

UK interbank exposures: systemic risk implications

Simon Wells, Financial Industry and Regulation Division, Bank of England

A well-functioning interbank market is essential for efficient financial intermediation. But in exceptional circumstances, interlinkages between banks may provide a channel through which financial difficulties in an individual bank can be propagated to other banks. As part of theoretical efforts to understand these mechanisms, this article explores, in a stylised way, the effect of one type of extreme event – the sudden and unexpected insolvency of a single bank.

INTERLINKAGES BETWEEN financial institutions, and especially between banks, are essential for efficient financial intermediation. Of particular importance is a well-functioning interbank market, which provides an effective means of transferring liquidity from financial institutions with a cash surplus to those with a cash deficit. These transfers often occur through unsecured loans and deposits, which are straightforward and cost-effective.

But in the process of these transactions, banks can acquire significant exposures to each other. During normal market conditions, this poses no problems. But in exceptional circumstances, these interlinkages could be a channel through which problems in one bank spread across the system. Understanding the channels through which financial crises are transmitted is important in seeking to maintain financial stability, and uncovering the direct links between banks is a key element of this.

Links between banks can be of different kinds. Direct exposures mean that the failure of one could cause significant losses to its counterparties – potentially raising questions about their capital adequacy. Spill-overs would also arise if the failure of one bank led to market doubts about the soundness of other banks involved in similar activities.

This article focuses on the direct exposures between banks via the interbank market. In particular, it examines the potential for a sudden insolvency of a

single bank to trigger the failure or weakening of other banks. In reality, of course, problems in individual banks usually evolve over a long period, allowing their counterparties time to reduce their exposure. Furthermore, insolvencies can be triggered by disturbances that have a wider impact, simultaneously affecting more than one financial institution. The event considered in this article – a sudden single bank failure – is, therefore, very unlikely.

Previous studies of foreign banking systems have analysed the scope for these effects by simulating the unexpected failure of a single bank. Sheldon and Maurer (1998) consider the Swiss banking system and find that spill-over effects would mainly arise following the failure of one of the four largest Swiss banks, and then the effects would be severe: very few banks could survive such a shock. Upper and Worms (2002) analyse German banks. They find that an insolvency shock to a single bank usually triggers some additional failures, but the banks affected typically account for less than 1% of total banking system assets. In more extreme cases, however, the failure of a larger bank triggers spill-over affecting banks that account for more than 75% of total banking system assets¹. In contrast, a study of US banks by Furfine (1999) finds that the systemic effect of a major bank failure is small. The failure of the most significant bank typically affects just a few other banks, which account for less than 1% of total banking system assets. But Furfine's study considers only a small subset of total interbank exposure².

1: The average effect is very small relative to the maximum effect. This is, in part, because the German banking system is populated by a large number of small co-operative banks. Should these fail, there is little impact on other banks.


2: Specifically, Furfine (1999) focuses on Federal Funds exposures. These account for around 14% of total interbank exposures and may therefore understate the risk of spill-over.

The UK interbank market

The UK interbank market is highly concentrated: over 70% of total lending between banks operating in the UK is accounted for by only 15 institutions. The market is, therefore, characterised by ‘tiering’, ie a small number of large banks transact with each other and a greater number of smaller institutions, which place excess funds with the larger banks. In principle, this creates a potential for feedback effects between the relatively few dominant institutions and the larger number of small banks. In addition, there are large exposures amongst the biggest banks, which have a central role in the payments system.

London’s position as an international financial centre means that foreign banks have a significant involvement in the UK interbank market. Through branches located in the UK, they account for around half the total lending in the UK interbank market³. Moreover, the amount of lending from banks resident in the UK (ie UK-registered banks and branches of foreign banks situated in the UK) to other UK-resident banks is less than the amount lent to banks situated overseas. Unsecured interbank lending accounts for around 27% of UK-resident banks’ total assets. Of this, only 36% is to other UK-resident banks (Chart 1)⁴.

Chart 1:
UK-resident banks’ unsecured interbank lending as a share of total banking sector assets^{(a)(b)(c)}



Source: Bank of England.

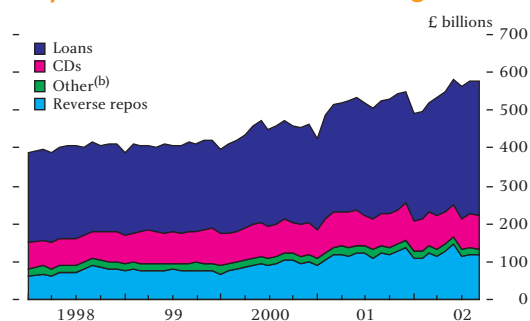
(a) All currencies.

(b) Unconsolidated, unsecured lending by UK-resident banks. Amounts outstanding Sep. 2001.

(c) Amount outstanding = £3,629 billion.

Over time, and particularly since the opening of the gilt repo market in the mid-1990s, there has been a shift from unsecured to secured exposures. To some extent, this has lowered counterparty risks because the loans are fully secured. Likewise, the growing use of collateralisation in swap markets has reduced direct exposures between banks. But much interbank lending remains unsecured (Chart 2). The unsecured interbank market has retained an important role in moving funds within the banking system in a flexible way. Banks with large correspondent bank networks and customer bases can end up with sizeable shortages at end day, which are met by transferring funds from banks in surplus.

Chart 2:
Components of UK interbank lending^(a)



Source: Bank of England.

(a) Lending by UK-resident banks to other UK-resident banks (all currencies).

(b) Commercial paper and bank bills.

Data limitations

The scope for understanding the way in which crises might potentially propagate within this structure is, however, limited by the available data. The extent of any spill-over effects would depend on the precise pattern of interbank exposures⁵. So to understand fully how the failure of one bank can affect the position of others would require a complete matrix of the bilateral exposures, arising from all of the activities of all banks that interact with each other within the global banking system. For the purpose of this article, this would mean observing the bilateral interbank exposures of all UK and foreign banks that either participate directly in the UK market, or interact with other banks that do.

3: The role of foreign banks in the UK interbank market is discussed in Box 11 of the June 2002 *Review*.

4: Total is unconsolidated and a significant element is intra-group lending.

5: A microeconomic model of spill-over in the interbank market, which demonstrates the importance of the exact interlinkages, is proposed in Allen and Gale (2000).

In practice, data are not readily available for banks that do not operate in the UK market directly. So the analysis is restricted to UK-resident banks. Even then, however, it is not possible to observe a complete matrix of bilateral exposures, describing the direct links between each and every bank resident in the UK. Rather, information on bilateral exposures is limited to the coverage of the large exposures data collected by the UK's Financial Services Authority (FSA).

These data are collected on a consolidated basis, ie each bilateral exposure reflects the combined exposure of all the reporting bank's branches and subsidiaries – including those located outside the UK – to all entities in another banking group. They include off-balance sheet as well as on-balance sheet exposures, capturing counterparty exposures under derivative contracts, contingent liabilities like guarantees and commitments, and other undrawn facilities. They do not cover exposures arising intra-day from payment and settlement activity, and for derivative transactions they cover only mark-to-market exposure, not potential future exposures that might arise as market prices change. For UK-owned banks, the data detail the size and counterparty for each of the bank's 20 largest exposures and any other exposures exceeding 10% of its Tier 1 capital. But no equivalent data are available on the large exposures of UK branches of foreign banks. Although useful, therefore, large exposure information falls well short of providing a complete map of the interactions between all banks operating in the UK.

An alternative approach is to start from data on each bank's aggregate exposure to the interbank market, and to try to find some way of estimating a breakdown into bilateral exposures. Data for each bank's total money market lending to (and deposits placed by) the UK interbank market as a whole are collected by the Bank to produce statistics on monetary aggregates and the UK-resident banking system's assets⁶. The coverage, however, differs from that of the large exposures data in a number of

important respects. It is, for example, unconsolidated: it includes only the exposures of the particular bank, not the banking group. Furthermore, focusing entirely on exposure to other banks operating in the UK means that only exposure to local branches and subsidiaries of foreign banks is captured. Moreover, it does not include off-balance sheet exposures and some other types of exposure. Nevertheless, the figures cover more than 75% of unsecured on-balance sheet interbank lending in the London market⁷.

Even though there are differences between the two data sources, the large exposures data may be used in a rather rough and ready way to adjust any estimates of bilateral exposures derived from the aggregate data. Refining the estimates in this way means that they reflect, to some extent, the pattern of activity implied by the large exposure data.

Estimating bilateral exposures

To estimate a matrix of bilateral exposures, the aggregate data are used as follows.

Data for the 24 largest banks are included individually. The remaining UK banks are grouped together, as they account for less than 1% of total interbank lending amongst UK-owned banks. Foreign banks are also grouped together, according to domicile⁸.

A matrix of bilateral exposures between these groups is then estimated under two sets of stylised assumptions about how banks distribute their aggregate interbank lending and borrowing across other individual banks, or groups of banks. Following previous studies in this area (Sheldon and Maurer (1998) and Upper and Worms (2002)), the first set of estimates assumes that banks seek to spread their borrowing and lending as widely as possible across all other banks. In practice, this involves modelling the exposure of bank A to bank B, say, as increasing both with bank A's total interbank lending and bank B's total interbank borrowing⁹. Hence, these exposures reflect the relative importance of each institution in the interbank

6: A comprehensive description of the 'monetary' data can be found in Gracie and Logan (2002).

7: Interbank exposures will also arise in payments, foreign exchange settlement and derivative positions. Data on these are limited and are therefore excluded from the analysis. This potentially ignores important channels of spill-over. On the other hand, the model does not take account of on-balance sheet netting agreements and thus may include exposures that would be netted against other exposures, if an intra-day insolvency occurred. The systemic risk posed by different categories of exposure is discussed in Michael (1998).

8: The groups are: Emerging market, French, German, Japanese, Other developed, Other EU, Swiss and US.

9: The annex describes the estimation method in more detail.

market via the size of its total borrowing and lending. This assumption rules out the possibility of 'relationship banking' ie a bank preferring some counterparties to others.

To adjust for this, the second model assumes that concentrations in the interbank market are reflected in the pattern of the large exposures data. The Bank drew on data on these exposures for 21 of the 24 UK-owned banks that enter the model and four of the overseas groups. As the large exposures data are readily available only for a single snapshot, both models are estimated using end-2000 data which, given end-of-year liquidity management, may be unrepresentative.

The propagation mechanism

Given these estimated bilateral exposures, the effect of a failure by an individual institution or group of institutions can be traced through the stylised banking system. If the initial failure leads to the insolvency of other banks, then the implications of these second-round failures can also be examined. For this purpose, it is assumed that the banks take no remedial action (for example, raising extra capital from shareholders); hence, the aim is to trace through *mechanically* the direct effects of exposures.

Each failure is assumed to be idiosyncratic to a particular bank, representing, for example, the effect of fraud. Given some assumption about the proportion of loss that is not recoverable (loss-given-default), any bank that has an exposure to the failing institution that exceeds its holdings of Tier 1 capital¹⁰ is also assumed to be insolvent¹¹. This definition of failure is to some degree arbitrary. In reality, great uncertainty surrounds the level of loss that would push a bank into insolvency. But a formal definition is required for the purpose of this kind of stylised exercise. Algebraically, assuming the insolvency of bank j , bank i also fails if

$$\theta x_{ij} > c_i, \quad (1)$$

where x_{ij} denotes the exposure of bank i to bank j , c_i is bank i 's Tier 1 capital and θ is the loss-given-default ratio. Assuming that the failure of bank j did trigger

the failure of bank i , further effects are examined on a round-by-round basis. In the third round, for example, bank k fails if its combined exposure to banks i and j exceed its Tier 1 capital, ie if

$$\theta(x_{kj} + x_{ki}) > c_k. \quad (2)$$

The loss-given-default ratio, θ , is assumed to be common to all banks. A large bank failure has not happened for many decades in the UK, making it extremely difficult to estimate the loss-given-default. As a guide, a Bank study of recoveries by the UK Deposit Protection Fund in the early 1990s suggests a median loss-given-default of 35% for failed UK banks¹². But the sample contains just 14 banks and individual values vary greatly (from 0% to 100%). Moreover, the sample includes only small banks, and so the estimate may well not be suitable for large banks. Also, these recovery rates are not adjusted for the time taken to achieve them. Therefore a study of this kind needs to consider the possibility of a higher loss-given-default rate. This is because even though a bank might be able to achieve a relatively high recovery rate over the long run, there will almost inevitably be uncertainty about eventual losses so that a bank with much of its capital at risk may be unable to continue to operate. Given this uncertainty, and following Furfine (1999) and Upper & Worms (2002), results are presented for a range of loss-given-default rates.

Results

The results are reported in two stages. First, the incidence of multiple bank failures is reported following each of the idiosyncratic bank insolvencies. This is presented for each of the two models in turn, thereby highlighting the effect of introducing the large exposures data into the estimate of the matrix of bilateral exposures. Second, the capital losses experienced by surviving banks are reported for both models because this indicates the extent of overall weakening of the banking system.

Multiple bank failures: Model I

The model can simulate the knock-on effects of 33 insolvencies¹³. Simulating each insolvency under the assumption that banks spread borrowing and

10: Tier 1 capital is defined under the Basel Accord on bank capital standards. It mainly comprises equity capital and is regarded as a high quality capital buffer.

11: Only the effects on 24 medium to large UK banks are examined, reflecting those institutions for which consolidated Tier 1 capital data are readily available. Second round failures to the remaining small UK-owned banks and foreign banks are not captured.

12: See Jackson (1996) for a review of deposit protection and bank failures in the UK.

13: The insolvency of 24 UK-owned banks, the simultaneous failure of all smaller UK-owned banks, and the failure of each of the eight groups of foreign banks.

lending as widely as possible suggests that, although an idiosyncratic failure of one bank could cause multiple failures of other banks, it is the exception rather than the rule. Even if none of the exposure is recovered (ie the loss-given-default is 100%), the insolvency of a single bank triggers additional failures in only four of the 33 cases (Table 1). With a lower, more realistic, level of loss-given-default, even fewer events lead to the failure of other banks in the stylised system.

Table 1:
Multiple failures in Model I

Loss-given-default (per cent)	Cases of multiple failures ^(a)	Balance sheet assets affected (per cent)	
		Median case ^(b)	Worst case ^(c)
100	4	8.8	25.2
80	4	1.0	6.7
60	3	0.0	6.7
40	2	0.0	0.0
20	0	0.0	0.0

Source: Bank calculations.

(a) Out of a possible 33 cases.

(b) Conditional on multiple failures occurring, the median impact in terms of aggregate balance sheet assets.

(c) The case of multiple failures that gives rise to the largest impact on aggregate balance sheet assets.

To analyse the severity of each case of spill-over, the size of the banks involved is measured by total assets¹⁴. Table 1 shows the percentage of UK-owned banks' balance sheet assets accounted for by the banks that fail due to spill-over effects. In particular, it reports the 'worst' case (ie the case of spill-over that affects the largest proportion of total balance sheet assets) and the 'median' case, which, conditional on multiple failures occurring, shows the median impact in terms of balance sheet assets.

The results are highly dependent on the assumed rate of loss-given-default. But, at all levels of loss-given-default, spill-over only arises following the insolvency of a large bank. Moreover, the failures involve a relatively small percentage of banking assets (9% in the median case of spill-over, even if loss-given-default is 100%), highlighting the small size of the banks failing due to direct exposure. On the other hand, in the worst insolvency case, much larger banks are involved and up to 25% of banking assets could be affected.

In the majority of cases, most knock-on insolvencies occur as a direct result of exposure to the initial

failure. This reflects the fact that, for the most part, only small banks are affected. Only in the more extreme cases do the spill-over effects continue for several rounds.

Multiple bank failures: Model II

A comparison of the exposures estimated in Model I with the reported large exposures suggests that the former understates the exposure of the biggest UK-owned banks to foreign banks. This may be because the data underpinning the Model I estimates do not include the exposure of overseas branches and subsidiaries of UK banks. Nor do they include exposure to non-UK parts of foreign banks, and also exclude derivative exposures. Alternatively, it may also reflect the inappropriateness of the assumption of 'wide' dispersion. In contrast, the reported large exposures between the biggest UK-owned banks are, on average, slightly lower than those estimated in Model I.

Incorporating the pattern of the large exposures into the estimates therefore increases the average exposure of big UK-owned banks to foreign banks and reduces exposures between large UK-owned banks. This shift has some interesting consequences for spill-over. The size of the exposures to foreign banks introduces the possibility of importing financial distress: in contrast to Model I, in Model II it is possible for insolvency in a group of foreign banks to trigger the direct failure of UK-owned banks.

This increases the number of insolvencies that trigger additional failures, although the average size, measured in terms of banking system assets affected, is smaller (Table 2). Under the extreme assumption of 100% loss-given-default, knock-on failures are experienced in nine of the possible 33 cases. But, in terms of size, the 'worst' case is greatly reduced, reflecting lower exposures between large UK-owned banks. And five of the nine cases involve the failure of just one small bank, which is reflected in the median case affecting only 0.1% of total assets. Against this, relative to Model I, more banking system assets are affected in the 'worst' case for loss-given-default rates of between 60% and 90%.

14: 'Total assets' refers to the aggregate consolidated balance sheet assets of the 24 UK-owned banks in the model. Recall that, in the model, only these 24 banks can fail due to direct exposures.

Table 2:
Multiple failures in Model II

Loss-given-default (per cent)	Cases of multiple failures ^(a)	Balance sheet assets affected (per cent)	
		Median case ^(b)	Worst case ^(c)
100	9	0.1	15.7
80	7	0.0	15.7
60	6	0.0	15.7
40	3	0.0	0.0
20	0	0.0	0.0

Source: Bank calculations.

(a) Out of a possible 33 cases.

(b) Conditional on multiple failures occurring, the median impact in terms of aggregate balance sheet assets.

(c) The case of multiple failures that gives rise to the largest impact on aggregate balance sheet assets.

Weakening

As mentioned previously, the definition of knock-on failure used in the preceding analysis is somewhat crude. In reality, a sufficiently large loss might cause a bank to fail, even if it does not completely wipe out its Tier 1 capital. It may prevent the bank from operating normally due to indirect effects such as credit rating downgrades and/or a rush withdrawal of deposits. Therefore, from a financial stability perspective, it is useful to characterise the distribution of losses realised by banks that do not fail but which do suffer a large loss of capital. For each 'worst' case of outright failure, this distribution is shown in Chart 3 for Model I and Chart 4 for Model II.

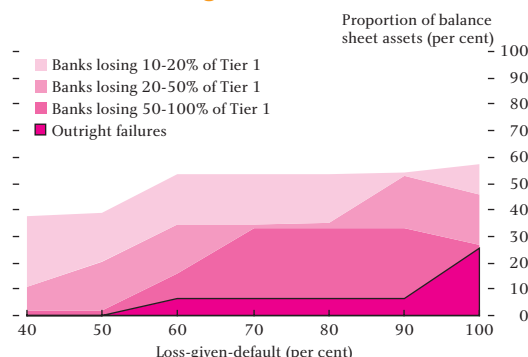
In Model I, the failure of one bank can trigger significant losses even at low levels of loss-given-default. To see this, suppose that loss-given-default is 40%. Chart 3 shows that there is a negligible amount of outright failures and only one or two small banks lose more than half their Tier 1 capital. However, some larger banks lose more than 20% of their Tier 1 capital, and together these banks account for 11% of total banking system assets. Further, banks accounting for 38% of total assets lose more than 10% of Tier 1 capital.

For loss-given-default rates higher than 60%, Model II implies a similar distribution of losses – banks accounting for around 64% of total balance sheet assets lose more than 10% of their Tier 1 capital. For lower levels of loss-given-default, the losses realised by the surviving banks is substantially reduced.

Conclusions

The interbank market, while essential for transferring funds between banks, is a channel through which problems experienced by one bank could have a direct impact on other banks in the system. In the

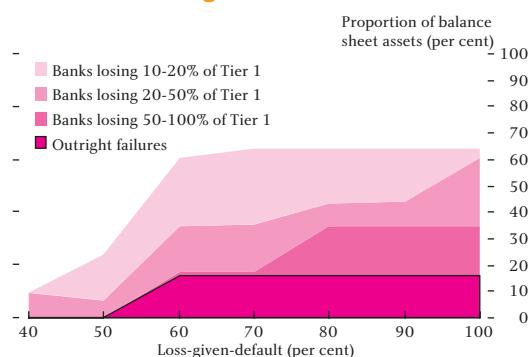
Chart 3:
Model I: Proportion of banking system assets accounted for by failing banks in 'worst' case and losses of surviving banks^(a)



Source: Bank calculations.

(a) The case of multiple failures that gives rise to the largest impact on aggregate balance sheet assets.

Chart 4:
Model II: Proportion of banking system assets accounted for by failing banks in 'worst' case and losses of surviving banks^(a)



Source: Bank calculations.

(a) The case of multiple failures that gives rise to the largest impact on aggregate balance sheet assets.

UK, there is a large interbank market and, in the course of their business, banks can acquire significant exposures to each other.

A complete analysis of the exposures in the interbank market is not feasible due to data limitations. Data on banks' aggregate exposure to the entire UK interbank market are incomplete in terms of the instruments they cover. Bilateral data, on the other hand, cover more instruments but have a different geographical coverage, and only include exposures that exceed a certain threshold.

Within the constraints of the available data, only a stylised model of the interbank market can be set up. One approach, which assumes banks seek to spread exposure as widely as possible, suggests that if a

multiple bank failure were to occur, it would most likely be triggered by the assumed insolvency of a large UK-owned bank. Large UK-owned banks generally have high credit ratings, ie their probability of default is generally low. This suggests that such a shock to the system is very unlikely.

Incorporating concentrations implied by the pattern of the large exposures data opens up the possibility that the insolvency of a large foreign bank could cause multiple bank failures in the UK system. However, when multiple failures do occur, the systemic implications seem to be somewhat less than under the simpler approach.

In the event of the failure of a large bank, there does appear to be the potential for a substantial weakening

in the capital position of a number of other banks. This is important since large banks rely on high credit ratings in order to participate in certain markets. There could therefore be significant spill-over effects by this route, even in the absence of outright failure.

This exercise is subject to important caveats. No account is taken of any remedial action by banks. Nor is any allowance made for netting of exposures, which adds to the uncertainty surrounding the appropriate level of loss-given-default. And analysis has focussed on exposures in the UK interbank market whereas UK banks are active in many financial centres around the world. But exercises such as this help to underline the importance of a sound banking industry to systemic stability and an orderly payments system.

References

- 1: Allen, F and Gale, D (2000) 'Financial Contagion', *Journal of Political Economy*, Vol. 108, No 1, pages 1–33.
- 2: Censor, Y and Zenios, S A (1997) 'Parallel Optimization', Oxford University Press.
- 3: Fang, S C, Rajasekera, J R and Tsao, H S J (1997) 'Entropy Optimization and Mathematical Programming', Kluwer Academic Publishers.
- 4: Furfine, C H (1999) 'Interbank Exposures: Quantifying the Risk of Contagion', BIS Working Papers No 70, Basel, June.
- 5: Gracie, A and Logan, A (2002) 'UK Bank Exposures: Data Sources and Financial Stability Analysis', *Bank of England Financial Stability Review*, June.
- 6: Jackson, P (1996) 'Deposit Protection and Bank Failures in the United Kingdom', *Bank of England Financial Stability Review*, Autumn.
- 7: Michael, I (1998) 'Financial Interlinkages and Systemic Risk', *Bank of England Financial Stability Review*, Spring.
- 8: Sheldon, G and Maurer, M (1998) 'Interbank Lending and Systemic Risk: An Empirical Analysis for Switzerland', *Swiss Journal of Economics and Statistics*, Vol. 134 (4.2), pages 685–704.
- 9: Upper, C and Worms, A (2002) 'Estimating Bilateral Exposures in the German Interbank Market: Is there a Danger of Contagion?', Deutsche Bundesbank Discussion Paper 09/02.

Annex: Estimating bilateral exposures with incomplete information

The aim is to estimate a complete network of bilateral exposures using each bank's total interbank assets and liabilities. For N banks, the network of exposures can be summarised by an $(N \times N)$ matrix,

$$X = \begin{bmatrix} x_{1,1} & \cdots & x_{1,j} & \cdots & x_{1,N} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{i,1} & \cdots & x_{i,j} & \cdots & x_{i,N} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ x_{N,1} & \cdots & x_{N,j} & \cdots & x_{N,N} \end{bmatrix}$$

$$\text{with } \sum_{j=1}^N x_{ij} = a_i, \quad \sum_{i=1}^N x_{ij} = l_j,$$

where x_{ij} denotes bank i 's exposure to bank j , and a_i and l_j are, respectively, bank i 's interbank assets and liabilities. Without an assumption about the distribution of the bilateral exposures the matrix, X , cannot be identified. In the absence of further information, a sensible approach is to assume that banks seek to maximise the dispersion of their interbank activity¹⁵. In this case, a bank's choice of who to lend to is determined solely by the distribution of borrowing across banks. Similarly, its choice of who to borrow from is determined by the distribution of lending across banks. The source and destination of credit are independent and, following appropriate normalisation, the individual exposures are given by the simple solution $x_{ij} = a_i l_j$.

This simple solution does not make use of all available information. For example, it implies the undesirable result that a bank may have an exposure to itself. Fortunately, it is straightforward to impose additional constraints on the interbank structure. Given a prior estimate, X^0 , of the structure, one can solve a minimisation problem to find a matrix that gets as 'close' as possible to the prior, subject to the row and column adding-up constraints. A suitable measure of 'closeness' is the cross-entropy between

the two matrices¹⁶. The appropriate interbank network is the solution to the problem

$$\min \sum_{i=1}^N \sum_{j=1}^N x_{ij} \ln \left(\frac{x_{ij}}{x_{ij}^0} \right)$$

$$\text{subject to } \sum_{j=1}^N x_{ij} = a_i$$

$$\sum_{i=1}^N x_{ij} = l_j$$

$$x_{ij} \geq 0$$

with the conventions that $x_{ij} = 0$ if, and only if, $x_{ij}^0 = 0$, and $\ln(0/0) = 0$. Problems of this type can be solved using a matrix balancing algorithm known as the RAS algorithm¹⁷.

This article compares two estimates of interbank structure. In the first model, the additional information is simply that banks do not have exposures to themselves: the prior estimate of the interbank structure is given by the maximum dispersion solution, but with appropriate zero entries on the leading diagonal. The second estimate incorporates information from the large exposures data. Algebraically, the initial estimates are given by

$$x_{ij}^{0,I} = \begin{cases} 0, & \text{if } i = j \\ a_i l_j, & \text{otherwise} \end{cases} \quad \text{and}$$

$$x_{ij}^{0,II} = \begin{cases} \frac{E_{ij}}{\sum_{j=1}^N E_{ij}} a_i, & \text{if bank } i \text{ reports large exposures} \\ x_{ij}^{0,I}, & \text{otherwise} \end{cases}$$

where E_{ij} represents the large exposure to bank j reported by bank i .

15: Technically, this involves selecting the distribution that has maximum entropy. To understand why this is a sensible approach, consider the problem of selecting a distribution for the outcome of rolling a dice. Unless one has information that the dice is loaded in some way, the distribution that places equal weight on each outcome should be selected. But this distribution also maximises the uncertainty, or entropy, about the outcome. Therefore, in the absence of information about concentrations in the interbank market, the maximum entropy distribution is chosen.

16: See Fang et al (1997) for more details on entropy and cross-entropy.

17: See Censor and Zenios (1997) for a description of the RAS algorithm. This algorithm is also applied to estimate interbank market structure in Upper and Worms (2002).