of the 104 failures of (mainly large) banks considered by Goodhart and Schoenmaker (1995) involved a bailout of one form or another. However, while a policy of indiscriminate bailouts will surely prevent contagion, it is also likely to undermine market discipline. Striking the balance between preventing systemic crises on the one hand and limiting moral hazard on the other requires a good knowledge of the implications for the stability of the system as a whole if a bank is

[1] A previous version of this paper was circulated under the name “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”. I am grateful to Claudio Borio, Agnes Lublóy, Gregory Nguyen and Nikola Tarashev as well as seminar audiences at the Collegium Budapest’s Workshop on Systemic risk in the financial sector in October 2005, the 2006 Complexity Meeting in Aix-en-Provence, the Bank of Canada, the Centre of Central Banking Studies and the BIS for many useful comments. The views expressed in this paper are the author’s own and do not necessarily represent those of the Bank for International Settlements.

[2] Contagion in the payment system has been studied by a number of authors, starting with the seminal contributions of Humphrey (1986) and Angelini, Mariesca and Russo (1996). A scenario in which a bank defaults on its FX settlement obligations has been considered by Blavarg and Nimander (2002).

[3] See de Bandt and Hartmann (2001) for an extensive discussion of different channels of contagion. They also review a different literature that focuses on contagion between markets, rather than institutions.

[4] Estimates for the costs of financial crises are given by Boyd, Kwak and Smith (2005), Dell’Ariccia, Detragiache and Rajan (2004), Hoggarth, Reis and Saporta (2001) and Cerra and Saxena (2007).

allowed to fail. In other words, supervisory authorities have to assess the danger of contagion associated with the breakdown of individual banks in order to take the appropriate decisions when managing a crisis.

Figure 1: Interbank lending
as % of total assets, end-June 2005

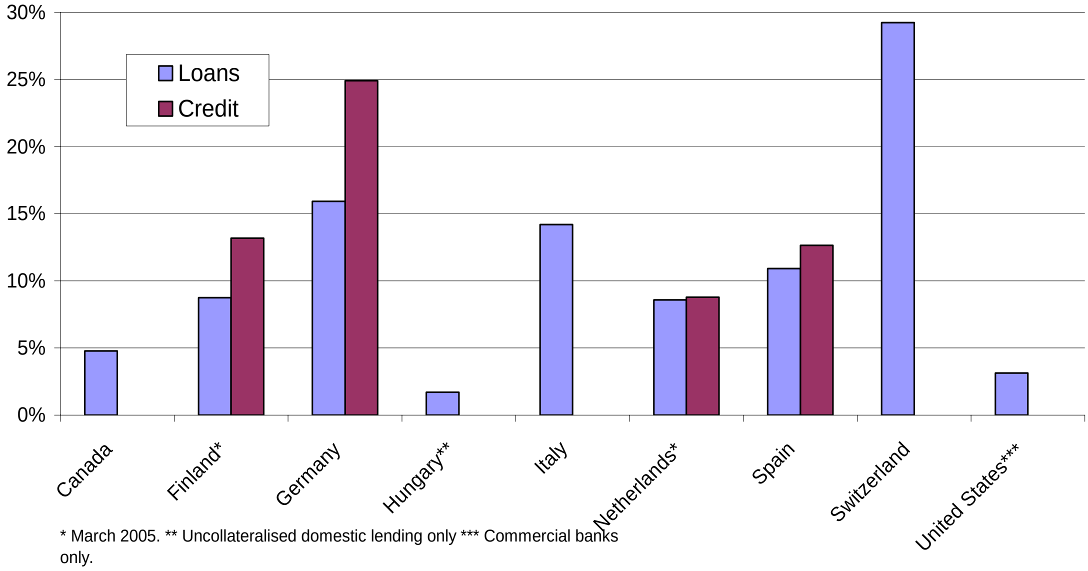

* March 2005. ** Uncollateralised domestic lending only *** Commercial banks only.
Source: BIS database, Lublóy (2005)

The focus on interbank lending stems from the fact that such exposures tend to be both large and lumpy. Loans to banks make up a large proportion of banks' balance sheets in many countries, often exceeding capital. For example, at the end of June 2005 interbank credits accounted for 29% of total assets of Swiss banks and 25% of total assets of German banks (figure 1). In other countries, the corresponding figures were lower, but, with the possible exception of the United States[5] and Canada, loans to other banks still exceeded book capital. In addition, interbank credit exposures do not only account for a large part of banks' balance sheets, they also tend to be more granular than exposures to non-banks, which further adds to the danger of contagion.

Unfortunately, analytical results on the relationship between market structure and contagion have been obtained only for a limited number of highly stylised structures of interbank markets, which are of limited use when it comes to assessing the scope for contagion in real-world banking systems. The theoretical models of Allen and Gale (2000) and Freixas, Parigi and Rochet (2000) show that the scope for contagion depends on the size of interbank exposures relative to capital as well as on the precise pattern of such linkages. Contagion is less likely to occur in what Allen and Gale term a complete structure of claims, in which every bank has symmetric exposures to all other banks. Incomplete structures, where banks are exposed only to a few neighbouring institutions, are shown to be more fragile. Finally, the scope for contagion in a system with money-centre banks, where the institutions on the

[5] The figures on interbank lending in the United States are not comparable as they do not include lending from the Federal Home Loan Banks, which amounted to $581 billion (equivalent to about 7% of commercial banks' total assets) at the end of 2004. Loans from similar institutions are counted as interbank lending in other countries.

periphery are linked to banks at the centre but not to each other, crucially depends on the precise values of the model's parameters (Freixas, Parigi and Rochet (2000)).

Given the scarcity of theoretical results, researchers have increasingly turned to computer simulations to study contagion. One strand of the literature (eg Thurner, Haner and Pichler (2003), Iori, Jafarey and Padilla (2006) and Nier, Yang, Yorulmazer and Alentorn (2007)) analyses complex artificial networks with the aim of detecting patterns which could make them prone to contagion. Other researchers simulated the effect of failures of individual institutions on the stability of the financial system using data on actual interbank exposures. The present paper surveys this latter strand of the literature.

Alternative methodologies to study contagion between banks often rely on asset price movements or deposit flows after a bank has been hit by a shock. [6] Such event studies can, by their very nature, only be done after a disruptive shock has been observed, which limits their applicability to shocks that have not yet occurred. In addition, the fact that most important banks have been bailed out rather than allowed to fail further limits the use of event analysis in the study of contagion.[7] Counterfactual simulations, by contrast, allow researchers to more or less freely specify the scenario they are interested in, without regard to whether similar events have happened in the past. Obviously, this comes at the price of making some very strong assumptions, the implications of which will be discussed below.

To give a brief summary of the findings of this survey, the simulations published so far suggest that contagion due to lending in the interbank market is likely to be rare. However, if contagion does take place, the costs to the financial system could be very high, destroying a sizable proportion of the banking system in terms of total assets. That said, it is not clear whether some of these more extreme results are the consequence of the very strong assumptions underlying the simulations. In particular, none of the simulations is based on a model that incorporates more than an extremely rudimentary behaviour by banks or policymakers.

The paper is structured as follows: Section 2 introduces data sources and the simulation methodology, with particular emphasis on the various modelling assumptions that might affect the outcome of the simulations. Sections 3 presents the results of studies that consider idiosyncratic bank failures, while section 4 summarises papers focusing on aggregate shocks. Section 5 discusses the link between insolvency and illiquidity. Section 6 assesses the reliability of counterfactual simulations of contagion and its implication for the use of such models in financial stability analysis. A final section concludes and identifies topics where further research is necessary.

2. Simulations methodology: a primer

This section introduces the methodology underlying counterfactual simulations of contagion. It first discusses how a matrix depicting bilateral exposures in the interbank market can be constructed from various data sources, and how the assumptions on the distribution of lending could affect the outcome of the simulations. The second part of the section reviews the way in which contagion is simulated.

---

[6] See de Bandt and Hartmann (2001) for a survey of the former and Schumacher (2000) and Iyer and Peydró Alcalde (2006) for examples of the latter approach.

[7] The last time that the failure of a single bank came close to producing a systemic crisis was in 1974, when the breakdown of Bankhaus Herstatt disrupted activity in the international interbank market for several weeks and apparently came close to causing a gridlock in the US payments system (Davis (1995)).

2.1 Constructing a matrix of interbank claims

An essential ingredient of any structural model for contagion is a notion of the links along which contagion may take place. In epidemiology, these links may represent physical contact and in international finance trade linkages. In our case, they represent credit exposures in the interbank market. The structure of such relationships can be represented either graphically,[8] or in matrix form. The latter approach turns out to be more useful for simulations of contagion.[9]

Suppose there are N banks that may lend to each other. In this case, the interbank market can be represented as an N x N matrix

$$\begin{array}{cc|c} & & \Sigma_i \\ \boldsymbol{X} = & \begin{bmatrix} 0 & \cdots & x_{1j} & \cdots & x_{1N} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{i1} & \cdots & 0 & \cdots & x_{iN} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{N1} & \cdots & x_{Nj} & \cdots & 0 \end{bmatrix} & \begin{array}{c} a_1 \\ \vdots \\ a_i \\ \vdots \\ a_N \end{array} \\ \hline \Sigma_j & \begin{array}{ccccc} l_1 & \cdots & l_j & \cdots & l_N \end{array} & \end{array}$$

where x_{ij} is the credit exposure of bank i vis-à-vis bank j. The row sums $a_i = \sum_j x_{ij}$ and column sums $l_j = \sum_i x_{ij}$ are bank i's total claims on other banks and bank j's liabilities in the interbank market, respectively. The zeros on the diagonal are due to the fact that banks do not lend to themselves.

The construction of matrix X crucially depends on the availability of data, which differs across countries and over time. The following subsections discuss how to use, and combine, different sources of data to obtain a systemwide matrix X.

2.1.1 Estimating X from credit register data

The construction of X is straightforward if information on all individual bilateral exposures is available from credit registers and supervisory reports. In some countries, for instance in Italy, Hungary, and Mexico, such reports cover all loans that banks extend to each other, in which case all elements of X are identified.

More often than not, however, credit registers or supervisory reports cover only exposures exceeding a threshold that is defined either in terms of the absolute amount of the exposure or as a fraction of the lender's capital. In the absence of any adjustment for missing smaller exposures, using such data could result in an underestimation of the scope for contagion, in particular for small banks. Other shortcomings of credit registers that may distort the simulation results are the reporting of credit lines instead of actual exposures (eg in the Netherlands) or the exclusion of off-balance sheet items (eg in Belgium, Austria, and the

---

8 See Boss, Elsinger, Summer and Thurner (2004), Iori et al (2005), Müller (2006) and Lublóy (2005).

9 The Appendix of Degryse and Nguyen (2005) show how the stylised models of interbank markets considered in the theoretical literature can be shown as graphs and as matrices.

Netherlands). However, even such partial data might be useful, in particular if combined with balance sheet information (subsection 2.1.3).

2.1.2 Estimating bilateral exposures from balance sheet data

A source of information that is widely available is banks' balance sheets. In contrast to credit registers they do not identify point-to-point exposures (ie the individual elements of X), but only contain information on total interbank lending and borrowing of the reporting institution. Nevertheless, this data can still be used to draw inferences on bilateral exposures, although the researcher has to make assumptions on how banks spread their interbank lending.[10]

It has become standard to assume that banks spread their lending as evenly as possible given the asset and liability positions reported in the balance sheets of all other banks. In technical terms, this corresponds to maximising the entropy (ME) of interbank linkages (see appendix for details). The concept of entropy originates from physics and was introduced into the contagion literature by Sheldon and Maurer (1998). Upper and Worms (2004) and Elsinger, Lehar and Summer (2006a) extended ME to handle zero entries on the diagonal of the matrix. A requirement for ME estimation is the availability of the balance sheets of all potential counterparties for a given balance sheet item. In practice, this has limited the use of this methodology to lending between domestic institutions.

ME is intuitively appealing as the concept is well-founded in information theory, where it denotes the most likely outcome given the a priori knowledge about an event. In the present context, this corresponds to the most likely structure of lending given the row and column sums of the interbank matrix as well as any other pieces of information that has been incorporated in the estimation programme as a constraint. From a practical point of view, ME yields a unique estimate of X, which is important since there might be an infinite number of alternative matrices with the same row and column sums.

Despite these attractive properties, there are at least three reasons why ME might not be a particularly good description of reality. First, fixed costs for screening of potential borrowers and monitoring loans may render small exposures unviable. ME, by contrast, will always return positive x_{ij}'s as long as both a_i and l_j are non-zero. Similarly, relationship lending may limit the number of counterparties of any one bank and could thus lead to a higher degree of market concentration than suggested by ME.[11] Finally, ME results in all banks holding essentially the same portfolio of interbank assets and liabilities, differing only by size and by the fact that no bank has any claims on itself.[12]

ME biases the exposure matrix X towards a "complete structure of claims", to use the terminology of Allen and Gale (2000), and should therefore raise the threshold for a shock leading to contagion. However, if contagion does occur, it may be more severe than in systems with less evenly spread exposures. Since ME will never return exposures that are

---

[10] This is necessary because combining the balance sheets of all banks of a system results in an underidentified system. The matrix X has N^2 elements, but balance sheets give N asset positions and N liability positions, corresponding to the row sums a_i and column sums l_i of, respectively. In addition, we know that the elements on the diagonal of X are zero as banks do not lend to themselves. This leaves us with N^2-$3N$ degrees of freedom.

[11] Cocco, Gomes and Martins (2005) show that interbank lending relationships are important in the Portuguese money market.

[12] These limitations become less of a problem if X is made up of several submatrices corresponding to different maturity buckets or types of exposures. For example, German banks are required to break down their interbank assets and liabilities into several maturity bands and single out exposures to counterparties belonging to the same "pillar" of the banking system. This enables Upper and Worms (2004) to estimate a total of 25 matrices, which they add up to a single, systemwide matrix that is used in their simulations.

precisely zero unless either the row or the column sums are zero or the element is explicitly constrained to be zero on the basis of outside information, it cannot recognise the barriers to contagion represented by disconnected structures. ME, by itself, will also be unable to reproduce money centre systems, where a cluster of small banks forms around large banks.

Evidence on how ME affects the findings on contagion is provided by Mistrulli (2006) and Degryse and Nguyen (2007). For Italian data dating from end-2003, ME leads to an underestimation of contagion for low losses-given-default (LGDs) and to an overestimation for high LGDs (figure 7 in Mistrulli 2006). For Belgian data from end-2002, contagion is more severe in simulations using a matrix obtained by ME than in those based on information from the credit register, although this may also have to do with the relatively high cut-off point of 10% of own funds above which banks have to report their exposures.

2.1.3 Combining balance sheet data with other sources of information

One advantage of ME is its ability to incorporate additional sources of information, eg from credit registers. This is particularly easy if exposures between two banks are known. In this case, it is possible to deduct the known value from both *ai* and *lj* and then use the RAS procedure to estimate the unknown elements of *X* only. Sometimes the exact exposures are not known but there is information on the maximum size they could take, e.g. because of regulatory constraints. The extension of the RAS algorithm by Blien and Graef (1991) is able to accommodate such inequality constraints.

The use of linear constraints to handle additional sources of information runs into difficulties if balance sheets and credit registers cover different types of exposures. For example, the British large exposures data analysed by Wells (2002) and (2004) captures uncollateralized positions only and includes off-balance sheet exposures such as derivatives or contingent liabilities. By contrast, banks' balance sheets report book loans only, without distinguishing between collateralised and uncollateralized exposures. Wells deals with this problem by assuming that banks' book loans are distributed across counterparties identically as the large exposure data. He then minimises the distance (cross-entropy) of *X* to a matrix constructed from the credit register.[13] Van Lelyveld and Liedorp (2006) follow a similar approach.

2.1.4 Estimating X from payments data

Exposures in the money market can also be estimated from payments data. This approach has been pioneered by Furfine (2003) for the federal funds market. The basic idea is quite simple: Any loan with a maturity of, say, one day, involves both a transfer of funds from the lender to the borrower on day zero and a payment of opposite sign on day one. Since loans are usually denominated in round amounts and interest is capitalised at repayment, one simply has to search all transactions of a large-scale payment system for possible repayments and then identify whether there has been a payment of the same amount minus interest but the opposite sign on the previous day.

The reliability of such estimates depends on whether interbank loans are standardised in a way that allows them to be filtered out of payments data, and on whether payments are routed through the same system. In Denmark, all conditions appear to be fulfilled, and Amundsen and Arnt (2005) are able to fully match the exposures reported by banks on a number of control days. In other countries, however, the method may be less reliable. For instance, Demiralp, Preslopsky and Whitesell (2004) find that some US banks split interest rate payments from the repayment of the principal, which introduces a substantial downward bias into Furfine's data. In Germany, the RTGS+ large-scale payment system coexists with

---

[13] See the Appendix, in particular footnote 27, for technical details.

several private systems, for which data is not available. Researchers at the Deutsche Bundesbank could identify an average of roughly 600 transactions per day, which seems few for to the more than 2000 banks that existed in Germany at the time. Another drawback of payment data is that the method can only be used to identify exposures that have already been paid back, which limits its suitability to the short end of the maturity spectrum. Again, how much this matters depends on which country one looks at.

An advantage so constructing *X* from payments data is that it gives a daily series of exposures compared to the monthly or quarterly observations at which other data are available. Payments data are therefore not likely to be plagued by window dressing and can be used in event studies which require a precise dating of exposures.

2.2 Building block # 2: Simulation methodology

Once the matrix of interbank linkages is in place, the researcher has to specify the type of shock whose potential for triggering contagion is analysed. The simplest approach is the sequential, or round-by-round, algorithm for simulating contagion, introduced into the literature by Furfine (2003) and used in most subsequent studies. It involves the following steps:

1. A bank *i* fails by assumption.
2. Any bank *j* fails if its exposure versus *i*, *xji*, multiplied by an exogenously given parameter for loss-given-default (LGD), exceeds its capital *cj*.
3. A second round of contagion occurs if there is a bank *k* for whom $LGD\left(x_{ki} + x_{kj}\right) > c_k$. Contagion stops if no additional banks go bankrupt. Otherwise a third round of contagion takes place.

It turns out that the LGD-parameter is crucial for whether or not contagion arises (see section 3). While losses-given-default in the second and higher rounds of contagion could, at least in principle, be computed from balance sheet data, few authors choose to endogenise LGDs. Instead, they fix the LGD at a specific value and keep it constant across banks and rounds of contagion. Acknowledging the paucity of our current knowledge of losses-given default, they usually perform robustness checks by trying out a large range of values.

Endogenising LGDs is far less common. Most authors that follow this approach use the clearing algorithm developed by Eisenberg and Noe (2001) that was extended and introduced into contagion analysis by Elsinger, Lehar and Summer (2006a), although this is not strictly necessary as the sequential procedure could easily be adapted it previous rounds are "revisited" as new losses lower recovery rates on past failures. Endogenising LGDs data is appealing, but it requires a series of assumptions that are far from innocuous. These concern (i) netting arrangements, (ii) the administrative costs of bankruptcy, (iii) the availability (and value) of collateral, (iv) credit risk transfer through off-balance sheet instruments, (v) the seniority of interbank relative to other claims, (vi) the market value of the defaulting bank's assets as well as the uncertainty associated with it, and (vii) the time path of and discount rate applied to recoveries.

Of these issues, netting agreements and bankruptcy costs are probably the easiest to deal with. Upper and Worms (2004), Elsinger, Lehar and Summer (2006a) and Degryse and Nguyen (2007) performed robustness checks using net instead of gross exposures and found conflicting evidence on the extent to which netting reduces the scope for contagion. In Upper and Worms (2004), netting led to a drop in the severity of the worst case of contagion from 76% of total assets to less than 10%. Similarly, in Degryse and Nguyen (2007), netting reduced the already low degree of contagion after the failure of a domestic bank even further. By contrast, Elsinger, Lehar and Summer (2006a) found that netting had only a small impact on contagion. They also incorporated bankruptcy costs and found that these substantially, and in a non-linear fashion, increase the incidence of contagion.

Incorporating collateral, credit risk transfer and the seniority structure of claims into the analysis has proved to be more difficult, although this is mainly due to a lack of data. Few of the available sources of data on interbank lending distinguish between collateralised and uncollateralized positions. As a consequence, researchers have often not been able to take collateralisation into account, except indirectly through the choice of lower LGDs.

For the same reason they also have left out of the analysis any type of credit risk transfers. In the past, this has probably not affected the results too much, as only a small number of banks have been active in such markets. For example, according to Elsinger, Lehar and Summer (2006a) about one quarter of all Austrian banks hold capital against positions in any kind of derivatives, but such positions were very small for all but a few banks. Similarly, Minton, Stulz and Williamson (2005) found that only about 6% of US banks hold credit derivatives. However, these tend to be fairly large banks, which are likely to be more relevant for contagion than smaller institutions. In addition, the market for credit derivatives and other instruments to transfer credit risk is growing rapidly. According to BIS data, the notional volume of credit default swaps increased from virtually zero at the turn of the millennium to approximately one third of the volume of domestic credit in the G10 countries in mid-2005. This sharp growth makes the assumption that no transfer of credit risk occurs increasingly untenable, even though lack of data means that there is no alternative.

Concerning the seniority of interbank claims, Elsinger, Lehar and Summer (2006a) and (2006b) assume they are junior to claims by non-banks. However, conversations with bank supervisors in Germany have shown that at least for that country this is a less reasonable assumption than an equal sharing of losses among all creditors. The situation in other countries may be different, but it will probably be difficult to reach any general conclusions.

A more substantial problem, and one that cannot be solved with more and better data, is the uncertainty associated with the market value of collateral, banks' assets more generally and the time path of recovery. Computing LGDs from balance sheet data implicitly assumes that book values can be realised under conditions of stress. This assumption has been criticised by Cifuentes, Ferruci and Shin (2005), who argue that fire sales will depress asset prices, providing a further channel of contagion. A related assumption is that recoveries are instantaneous. Anecdotal evidence suggests that this is not the case, and that merely looking at the ex post costs is misleading. For example, the Frankfurter Allgemeine Zeitung published an article in 1999 saying that the Herstatt's creditor banks had by then obtained 72% of their claims, but this was 25 years after the failure! In the case of BCCI, press reports at the time suggested that creditors expected to lose almost all of their exposures, but in the end recovered about one half. Not allowing for lags in recovering loans and the uncertainty surrounding the value of the failing banks' assets may lead to an underestimation of the scope for contagion. However, this issue could be addressed by assuming very high LGDs, as in the short-run scenario of Elsinger, Lehar and Summer (2006a), which provides an upper bound for the impact of uncertainty on the scope for contagion.

But even in light of the difficulties associated with endogenising losses-given default, doing so seems useful, as it provides some cross-sectional dispersion of LGDs. This is illustrated by the fact that although the median LGD of 35% obtained from the simulations of Elsinger, Lehar and Summer (2006a) is no far from the losses incurred in previous bank failures,[14] there are enormous variations around this value. For example, the 10% quantile of the LGD distribution is 7%, and the 90% quantile is 100%. Further evidence that

---

[14] James (1991) found that the average loss realised in bank failures in the mid-1980s United States was 30% of the book value of the bank's assets. In addition, creditors had to bear administrative and legal costs of a further 10%. Kaufman (1994) argues that the losses to creditors of Continental Illinois would have been a mere 5% of the face value of their loans, had the bank not been bailed out.

endogenising LGDs is useful is provided by Degryse and Nguyen (2007), who find that ignoring the cross-sectional dispersion of LGDs leads to an underestimation of contagion. It is not clear, however, how general this result is.

3. Contagion due to idiosyncratic shocks

The setup presented in the previous section can deal with any number of banks triggering contagion. Nevertheless, despite this intrinsic flexibility the vast majority of papers have focused exclusively on the failure of individual banks. Since simulations are easy to run and researchers often have diffuse priors concerning the choice of trigger bank and LGD, most modellers let each bank fail one at the time, and computed the impact on other banks for a broad range of LGDs. A summary of the results concerning single bank failures is given in figure 2. The x-axis plots losses-given-default and the y-axis shows the proportion of the banking system, measured by the share in total assets that is destroyed by contagious defaults (ie excluding the trigger bank).

Figure 2: Contagion due to idiosyncratic shocks

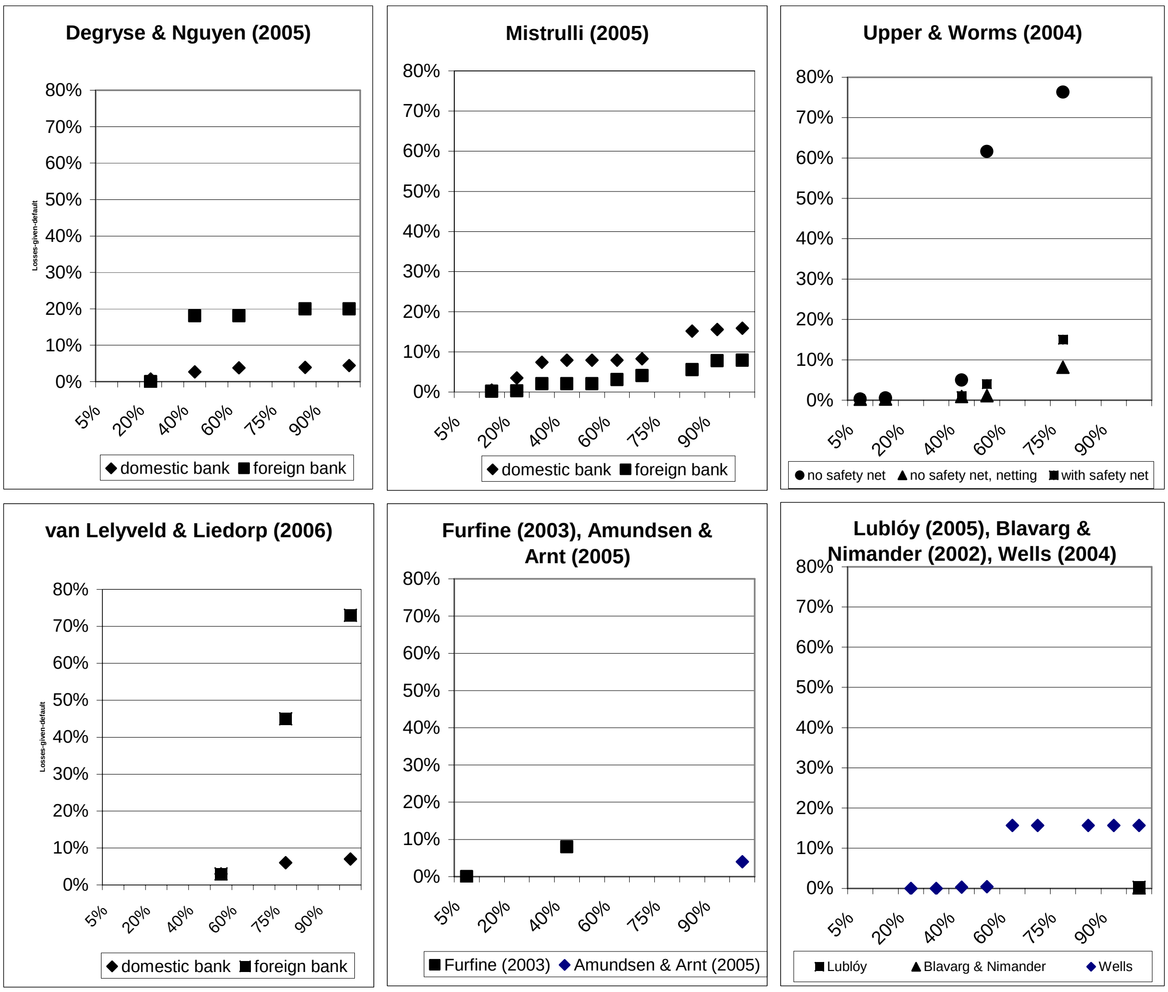

% of total assets of the banking system

Given the differences in the structure of the banking systems of the various countries and differences in the methodologies used, it is not surprising that few clear-cut results emerge. A first glance at figure 2 suggests that the danger of contagion is greatest in Germany and the Netherlands, where it may destroy institutions accounting for as much as three quarters of the banking system's total assets (Upper and Worms (2004), Van Lelyveld and Liedorp (2006)). However, a closer look reveals that both scenarios actually have a probability of zero and that they are therefore devoid of any practical relevance. In the Dutch case, the "bank" triggering the catastrophic results actually represents the aggregated banking system of Europe (except the Netherlands).[15] In Germany, the financial safety net in place at the time (end-1998) rendered the worst case scenario impossible. Allowing for guarantees from the state and from other banks reduces contagion in the worst-case scenario to 15% of the German banking system. This is of a similar order of magnitude as the results obtained by Degryse and Nguyen (2007) for Belgium (20% of total assets), Mistrulli (2005) for Italy (16%), and Wells (2004) for the UK (16%). While below the apocalyptic scenarios discussed above, these numbers are substantial by any standard, especially if one considers that most surviving banks loose a substantial proportion of their capital.

By contrast, little scope for contagion was found by Blavarg and Nimander (2002) for Sweden,[16] Lublóy (2005) for Hungary, and Sheldon and Maurer (1998) for Switzerland. Furfine (2003) and Amundsen and Arnt (2005) also report only a limited scope for contagion, but their samples are limited to overnight transactions and hence do not provide a full picture of interbank lending.

4. Contagion due to aggregate shocks

The studies reviewed so far have implicitly assumed that the shock triggering contagion has no effects on the health of the other banks except through losses on their exposures to the failing institutions. This may not be a bad assumption in the case of fraud or if the bank hit by the shock has a completely different risk profile than other banks. While such cases are not unheard of,[17] they represent only a small number of all bank failures. By contrast, the available evidence suggests that the vast majority of failures result from shocks that hit several banks simultaneously.[18] Such shocks may weaken the resiliency of the remaining banks and may thus increase the risk of contagion.

Failures due to common shocks could in principle be handled by the same tools as the ones used to analyse the effects of idiosyncratic shocks. Instead of letting individual banks fail, researchers have to specify groups of banks that are likely to fail together and follow the same procedure as for individual bank failures. However, this approach makes sense only if it is possible to define meaningful groupings of banks, for example based on their exposures to particular sectors. Perhaps for this reason, it has, to my knowledge, only been used once.

---

[15] By contrast, contagion due to the failure of a domestic institution (foreign institutions are aggregated by regions) may affect at most 7% of total assets.

[16] None of the four major banks considered failed due to contagion after the failure of a major debtor, although there was one instance where a bank lost all its tier I capital following losses on FX settlement.

[17] Examples are the failures of Baring and BCCI, respectively. The former was brought down by losses piled up (and hidden) by a single trader in Singapore, while the latter had a very different business model and organisational structure than other banks.

[18] Caprio and Klingebiel (1996) provide an extensive list of financial crises and their causes. See also Basel Committee on Banking Supervision (2004).

Guided by results of stress tests undertaken on individual bank portfolios, Lublóy (2005) grouped banks according to their FX exposures let all banks in a given category fail jointly.

An alternative methodology in which multiple failures arise endogenously in response to aggregate shocks has been suggested by Elsinger, Lehar and Summer (2006a) and (2006b). In the first paper, they embed a matrix of interbank linkages of the Austrian banking system in a risk management model covering both market and credit risk.[19] They then performed Monte Carlo simulations by drawing from the distributions of the risk factors and computing the effect on each bank's capital. If banks became insolvent, they tested for the scope for contagion to other institutions, which may already be weakened by the shock to their remaining assets. In contrast to simulations of idiosyncratic failures, their approach provides estimates of the probability in addition to estimates on the severity of contagion. In the second paper, they model the probability of default by individual banks with a multivariate Merton model that allows for correlated shocks.

The results of both papers indicate that contagious failures are rare compared to failures due to losses on exposures to non-banks. That said, if contagion does happen, it could affect a large part of the banking system. An earlier version of the first paper reports that the worst case of contagious defaults affected 37% of the banking system, measured by the failing banks' share in total assets. Moreover, fundamental failures and contagion are not independent, as contagion is much more likely in an environment where banks have already been weakened by common shocks. The second paper shows that ignoring the correlation structure of the processes driving banks' distances to default and interbank linkages results in a considerable underestimation of the probability of a systemic crisis. That said, it appears to be more important to take into account correlations in the banks' market values than exposures in the interbank market.

5. Insolvency and illiquidity

The studies reviewed so far were only concerned with cases in which contagion arose as consequence of the insolvency of the trigger bank(s). Liquidity entered the models only through the back door, via its effect on losses-given-default. However, illiquidity may not only amplify contagion, it may even cause it. An interesting simulation by Müller (2006) considers the effect on solvency and liquidity of a complete unwinding of all interbank lending. Although all banks were solvent ex ante, some institutions found that they did not have enough liquid assets to fully repay their obligations and defaulted. These defaults then led to the insolvency of creditor banks. In an extension of her base scenario, Müller analysed how the ability to draw on credit lines affected the scope for contagion. In principle, credit lines could have two opposing effects. On the one hand, they provide a source of liquidity and reduced the likelihood of banks not being able to meet their commitments, thus leading to fewer contagious failures. However, this introduces a liquidity shock at banks that have to provide the extra liquidity, which itself could lead to contagion. In Müller's simulations, the first effect dominated and the existence of credit lines reduced the scope for contagion.

Drawing on credit lines is only one of several actions that banks may take when confronted with the failure of a debtor. Perhaps the most obvious action is to sever as many of the links to the failing institution as possible, assuming that there is some time between the moment a bank learns about a failure and the moment claims are frozen. Most simulations rule out such behaviour by assuming that contagion is instantaneous, ie without any warning

[19] More recently, a similar approach has been used by Danmark Nationalbank (2007).

period.[20] Degryse and Nguyen (2007) test for the potential of contagion using interbank exposures arising from exposures with maturities of 8 days or more. Quite surprisingly, this does not affect the results very much, despite the dominance of short-term lending in the Belgian interbank market. However, while this approach solves one problem, it opens up another, namely that the unwinding of short-term lending may itself lead to contagion, as in Müller (2006).

6. How useful are counterfactual simulations of contagion?

The results of the literature reviewed in this paper could perhaps best be summarised as indicating that contagion due to exposures in the interbank loan market is an unlikely event in the sense that it happens in only a small number of the scenarios considered, but that it could have substantial effects on the health of the banking systems of many countries if it does occur. Beyond this broad picture, counterfactual simulations may offer important insights on which institutions are critical for financial stability and how the structure of the interbank market affects the scope for contagion. In principle, these models can be used in a variety of settings, such as stress testing, allocating scarce supervisory resources, analysing the costs and benefits of regulation, or crisis management. Of course, all these potential uses require the simulations to provide an accurate and timely picture of the scope for contagion, although some put higher requirements on accuracy than others. This section attempts to provide a metric on which the accuracy of the simulations can be assessed. This leads into a discussion of the potential uses of such simulations in the financial stability analysis of a central bank or regulatory authority.

6.1 How accurate are these results?

Running counterfactual simulations involves a large number of assumptions, some of which might bias the results into one direction or other. Based on the methodological discussion above, table 1 provides a list of potential sources of bias, distinguishing between the incidence of contagion (the possibility that contagion might happen) and the severity of contagion (the share of the banking system that might be subject to contagious failures). Since the biases stemming from the various assumptions can go either way, it is not possible to say whether, taken together, they result in an overestimation or an underestimation of contagion.

To which extent do these potential sources of bias undermine the usefulness of counterfactual simulations in financial stability analysis? Answer this question requires a metric for the reliability of their results. Unfortunately, conventional statistical measures such as the goodness of fit or the mean-squared forecast errors are of little use in this regard, given the rarity of bank failures and widespread government intervention. In the absence of a meaningful reliability measure, the robustness of the results to variations in the simulation methodology and in the underlying data provides probably the best criterion for the usefulness of such simulations.

Robustness checks form an important part of many of the papers reviewed in the previous sections. For example, Wells (2004), van Lelyveld and Liedorp (2006), Degryse and Nguyen (2007) and Mistrulli (2006) performed simulations on different datasets, one consisting of bilateral exposures estimated from balance sheet data using ME, and one based on credit

[20] Alternatively, one may assume that contagion takes place on the same day, ie before overnight loans could be recalled.

registers. They find that the order of magnitude of contagion is similar for different datasets, although the precise number of banks affected might differ. Similarly, the scope for contagion does not appear to depend on the date for which the simulations are performed. Unfortunately, the authors do not discuss whether the list of critical banks is robust to the choice of dataset. There is much less evidence of how sensitive the results are to changes in the simulation methodology. In particular, it is not clear what impact the exogeneity of the LGD has on the estimated incidence and scope for contagion or on the identity of critical banks.

Table 1: Potential sources of bias

Source of bias	Direction of bias		Potential remedies (and their side effects)
	Incidence	Severity	
Maximum entropy	-	+	Collect data on bilateral exposures
Reporting floors for credit register data	0[1]	-	Better data
No netting	+	+	Assume full netting (may result in underestimation if full netting cannot be enforced)
No collateral	-	-	Collect data on collateral
No bankruptcy costs	-	-	Adjust LGD
No credit risk transfer	+/-	+/-	Better data
Constant LGDs (if exogenous)	?	?	Endogenise LGD (involves a large number of alternative assumptions, see section 2.2)
Interbank claims junior to claims from nonbanks (if LGD endogenous)	+	+	Better data
No uncertainty about asset values	-	-	Micro founded model
Immediate recovery	-	-	Assume 100% LGD (leads to overestimation of effects)
Failures are not anticipated, banks cannot react	?	?	Exclude short-term assets, incorporate credit lines, micro founded model
Authorities do not react	?	?	Can easily be incorporated
No safety net	+	+	Incorporate guarantees (potential overestimation if guarantees are not fully covered, underestimation if guarantees lead to contagion)

+ overestimation, - underestimation, 0 no significant bias, +/- can go both ways, ? not clear

[1] Based on the assumption that the failure of small banks does not trigger contagion.

6.2 Potential uses of counterfactual simulations

Taken together, the available robustness tests indicate that the glass is either half full or half empty, depending on the perspective. On the one hand side, it counterfactual simulations do seem to give a rough indication on whether or not contagion could be an issue. If they remind policy makers that the fact that contagion was not observed in the past need not mean that contagion could not happen, then they have already made a big contribution.

However, counterfactual simulations do not only tell us that contagion might be possible, but they could also help to identify which banks are critical to the stability of the system. This is particularly important since the criticality of a bank is not only determined by its size or the structure of its balance sheet, which can be gauged from balance sheet data, but from the interaction of the magnitude of its interbank liabilities, its exposure to other banks, its capital and its precise location in the interbank network. Unlike any other methodology, counterfactual simulations are able to account for all of these factors simultaneously, thus offering new insights. For example, in their analysis of the Mexican banking system, Guerrero-Gómez and Lopez-Gallo (2004)) found small banks whose failure could trigger contagion of other small banks (although never the failure of large institutions).

In practice, the use of counterfactual simulations to identify critical institutions does not depend so much on whether they predict the extend of contagion with any reasonable degree of accuracy, but on whether the list of critical institutions is robust, ie that their identity does not vary if the simulation methodology is changed. While there is little published evidence in this regard, my own work on German data suggests that there are some banks that pop up regularly no matter how the model is specified.

By contrast, there are good reasons to be more sceptical regarding the use of the existing models in stress testing, in cost-benefit analysis or in assessing policy options during crises. First, the assumption that banks do not react after a shock has hit the system means that they can only be used to model events that are both unforeseen and take place within a very short period of time. This seriously limits their use in stress testing, which, almost by definition, is concerned with periods of rapidly changing market conditions in which banks tend to react very quickly. It is difficult to envisage any progress on this front unless models are built from first principles and incorporate strategic behaviour by the main actors. The second limitation of counterfactual simulations in policy analysis is provided by the absence of meaningful probability estimates, except in the Monte Carlo analysis of Elsinger, Lehar and Summer (2006a).

7. Conclusions and suggestions for further research

Counterfactual simulations of contagion may be plagued by a series of shortcomings, but they provide as yet the only way of estimating the potential for contagious defaults in a real-world banking system that can distinguish between different channels of contagion. However, while the models have improved considerably since the first of such studies was undertaken approximately ten years ago, there is still a long way to go until they become an integral part of the toolbox of any authority responsible for financial stability.

In part, the usability of counterfactual simulations has been limited by insufficient data. For this reasons, simulations have not been able to fully account for some important features of real-world interbank markets such as collateralisation, differing seniorities and the transfer of credit risk. Better data would allow researchers to capture their effects, thus rendering the estimates much more reliable.

A second area in which improvements could be made is the specification of the scenarios leading to contagion. Most studies, with the prominent exception of those by Elsinger, Lehar and Summer (2006a) and (2006b), have focused on the failure of single banks for idiosyncratic reasons. This is not the scenario that is of most relevance for supervisors. Instead, future work should consider the effect of common shocks on the stability of the banking system. In addition, any use of such models in policy work would require measures of the probability of the scenarios that may lead to contagion. It is difficult to justify costly remedial actions unless there is some information on the expected benefits.

A more fundamental problem is the absence of optimising banks. Several recent advances in economic theory could provide the behavioural foundations that are necessary to capture strategic behaviour by banks and authorities alike. For example, Iyer and Peydro-Alcalde (2005) model the interaction between losses due to defaults in the interbank market and deposit withdrawals. The role of fire sales, which could add to the losses on interbank lending, is explored by Cifuentes, Ferrucci and Shin (2005). Including such channels in counterfactual simulations would represent a major advance and could considerably improve their applicability for a large range of policy questions.

Appendix:
Maximising the entropy of the interbank lending matrix *X*

With the appropriate standardisation, interbank assets *a* and liabilities *l* can be interpreted as realisations of two marginal distributions, $f(a)$ and $f(l)$, and bilateral exposures *xij* as realisations of their joint distribution, $f(a,l)$. If $f(a)$ and $f(l)$ are independent, then $x_{ij} = a_i l_j$. Unfortunately, the resulting matrix *X* has the unappealing feature that the elements on the main diagonal that are non-zero if a bank is both lender and borrower, ie that banks lend to themselves. This problem does not necessarily disappear as the number of banks increases if interbank lending or borrowing is relatively concentrated. We therefore need to modify the independence assumption by setting $x_{ij} = 0$ for $i = j$.[21] This should be done by departing from the assumption of independence as little as possible. More formally, this means that we have to minimise the relative entropy of *X** with respect to a matrix *X* with elements $x_{ij} = a_i l_j$ for $i \neq j$ and zero for $i = j$:[22]

$$\min_{\mathbf{x}^*} \mathbf{x}^{*\prime} \ln \frac{\mathbf{x}^*}{\mathbf{x}}$$

$$\text{s.t. } \boldsymbol{x} \geq 0 \text{ and } \boldsymbol{Ax} = [a', l]',$$

where *x** and *x* are $(N^2 - N) \times 1$ vectors containing the off-diagonal elements of *X** and *X*, respectively, *a* and *l* are the marginals, and *A* is a matrix containing the adding-up restrictions $a_i = \sum_j x_{ij}$ and $l_j = \sum_i x_{ij}$. Since the objective function is strictly concave, programme (*) yields a unique solution for the structure of interbank lending *X** and can be solved numerically with the RAS algorithm that is commonly used in computing input-output tables.[23]

---

[21] Setting the elements on the diagonal equal to zero also reduces the number of coefficients to be estimated to $N^2 - 3N$ by imposing more structure on *X*.

[22] *X* could also represent a matrix constructed from credit register data (see section 2.1.3).

[23] See Blien and Graef (1991).

References

Allen, F and D Gale (2000): “Financial Contagion, Journal of Political Economy”, no 108(1), pp 1–33.

Amundsen, E and H Arnt (2005): “Contagion Risk in the Danish Interbank Market”, Danmark Nationalbank, Working Paper 2005-25.

Angelini, P, G Mariesca and D Russo (1996): “Systemic Risk in the Netting System“; Journal of Banking and Finance, 20: pp 853–68.

Basel Committee on Banking Supervision (2004) “Bank Failures in Mature Economies“, *Working Paper* no 13.

Blavarg, M and P Nimander (2002): “Inter-bank exposures and Systemic Risk“, Sveriges Riksbank, Economic Review, no 2/2002, pp 19–45.

Blien, U and F Graef (1991): “Entropy Optimization in Empirical Economic Research “, Jahrbücher für Nationalökonomie und Statistik, no 208(4), pp 399–413.

Boss, M, H Elsinger, M Summer and S Thurner (2004): “The Network Topology of the Interbank Market“, Oesterreichische Nationalbank, Financial Stability Review, no 7, pp 84–95.

Boyd, JH, S Kwak and B Smith (2005): “The Real Output Losses Associated with Modern Banking Crises“, Journal of Money, Credit and Banking, no 37(6), pp 977–99 .

Caprio, G, Jr and D Klingebiel (1996): “Bank Insolvencies: Cross-Country Experience“, World Bank, Policy Research Working Paper 1620.

Cerra, V and S C Saxena (2007): “Growth Dynamics: the Myth of Economic Recovery”, *BIS Working Paper* no 226.

Cifuentes, R, G Ferrucci and H S Shin (2005): “Liquidity Risk and Contagion”, Journal of the European Economic Association, no 3(2-3), pp 556–66.

Cihák, M (2007): “Introduction to Applied Stress Testing”, IMF Working Paper WP/07/59.

Cocco, J F, F J Gomes and N C Martins (2005): “Lending Relationships in the Interbank Market“, mimeo.

Danmarks Nationalbank (2007): “Macro Stress Testing of the Financial System”, Financial Stability, 2007.

Davis, E P (1995): Debt, Financial Fragility and Systemic Risk, revised and expanded edition, Oxford University Press.

De Bandt, O and P Hartmann (2001): “Systemic Risk: A Survey“, in Goodhart, C A E, and G Illing (eds) Financial Crisis, Contagion, and the Lender of Last Resort: A Book of Readings, Oxford, Oxford University Press, pp 249–98.

Degryse, H and G Nguyen (2007): “Interbank Exposures: An Empirical Examination of Systemic Risk in the Belgian Banking System“, International Journal of Central Banking, no 3(2), pp 123–71.

Dell’Ariccia, G, E Detragiache and R Rajan (2004) “The Real Effect of Banking Crises“, International Monetary Fund, Working Paper 05/63.

Demiralp, S, B Preslopsky and W. Whitesell (2004): “Overnight Interbank Loan Markets”, mimeo.

Eisenberg, L and T H Noe (2001): “Systemic Risk in Financial Systems“, Management Science, no 47(2), p 236–49.

Elsinger, H, A Lehar and M Summer (2006a): “Risk Assessment for Banking Systems”, Management Science, no 52, pp 1301–14.

——— (2006b) “Using Market Information for Banking System Risk Assessment”, International Journal of Central Banking, no 2(1), p 137–65.

Freixas, X, B Parigi and J C Rochet (2000): “Systemic Risk, Interbank Relations and Liquidity Provision by the Central Bank”, Journal of Money, Credit and Banking, no 32(3), Part 2, pp 611–38.

Furfine, C H (2003): “Interbank Exposures: Quantifying the Risk of Contagion“, Journal of Money, Credit and Banking, no 35(1), pp 111–28.

Goodhart, C A E and D Schoenmaker (1995): “Institutional Separation between Supervisory and Monetary Agencies“, in Goodhart, The Central Bank and the Financial System, Macmillan.

Guerrero-Gómez, S and F Lopez-Gallo (2004): “Interbank Exposures and Systemic Risk Assessment: An Empirical Analysis for the Mexican Banking Sector”, mimeo.

Hoggarth, G, R Reis and V Saporta (2001): “Costs of Banking System Instability: Some Empirical Evidence“, Journal of Banking and Finance, no 26, pp 825–55.

Humphrey, D B (1986): “Payments Finality and the Risk of Settlement Failure“, in A Saunders and L J White (eds) Technology and the Regulation of Financial Markets: Securities, Futures and Banking, Lexington, MA: Lexington Books.

Iori, G, S Jafarey and F G Padilla (2006): “Systemic Risk on the Interbank Market”, JEBO, forthcoming.

Iori, G, G de Masi, O V Precup, G Gabbi and G Caldarelli (2005): “A Network Analysis of the Italian Overnight Money Market”, mimeo.

Iyer, R and J L Peydró-Alcalde (2005): “How Does a Shock Propagate? A Model of Contagion in the Interbank Market Due to Financial Linkages”, mimeo.

——— (2006): “Interbank Contagion: Evidence from India”, mimeo.

James, C (1991): “The Losses Realized in Bank Failures”, Journal of Finance, no 46, pp 223–42.

Kaufman, G (1994): “Bank Contagion: A Review of the Theory and Evidence“; Journal of Financial Services Research, no 8, pp123–50.

Lublóy, A (2005): “Domino Effect in the Hungarian Interbank Market”, mimeo.

Minton, B A, R Stulz and R Williamson (2005): “How Much Do Banks Use Credit Derivatives to Reduce Risk?”, NBER Working Paper no 11579.

Mistrulli, P E (2005): “Interbank Lending Patterns and Financial Contagion”, mimeo.

——— (2006): “Assessing Financial Contagion in the Interbank Market: A Comparison between Estimated and Observed Bilateral Exposures”, mimeo.

Müller, J (2006): “Interbank Credit Lines as a Channel of Contagion“, Journal of Financial Services Research, 29(1): pp 37–60.

Nier, E, J Yang, T Yorulmazer and A Alentorn (2007): “Network Models and Financial Stability”, Journal of Economic Dynamics and Control, 31: pp 2033–60.

Schumacher, L (2000): “Bank Runs and Currency Run in a System without a Safety Net: Argentina and the ‘Tequila’ Shock”, Journal of Monetary Economics, 46(1): pp 257–77.

Sheldon, G and M Maurer (1998): “Interbank Lending and Systemic Risk: An Empirical Analysis for Switzerland“, Swiss Journal of Economics and Statistics, 134(4.2): pp 685–704.

Thurner, S, R Hanel and S Pichler (2003): “Risk Trading, Network Topology, and Banking Regulation”, Quantitative Finance, 3: pp 306–19.

Upper, C.and A Worms (2004): “Estimating bilateral exposures in the German interbank market: Is there a danger of contagion?”, European Economic Review, 48(4): pp 827–49.

Van Lelyveld, I and F Liedorp (2006): “Interbank Contagion in the Dutch Banking Sector: A Sensitivity Analysis”, International Journal of Central Banking, 2(2): pp 99–133.

Wells, S (2002): “UK Interbank Exposures: Systemic Risk Implications“, Bank of England, Financial Stability Review, 13: 175-81.

——— (2004): “Financial Interlinkages in the United Kingdom’s Interbank Market and the Risk of Contagion“, Bank of England, Working Paper no 230.