

Financial interlinkages in the United Kingdom's interbank market and the risk of contagion

Simon Wells

Working Paper no. 230

Sterling Markets Division, Bank of England.

E-mail: simon.wells@bankofengland.co.uk

The views expressed in this paper are those of the author and do not necessarily reflect those of the Bank of England. I am grateful for comments from Nicola Anderson, Andrew Gracie, Patricia Jackson and two anonymous referees. This paper has also benefited from the comments of seminar participants at the Bank of England.

Copies of working papers may be obtained from Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone 020 7601 4030, fax 020 7601 3298, e-mail mapublications@bankofengland.co.uk.

Working papers are also available at www.bankofengland.co.uk/wp/index.html.

The Bank of England's working paper series is externally refereed.

Contents

| | |
|---|----|
| Abstract | 5 |
| Summary | 7 |
| 1. Introduction | 9 |
| 2. Literature review | 10 |
| 3. Estimating the matrix of bilateral exposures | 12 |
| 3.1 Basic method | 12 |
| 3.2 Defining contagion | 14 |
| 4. Data | 16 |
| 4.1 Balance sheet data | 17 |
| 4.2 Large exposures data | 18 |
| 4.3 Estimation in practice | 19 |
| 5. Results | 21 |
| 5.1 Model I: The benchmark case | 21 |
| 5.2 Increasing weight to intragroup exposures | 24 |
| 5.3 Model II: Incorporating data on large exposures | 25 |
| 5.4 Model III: Money-centre model | 29 |
| 5.5 Interbank contagion following a system-wide shock | 31 |
| 5.6 Stability through time | 32 |
| 6. Conclusions | 33 |
| Appendix | 35 |
| References | 39 |

Abstract

A well functioning interbank market is essential for efficient financial intermediation. But interbank exposures imply the possibility of direct contagion: the insolvency of a single institution may trigger multiple bank failures due to direct credit exposures. The complete network of interbank exposures that gives rise to this channel of contagion is not observable, making it difficult to assess the systemic risk it poses. This paper uses data on loans and deposits between UK-resident banks to estimate the distribution of bilateral exposures. The potential for contagion is examined by assuming the sudden failure of each individual bank and estimating the losses incurred to other banks as a result of the initial shock. This study suggests that, while a single bank failure is rarely sufficient to trigger the outright failure of other banks, it does have the potential to weaken substantially the capital holdings of the banking system. And, when the failure of a single bank does result in knock-on effects, their severity depends greatly on the maintained assumptions about the distribution of interbank loans and the level of loss given default. But data constraints mean that drawing definitive conclusions is difficult.

Summary

A well functioning interbank market is essential for efficient financial intermediation. But interbank credit exposures imply the possibility of direct contagion: the sudden insolvency of a single institution may trigger multiple bank failures due to direct credit exposures. This paper aims to examine the potential for direct contagion in the UK interbank market.

Economic theory suggests that the potential for direct contagion depends, to some extent, on the exact structure of the interbank market. The problem is that the precise network of interbank exposures is unobservable. So this paper uses available data to estimate bilateral exposures between UK-resident banks. The estimates are used to assess the potential for direct contagion by tracing the path of assumed insolvency shocks through the banking system. We simulate the failure of each individual bank in the model and estimate the losses suffered by other banks as a result of the initial shock. We assume that contagion occurs (ie a bank fails outright) if a bank suffers a loss that exceeds its Tier 1 capital holdings.

Analysis is performed on three alternative estimates of the UK interbank structure. In each case, data on the total borrowing and lending positions of each UK-resident bank with the entire UK system are used to estimate the complete map of bilateral exposures. The first model (the benchmark case) assumes interbank borrowing and lending is as widely dispersed as possible, given each bank's observed total interbank assets and liabilities. This estimate is not conditional on market structure and so may be a poor representation of reality. The second estimate incorporates information from a database of bilateral exposures reported by banks. These data do not provide a complete map of interbank exposures: they include only exposures exceeding a certain threshold for a sample of banks. Nevertheless, incorporating information from this database into the model may mean that it better reflects concentrations in the UK interbank market. The final model is a restricted version of the benchmark case, where smaller banks and foreign banks are assumed to transact only with large UK-owned banks. In this case, the large banks can be thought of as a money centre for all banks in the UK system.

Data constraints mean that it is difficult to draw reliable conclusions about the potential for contagion. First, only banks that are resident in the United Kingdom are modelled. This means that the estimates capture only the exposure of UK-owned banks to the UK branches and subsidiaries of foreign banks and not to entire foreign banking groups. Given London's position as an international financial centre, failing to capture the full extent of exposures to foreign banks

rules out a potentially important channel of contagion. Second, suitable data are only available for interbank money market loans and deposits. Although these capture a large part of unsecured interbank activity, exposures arising from other instruments (such as interbank holdings of Certificates of Deposit and financial derivatives) are not included. Third, each model of the interbank market is derived from partial information and we show that the results depend on the assumed distribution of lending across banks.

Despite these caveats, our results give some useful information on the general potential for contagion in the UK interbank market. We explore the effect of one type of extreme event – the sudden and unexpected insolvency of a single bank. Our results show that an insolvency shock, idiosyncratic to a given bank, can lead to a substantial weakening in the capital holdings of other banks, but in most cases does not result in additional (or knock-on) bank failures. But assuming complete loss given default, our stylised model suggests that, in extreme cases, a single bank insolvency could trigger knock-on effects leading in the worst case to the failure of up to one quarter of the UK banking system. At the same time, a further quarter of the banking system would suffer losses amounting to more than 10% of their Tier 1 capital. For loss given default levels of less than 50%, contagion affects, at worst, less than 1% of total banking system assets. However, even with low loss given default, a narrow shock can considerably reduce the capital reserves of many banks. And, if the initial shocks hit during a period where the banking system is already weakened (say during a period of large macroeconomic fluctuations), the effect of contagion can be larger.

1. Introduction

This paper examines the potential for interlinkages between UK banks to propagate financial difficulties experienced by an individual bank more widely through the financial system. Two such cases of contagion are considered. First, the likelihood of ‘domino effects’ is examined, defined when an insolvency shock to a single bank causes multiple bank failures, where these are assumed to occur whenever a bank has an exposure to a failing institution that exceeds its Tier 1 capital base. Second, the paper examines the extent to which the failure of one bank could cause a wider weakening of the UK banking system, where this occurs when the initial shock triggers a significant loss of capital holdings in surviving banks, but no further outright failures. In both cases, the potential for contagion is narrowly defined in the sense that the initial insolvency shock is idiosyncratic to the failing institution.

The study is motivated by the fact that a key element of financial stability is the ability of the financial system to absorb shocks without creating widespread disruption to the economy. Given that banks play an important role in the financial system, exploring the mechanism through which shocks can be propagated from bank to bank is valuable for understanding financial stability more generally. In particular, direct exposures arising through the interbank market could, in exceptional circumstances, provide a potentially important channel for contagion. But the risk posed by interbank contagion will depend on the exact structure of the interbank market and, for the United Kingdom, precise data on the complete structure do not exist. Therefore, we use available data and make various assumptions to form estimates of the complete pattern of bilateral interbank exposures. We then gauge the potential for contagion, as defined above, by simulating the sudden and unexpected failure of a single bank.

The paper is organised as follows. The next section gives a brief literature review. Section 3 describes the basic approach to estimating the matrix of bilateral exposures and discusses the definition of contagion used. Section 4 reviews the data, modifications to consolidate the data and the incorporation of information on bilateral exposures. The results are presented in Section 5. Section 6 concludes.

2. Literature review

A microeconomic model of the interbank market is formalised by Allen and Gale (2000). They find that if the system is ‘complete’ (ie fully connected, where each bank in the system borrows from, and lends to, all other banks in the system), the likelihood of contagion is reduced. Conversely, systems where there are concentrations in interbank borrowing and lending activity can be associated with a much higher risk of contagion. This suggests that to quantify the risk of direct contagion, one would like to observe all bilateral interbank exposures directly. The problem is that only limited data on bilateral exposures are available for many financial systems. One exception is the Swedish banking system in which the Riksbank regularly survey banks about their interbank exposures. In their analysis of these data, Blåvarg and Nimander (2002) conclude that direct contagion in the Swedish banking system is possible but unlikely.

Where complete data are not available, an alternative approach is required. One possibility is to examine only those interbank exposures where bilateral data are available. Furfine (1999), for example, identifies all bilateral interbank Federal Funds exposures between US banks for a two-month period in 1998. He finds that, with a loss given default rate of 40%, the failure of the most significant bank affects between two and six other banks. Even then, however, the systemic effect of these failures is relatively small because the failing banks affected typically account for less than 1% of total assets held by the US banking system. But since Federal Funds exposures account for just 14% of total interbank exposures, these results may understate the likelihood of contagion within the US interbank market as a whole.

Another possible approach is to use aggregate data for each bank’s interbank assets and liabilities (but with no breakdown by counterparty) where these data are more widely available. Sheldon and Maurer (1998), for example, use average values of short-term interbank transactions to estimate a matrix of bilateral exposures between Swiss banks, by assuming that loans and deposits are dispersed as widely as possible across the system. They find that contagion mainly occurs following the failure of one of the four largest banks in the system. Although the corresponding effects are severe, they argue that these large banks have the smallest probability of default, so that contagion as a result of direct interbank exposures does not pose a substantial threat to the financial system in Switzerland. Upper and Worms (2002) take a similar approach to analyse the German banking system. Their approach is less aggregated because they treat all

banks individually and they include interbank assets and liabilities with longer maturities at a single snapshot in time, December 1998.⁽¹⁾ They find that an insolvency shock to any given bank almost always causes some knock-on failures. But this is usually limited to small banks, which typically account for less than 1% of total banking system assets. In a few extreme cases, however, they find that a single bank insolvency could potentially cause contagion affecting over 75% of total banking system assets. Note that conditional on the observed value of each bank's total interbank activity, their identifying assumption is an example of a 'complete' interbank structure. As such, their estimates may be biased against domino effects.

As mentioned previously, this paper adopts a similar approach and estimates a matrix of bilateral exposures for all banks operating in the United Kingdom.⁽²⁾ Following the earlier studies, the first model considered here is close to a 'complete' structure of bilateral exposures, estimated using aggregate data on each bank's total lending to (and claims on) all other banks that participate directly in the UK market. The paper then extends the earlier studies in several important ways. First, while the 'complete' model provides a benchmark to build on, it is not very realistic since it is well known that concentrations exist in the interbank market. Therefore we seek to obtain more realistic results by augmenting this basic structure with incomplete data on bilateral exposures. We also restrict the interbank structure more formally, in particular, by assuming that the largest banks act as a money centre for all other banks in the system.

Another important extension of earlier studies is that, rather than focus exclusively on the domino effect of contagion (one failure leading to multiple failures), we also examine the extent of banking system weakening that follows a single idiosyncratic shock. Specifically, we look at the distribution of loss of regulatory capital suffered by all UK banks following a single bank failure. Related to this, we study the effect of an idiosyncratic shock that hits when the entire banking system is already in distress. This reflects the fact that, in reality, multiple bank failures are often observed following severe macroeconomic shocks. Finally, since the balance sheet data used are collected from banks regularly, we repeat the analysis biannually over a three-year period in order to assess the stability of results.

⁽¹⁾ Note that Upper and Worms also include collateralised loans, but claim that the effect on their results is negligible.

⁽²⁾ Some of the results from this paper were reported, in less detail, in Wells (2002).

simple solution $x_{ij} = a_i * l_j$.⁽⁵⁾ This implies that the amount lent from bank i to bank j , is increasing in both bank i 's share of total lending, and bank j 's share of total borrowing. Hence, these exposures reflect the relative importance of each institution in the interbank market via the size of its total borrowing and lending. It also follows that this specification will be close to a 'complete' structure since any bank that borrows even a small amount will spread it across all banks that lend. Hence, this method rules out the possibility of 'relationship banking' ie a bank preferring some counterparties to others.

To arrive at our first benchmark model, we then make one further adjustment, to take into account the fact that a bank cannot have an exposure to itself. (Note that if the i^{th} bank is both a lender and a borrower, the maximum entropy solution implies it has some exposure to itself, because, $x_{ii} = a_i * l_i$, is greater than zero.) Fortunately, it is a relatively straightforward exercise to update the maximum entropy solution to incorporate such additional information. We begin by constructing an initial estimate of the interbank structure, X^0 , with elements

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

In general, the matrix X^0 violates the adding-up constraints (2). We therefore need to find a new matrix, X , that lies close to X^0 , but satisfies these constraints, which we do by choosing the matrix, X , that minimises some 'distance' measure with respect to X^0 . A suitable 'distance' measure for this type of problem is the *cross-entropy* between the two matrices.⁽⁶⁾ Cross-entropy has been used for similar problems by Upper and Worms (2002) and Elsinger *et al* (2002). Following this approach, the benchmark interbank structure is given by the solution to

⁽⁵⁾ The intuition behind this result is that a bank's choice of which other banks to lend to, and which other banks to borrow from, can be treated as stochastically independent decisions. In effect, the x_{ij} 's can then be viewed as realisations of a joint distribution, $f(a, l)$, where the distributions of assets and liabilities across banks ($f(a)$ and $f(l)$ respectively) are the marginal distributions. The amount lent from bank i to bank j , is increasing in both bank i 's share of total lending, and bank j 's share of total borrowing.

⁽⁶⁾ Further details on cross-entropy can be found in Fang *et al* (1997).

$$\begin{aligned}
& \min \sum_{i=1}^N \sum_{j=1}^N x_{ij} \ln \left(\frac{x_{ij}}{x_{ij}^0} \right) \\
& \text{subject to } \sum_{j=1}^N x_{ij} = a_i \\
& \sum_{i=1}^N x_{ij} = l_j \\
& x_{ij} \geq 0
\end{aligned}$$

with the conventions that $x_{ij} = 0$ if, and only if, $x_{ij}^0 = 0$, and $\ln(0/0) \equiv 0$. Problems of this type can be solved using a matrix-balancing algorithm known as the RAS algorithm – this is outlined in the technical appendix. Further details can be found in Censor and Zenios (1997).

3.2 Defining contagion

Having obtained an estimate of interbank exposures, the potential for contagion is modelled by allowing the failure of each bank one-by-one. Any bank that has an exposure to the failing institution(s) that exceeds its holding of Tier 1 capital is then also interpreted as being insolvent. We assume that the initial insolvency is sudden and idiosyncratic, and that the affected bank takes no remedial action. Note that this provides very much a worst case scenario. In practice, we would expect banks to experience a gradual weakening rather than a sudden failure, and for this to be reflected in other banks cutting lines to the affected institution.

The maintained definition of contagion is purposely mechanistic, so that it is a straightforward exercise to examine the knock-on effects of the sudden failure of a single institution. For example, simulating the sudden insolvency of bank j , bank i also fails if its unrecoverable loss exceeds its Tier 1 capital, ie if

$$\theta x_{ij} > c_i \tag{3}$$

where c_i = bank i 's Tier 1 capital and θ is a loss given default ratio. The ‘domino’ effect of contagion can then be examined on a ‘round-by-round’ basis. The initial (exogenous) failure is considered as the first-round failure, so that subsequent failures implied by (3) are referred to as second-round failures. A bank fails in the third round if its combined exposure to all first and

second-round failures exceeds its Tier 1 capital. Suppose, for example, that the failure of bank j in round 1 did indeed trigger the failure of bank i in round 2. In round 3, bank k also fails if

$$\theta(x_{kj} + x_{ki}) > c_k \tag{4}$$

The loss given default ratio, θ , is assumed to be common to all banks. Assigning a value to this parameter is difficult for two reasons. First, bank failures are rare in the United Kingdom making it difficult to estimate the loss given default that would be realised by their creditors were they to fail. Second, even if we could estimate total recovery rates precisely, we might want to take into account the time taken to achieve them. This is because even if a bank can achieve a relatively high recovery rate over the long run, there will almost inevitably be uncertainty about eventual losses. And an affected bank with much of its capital at risk may be unable to continue to operate on the expectation of recoveries that are uncertain. Given these uncertainties about the appropriate level of recovery, and following Furfine (1999) and Upper and Worms (2002), results are presented in Section 5 for a range of loss given default rates.


Although it is difficult to estimate a reasonable loss given default for the first-round failure (the initial sudden and unexpected bank failure must be triggered by some highly unusual event), the same may not be true for knock-on failures. Given some idea of the affected banks' balance sheets and the seniority of interbank claims, we could, in theory, calculate by how much the banks failed. In turn, we could endogenise the loss given default rate for knock-on failures.⁽⁷⁾ But to do this, we must make a stand on bankruptcy costs and assume that resolution is instantaneous. As a consequence, the estimated loss given default rates may be almost as arbitrary as assuming the initial value to be constant to all banks. Therefore, for simplicity, we maintain the assumption of a loss given default rate that is common to all banks. Nevertheless, it seems likely that in reality second-round failures would have lower loss given default rates, since the unusual event that triggered the initial failure does not impact directly on all banks. In turn, our results may overstate the potential for contagion – particularly if the assumed loss-given-default ratio is high.

⁽⁷⁾ Elsinger *et al* (2002) take this approach.

4. Data

Consistent with London's role as an international financial centre, the UK interbank market is large and banks operating within the United Kingdom lend sizable amounts to banks located elsewhere. Chart 1 shows that interbank lending accounts for 29% of total Monetary and Financial Institution (MFI) assets, around three quarters of which is lending by UK-resident banks to non-resident banks.⁽⁸⁾ The sizable exposures between UK banks and non-resident banks means that, to obtain a comprehensive understanding of the scope for contagion within the UK banking system, bilateral exposures would have to be estimated for all banks within the global system, including those between overseas banks. This is because the failure of a foreign bank, which has no direct links with banks connected to the UK market, could affect the UK banking system indirectly if it caused difficulties for other institutions that do transact with banks operating within the United Kingdom. In practice, however, data are not readily available for banks that do not operate in the UK market directly. So the analysis in this paper is restricted to UK-resident banks, ie UK-owned banks and branches or subsidiaries of foreign banks located within the United Kingdom. Balance sheet data for these banks are collected regularly by the Bank of England.

Chart 1: UK-resident banks' unsecured interbank lending as a share of total assets (December 2000)



Source: Bank of England. Amount outstanding = £3,134 billion.

⁽⁸⁾ Much of the lending to non-resident banks represents lending by foreign banks located in the United Kingdom to their head office or other overseas parts of the same group.

4.1 Balance sheet data

The exposure matrices are estimated using data on UK-resident banks' money market loans and deposits with other UK-resident banks. Of course, other types of interbank exposure exist - most notably those arising from intraday payment and settlement, and those which arise from off-balance sheet instruments. Although an ideal study of direct contagion would include all categories of exposure, data on these are limited so they cannot be considered in this paper. While this potentially ignores important channels of contagion, the most significant credit exposures are likely to be captured in the interbank loans considered here - the data cover around 75% of total (on-balance sheet) unsecured interbank assets.⁽⁹⁾ Available data suggest that over-the-counter (OTC) derivative exposures are small relative to on-balance sheet interbank exposures. In particular, at end-2000, the total OTC derivative exposure of large UK-owned banks was around 20% of total interbank exposure. But not all of this exposure is to other banks – around half of it is to non-bank financial institutions.⁽¹⁰⁾ Further, the absence of information on other types of exposure is mitigated by various collateral and netting agreements. Exposures within the main UK payments systems are collateralised and for credit exposures on OTC derivatives it is increasingly common practice among banks to use netting agreements and margining arrangements to reduce credit risks. Of course, any remaining credit risk from off-balance sheet exposures will not be captured by our model. But acting against this, the results take no account of any netting agreements for unsecured on-balance sheet transactions.

Another drawback with these data is that they are unconsolidated. This is not ideal because the UK banking system is highly concentrated and the largest banking groups often have significant overseas subsidiaries and/or own other banks in the United Kingdom. Where banks belong to a larger banking group, a significant proportion of their interbank activity may be with other entities of the same group. What matters for a study of this type are the consolidated exposures between banking groups, since the entities of a large UK-owned banking group are likely to stand or fall together.


This paper estimates consolidated exposures between banking groups using unconsolidated data. The technique used is crude but straightforward. The first step is to estimate a matrix of bilateral exposures between all individual banks operating in the United Kingdom using the

⁽⁹⁾ The other 25% is accounted for by Commercial Paper and Certificates of Deposit.

⁽¹⁰⁾ For more information, see page 92 of the *Bank of England Financial Stability Review*, June 2002.

unconsolidated data. Then, where individual banks belong to the same banking group, their exposures are grouped together to form a set of pseudo-consolidated intergroup exposures. Figure 1 illustrates a simple case. The exposure of group A to group B is provided by the sum of money lent from each entity in group A, to each entity in group B (the thickly dashed lines). Note that since they are not covered by the raw data, these pseudo-consolidated exposures do not take into account activities related to the overseas subsidiaries of UK-resident banks.

Figure 1: The consolidation method



Using this method of consolidation, the estimated intragroup exposures (shown by the thinly dashed lines between A1 and A2 (and B1 and B2) in Figure 1) are dropped from the system. In estimating the unconsolidated matrix, no distinction is made between banks that belong to the same parent company and those which do not. This may be unrealistic if intragroup activity is more concentrated than intergroup activity, since this type of relationship banking is ruled out in the benchmark analysis. The assumption of wide dispersion may therefore underestimate the extent of intragroup exposures. In turn, this implies that exposure of the consolidated group to other banks may be overestimated. As there is no clear solution to this problem, in Section 5.2 we present a sensitivity analysis of our results by increasing the relative importance of intragroup exposures.

4.2 Large exposures data

Limited information on the bilateral exposures of UK banks can be obtained from the large exposures data collected by the United Kingdom's Financial Services Authority (FSA). These data differ from the interbank loan data in a number of important respects. They have been compiled from the data for a single reporting date, taken at end-2000, and only capture exposures that exceed a certain threshold. They 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 United Kingdom – to all entities in another banking group. They cover more categories of exposure, capturing off-balance sheet instruments such as counterparty exposures under derivative contracts, contingent liabilities like guarantees and commitments, and other undrawn facilities. But exposures arising intraday from payment and settlement activity are not included. 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. For branches of non-European Economic Area (EEA) banks, only the 20 largest exposures are reported. But there are no data available on the large exposures of branches of EEA banks. So, although useful, the large exposure information falls well short of providing a complete map of the interactions between all banks operating in the United Kingdom.

4.3 Estimation in practice

As the large exposures data are only readily available for a single point in time, the bulk of the analysis is conducted using exposures estimated with end-2000 data.⁽¹¹⁾ For this period, the balance sheet data are used to estimate exposures between 24 individual UK-owned banks and banking groups, which we categorise as 'Major', 'Medium' or 'Small', depending on their total assets. The first three rows of Table A (shaded) show that these banks account for around 60% of interbank loans and deposits made in the UK market. The remaining UK-owned banks ('Other UK Small' in Table A) are grouped together. Given the aim of this study, there is little gain to including all small banks individually, since they account for less than 1% of assets held by the UK banking system.

The UK-resident foreign banks are grouped together according to domicile. Although Table A shows that these groups account for a significant proportion of total interbank activity, they are not entered individually into the model. This is because the UK branches of foreign banks do not have their own separate allocation of capital – their capital position depends on that of the bank as a whole. So, in the context of the model, even if branches of overseas banks are entered individually it is difficult to assess whether or not they would fail as a result of direct interbank exposures. In other words, they could only cause contagion by acting as the source of the initial shock. Grouping overseas banks together has the advantage of simplifying the analysis and

⁽¹¹⁾ Note that in Section 5.6 we estimate exposures at different times to assess the stability of results.

permits the modelling of extreme events to foreign banking systems, ie an event that triggers the default of all branches of banks from a given domicile.

Table A: UK-system interbank loans and deposits in 2000 Q4 (all currencies)

| Group | No. of entries in matrix | Interbank assets | | Interbank liabilities | |
|-----------------|--------------------------|------------------|-------------------|-----------------------|-------------------|
| | | £ bn | Per cent of total | £ bn | Per cent of total |
| Major UK Banks | 10 | 142.1 | 59 | 147.2 | 59 |
| Med. UK Banks | 6 | 1.7 | 1 | 1.9 | 1 |
| Small UK Banks | 8 | 0.8 | 0 | 0.5 | 0 |
| Other Small UK | 1 | 2.7 | 1 | 0.8 | 0 |
| US | 1 | 22.1 | 9 | 18.9 | 8 |
| Swiss | 1 | 6.7 | 3 | 6.8 | 3 |
| Other Overseas | 1 | 9.3 | 4 | 3.1 | 1 |
| Other Developed | 1 | 14.8 | 6 | 14.3 | 6 |
| Japanese | 1 | 13.1 | 5 | 12.0 | 5 |
| French | 1 | 3.8 | 2 | 4.4 | 2 |
| German | 1 | 16.7 | 7 | 19.8 | 8 |
| Other EU | 1 | 8.3 | 3 | 19.8 | 8 |
| Total | 33 | 242.0 | 100 | 249.4 | 100 |

Source: Bank of England.

Given our definition of contagion, a necessary condition for a bank to be vulnerable to ‘domino’ effects is that its total interbank money market lending must exceed its Tier 1 capital holdings. Using 2000 Q4 data, this is the case for 8 of the 24 UK-owned banks that enter the model individually.⁽¹²⁾ This explains why it is particularly useful to widen the scope of the study by looking at the losses of banks which do not fail outright, but lose a substantial proportion (but not all) of their Tier 1 capital.

⁽¹²⁾ In Section 5.6, where we examine the potential for contagion risks in alternative time periods, the number of vulnerable banks may differ.

5. Results

5.1 Model I: The benchmark case

Having obtained an estimate of the pseudo-consolidated interbank exposures, the effect of an assumed sudden failure of each individual bank/banking group is simulated. In total, our model permits us to study the effects of 33 different shocks: the sudden insolvency of each of the 24 UK-owned banks/banking groups, the simultaneous failure of all other UK-owned banks, and the failure of the 8 aggregate groups of foreign banks. Following a shock, we record the losses of each bank and the number of additional banks that fail.

For a range of loss given default rates, Table B shows the number of times one of the 33 shocks triggers the failure of at least one additional bank (described in the table as ‘cases of contagion’). For low loss given default, contagion does not occur. But even if loss given default is high (more than 80%), the insolvency of a single bank only triggers additional failures following 4 of the 33 scenarios. Note, however, that each of these 33 scenarios may not be equally likely. In a regulated banking system, such as the United Kingdom, the regulator may have information on which banks are most likely to trigger domino effects. If so, it may regulate these banks differently – possibly reducing their probability of default. In turn, the probability of experiencing a shock that triggers large spillover effects may be low relative to shocks that do not trigger additional failures.

Table B also indicates the severity of contagion by reporting the total balance sheet assets of the banks that fail due to spillover effects as a percentage of the total balance sheet assets for all 24 UK-owned banks included individually in the model.⁽¹³⁾ In particular, it shows the ‘worst’ case (ie the case of spillover 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. Note that the results are highly dependent on the assumed rate of loss given default. When spillover does arise, the failures involve a relatively small percentage of aggregate banking system assets (9% in the median case of spillover, even if loss given default is 100%), highlighting the relatively small size of the banks failing due to direct exposure. On the

⁽¹³⁾ ‘Total assets’ refers to the aggregate consolidated balance sheet assets of the 24 UK-owned banks in the model.

other hand, in the worst insolvency case, larger banks are involved and up to 25% of banking assets could be affected.

Two other points are worth noting at this stage. First, in the benchmark model, direct failures only follow the insolvency of a large UK-owned bank. The failure of smaller banks and groups of foreign banks (on the basis of their exposures through UK branches) do not trigger the insolvency of other institutions. To some extent, this is to be expected – the clearing banks have a central role in the UK sterling money markets and payments systems. But, given the significant role of foreign-owned banks in the UK interbank market, it is somewhat surprising that a shock to one of foreign groups does not trigger the failure of any UK-owned banks. This may in part reflect the assumption of wide dispersion – if overseas-owned banks transact mainly with just one or two UK-owned banks, there may be significant concentrations that are not captured by the benchmark model. Further, if it were possible to capture the exposure to the entire foreign banking group, it is likely that the significance of the foreign banks would increase. The second point to note is that, in the majority of cases, 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 domino effects continue for several rounds.

Table B: Benchmark results: cases of contagion

| Loss given default (%) | Cases of contagion ^(a) | Balance sheet assets affected (%) | |
|------------------------|-----------------------------------|-----------------------------------|---------------------------|
| | | Median case ^(b) | Worst case ^(c) |
| 100 | 4 | 8.80 | 25.20 |
| 90 | 4 | 0.97 | 6.65 |
| 80 | 4 | 0.97 | 6.65 |
| 70 | 3 | 0.03 | 6.65 |
| 60 | 3 | 0.03 | 6.65 |
| 50 | 3 | 0.03 | 0.03 |
| 40 | 2 | 0.03 | 0.03 |
| 30 | 2 | 0.03 | 0.03 |
| 20 | 0 | 0 | 0 |

^(a) Out of a possible 33 cases.

^(b) Conditional on some contagion occurring, the median impact in terms of aggregate balance sheet assets.

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

The definition of contagion 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 could trigger ratings downgrades leading to collateral calls or a loss of deposits that

could, in turn, make the bank unviable. Therefore, it is useful to characterise the distribution of losses realised by banks that do not fail outright, but do suffer a large loss of capital.

Charts 2 and 3 indicate the losses incurred by surviving banks for each worst case of knock-on failure described in Table B. Specifically, the lower (black) portion of Chart 2 shows the number of banks that fail in each worst case. The upper (grey) portions show the number of banks that simultaneously lose a large amount of capital.⁽¹⁴⁾ Chart 3 maps the number of banks affected onto the proportion of aggregate banking system assets accounted for by these banks. The results show that 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 2 shows that, while only one small bank fails outright, a further three banks simultaneously lose more than 20% of their Tier 1 capital. And Chart 3 shows that these banks account for over 10% of total banking system assets. In all, six banks (including the one that fails), accounting for 38% of total assets, lose more than 10% of Tier 1 capital.

Chart 2: Number of banks affected in ‘worst case’ using benchmark exposures





Chart 3: Proportion of total assets affected in ‘worst case’ using benchmark exposures



Focusing on the worst case highlights only the extreme events. More generally, the benchmark results suggest that it is very rare for a single shock to result in the outright failure of other banks. When knock-on failures do occur, typically just one or two small banks are affected. But a single insolvency can cause a substantial capital loss to the surviving banks. And some small banks

⁽¹⁴⁾ Note that the charts do not report results for loss given default below 30%. This is because the ‘worst’ case is defined in terms balance sheet assets affected by contagion. As shown in Table B, below 30%, there is no contagion.

appear particularly vulnerable to direct contagion, even at relatively low levels of loss given default. Even in the worst case however, assuming a 40% loss given default rate (the value considered by Furfine (1999)), results in contagion affecting just 0.03% of total banking system assets. For higher loss given default rates, the failure of an individual bank can trigger more knock-on failures. In the very worst case, the failing banks account for 25% of banking system assets.

5.2 Increasing weight to intragroup exposures

Recall from Section 4 that, in the process of consolidating exposures, all estimated intragroup exposures are dropped from the system. This implies that, if activity between banks that belong to the same parent company is highly concentrated, the amount of intragroup lending associated with the benchmark model might be too low. In turn, estimates of intergroup exposures may be too high and the risk of contagion overstated. To ascertain whether or not this has a serious effect on our results, we conduct a sensitivity analysis by increasing the importance of intragroup activity over and above the level implied by the benchmark model. To do this, we form a new initial estimate of the interbank structure where, if a bank i belongs to a banking group I , the elements are given by

$$x_{ij}^{INTRA} = \begin{cases} (1 - \delta)x_{ij}^0 + \delta \frac{x_{ij}^0}{\sum_{j \in I} x_{ij}^0}, & \text{if } j \in I \\ (1 - \delta)x_{ij}^0, & \text{otherwise} \end{cases}$$

Here δ determines the additional weight given to intragroup lending. One extreme, $\delta = 0$, yields the benchmark case. At the other extreme, $\delta = 1$, a bank that belongs to a larger group is assumed to borrow and lend only with other members of the same group.

Table C: Increasing weight to intragroup exposures (worst case, loss given default = 100%)

| Delta (additional weight to intragroup) | Unconsolidated assets held intragroup (% of total) | No. of banks affected by contagion | Assets affected (% of total) |
|---|--|------------------------------------|------------------------------|
| 0 | 13 | 6 | 25.20 |
| 0.1 | 21 | 4 | 6.65 |
| 0.2 | 28 | 4 | 6.65 |
| 0.3 | 35 | 4 | 6.65 |
| 0.5 | 48 | 2 | 0.06 |

In the benchmark case ($\delta = 0$), intragroup exposures account for around 13% of total unconsolidated liabilities of the major banks. It is difficult to know whether this is a reasonable approximation of reality because there are no data detailing what proportion of total borrowing by a particular institution is from other banks in the same group. However, data on the value of total deposits taken from all related *companies*, ie not only related banks, are available. In 2000 Q4, intragroup deposits accounted for around 18% of total deposits for major banks. Although this figure is only indicative, it is slightly higher than the 13% implicit in the benchmark case. Table C shows how the worst case scenario (assuming complete loss given default) changes as the amounts of intragroup exposures are increased.

Increasing the proportion of assets held intragroup to 21% reduces the degree of contagion from 25% of total assets to 7% (the same level that was found with 90% loss given default in the benchmark case). Only when we assume the amount of intragroup lending is nearly one half of total lending does the potential for contagion become tiny (affecting, at worst, just two small banks). These results suggest that a moderate increase in exposures held between banks within the same group might rule out some extreme cases, but that contagion between groups is still possible.

5.3 Model II: Incorporating data on large exposures

Concentrations can also arise in intergroup activity, meaning that the benchmark (wide-dispersion) distribution of exposures may not be a good approximation of reality. To capture the effect of introducing concentrations, the approach taken here is to adjust the benchmark estimates so that they reflect the pattern of interbank activity implied by the large exposure data. As described in Section 4, directly comparing the estimated exposures with the reported exposures is difficult because the latter have different geographical coverage and include more categories of exposure than the data underpinning the model.

Table D: Bilateral exposures of UK-owned banks exceeding 50% of creditor's Tier 1 capital

| | | Model I: Benchmark | | Large exposures data | | Model II: Incorporating large exposures | | Model III: Money centre | |
|------------------------|----------------------------------|-----------------------|-------------------|-------------------------|-------------------|---|-------------------|----------------------------|-------------------|
| Exposure of ↓ | Exposure to → | Large UK- owned | Foreign -owned | Large UK- owned | Foreign -owned | Large UK- owned | Foreign -owned | Large UK- owned | Foreign -owned |
| Large UK- owned | No. of exposures > 50% Tier 1 | 8 | 3 | 0 | 12 | 7 | 8 | 7 | 9 |
| | Of which >100% | 4 | 0 | 0 | 4 | 3 | 2 | 3 | 3 |
| Medium UK- owned | No. of exposures > 50% Tier 1 | 0 | 0 | 6 | 0 | 0 | 0 | 0 | 0 |
| | Of which >100% | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 |
| Small UK- owned | No. of exposures > 50% Tier 1 | 10 | 4 | 11 | 8 | 7 | 7 | 10 | 0 |
| | Of which >100% | 6 | 0 | 9 | 6 | 5 | 4 | 8 | 0 |

Despite these difficulties, the large exposures data are a useful indicator of the underlying concentration of interbank exposures. For UK-owned banks, Table D shows the number of bilateral exposures that exceed 50% of the creditor bank's Tier 1 capital implied first by the benchmark model and then the large exposures data. Only exposures to major UK banks and foreign institutions are reported because exposures to smaller UK-owned banks are relatively small. Table D shows that, in general, the benchmark model implies exposures between major UK-owned banks that are larger than those reported. At the same time, it underestimates the exposure of UK-owned banks to foreign banks, relative to those implied by the large exposures data. This is perhaps unsurprising given that the reported exposures include those of branches and subsidiaries located abroad and cover more categories of exposure. Alternatively, it may also reflect the inappropriateness of the assumption of wide dispersion. As we shall see, adjusting the estimates to reflect the pattern of large exposures has some interesting implications for contagion.

The large exposures data contains suitable information for 21 of the 24 banks included in the benchmark model. To incorporate this information into the model, the initial estimate of the interbank structure is constructed as follows. If bank i is included in the sample of large exposures, its estimated exposure to bank j is determined by the size of its large exposure to j , relative to its total large exposures. If, on the other hand, bank i is not in the large exposures

sample, we simply use the values from the benchmark case. Formally, letting E_{ij} denote a large exposure reported by bank i to bank j , the initial interbank matrix $X^{0,LE}$ has elements

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

Minimising the cross-entropy with respect to $X^{0,LE}$ produces an estimate of bilateral exposures that reflects the pattern implicit in the large exposures data. But note that each bank's total borrowing and lending is exactly the same as in the benchmark model.⁽¹⁵⁾ Only the distribution of funds across banks is changed – the stock of interbank funds remains the same. Consequently, Model II is quite similar to the benchmark case, but moves one step towards the exposure distribution reported by the banks.

To see this, Table D shows the number of exposures exceeding 50% (and 100%) of the creditors Tier 1 capital after incorporating the large exposures data (henceforth Model II). Relative to the benchmark model, Model II implies that the UK-owned banks have higher exposure to the groups of foreign-owned banks. In part, this reflects the inclusion of the exposure of UK banks to the whole overseas group, not just its UK subsidiary. But exposures between major UK-owned banks remain higher than those implied by the large exposures data. This is why Model II should only be seen as a small step towards matching the pattern of the large exposures data.

Table E shows that the shift in emphasis has some interesting consequences for contagion. In particular, the increased exposure of UK-owned banks to foreign banks operating within the UK means that, 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 shocks that trigger additional failures, although the severity, measured in terms of banking system assets affected, is smaller. For example, under the assumption of 100% loss given default, contagion is experienced in 9 of the possible 33 cases. But, in terms of size, the 'worst' case is greatly reduced. And 5 of the 9 cases involve the failure of just one small bank, which is reflected in the median case affecting

⁽¹⁵⁾ As the large exposures data are consolidated, the bilateral exposures in Model II are estimated using each bank's total consolidated interbank borrowing and lending implied by the benchmark model.

only 0.06% of total assets. Against this, relative to the benchmark model, more banking system assets are affected in the worst case for loss given default rates of between 60% and 90%.

Table E: Model II: Cases of contagion incorporating large exposure data

| Loss given default (%) | Cases of contagion (out of 32) | Balance sheet assets affected (%) | |
|------------------------|--------------------------------|-----------------------------------|------------|
| | | Median case | Worst case |
| 100 | 9 | 0.06 | 15.66 |
| 90 | 9 | 0.03 | 15.66 |
| 80 | 7 | 0.04 | 15.66 |
| 70 | 6 | 0.03 | 15.66 |
| 60 | 6 | 0.03 | 15.66 |
| 50 | 4 | 0.03 | 0.04 |
| 40 | 0 | 0.03 | 0.03 |
| 30 | 0 | 0 | 0 |
| 20 | 0 | 0 | 0 |

Charts 4 and 5 show, respectively, the distribution of losses in each worst case of knock-on failure implied by Model II (defined in terms of number of banks affected and their balance sheet assets). For loss given default rates higher than 60%, Chart 5 shows that Model II implies a similar distribution of losses to that implied by the benchmark model – 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 are somewhat reduced.

Chart 4: No. of banks affected in ‘worst case’ incorporating large exposure data





Chart 5: Proportion of total assets affected in ‘worst case’ incorporating large exposures



5.4 Model III: Money-centre model

The benchmark model is an example of a highly connected structure because banks spread borrowing across all banks that lend, conditional on the importance of each bank in the market. Model II is less connected since the interlinkages reflect, to some extent, the pattern of the reported large exposures. The final model considered in this paper takes a more formal approach to describe a disconnected structure. Specifically, the major banks are assumed to act as a money centre for all other banks participating in the UK interbank market. As shown in Figure 2, smaller banks and foreign banks in this system must carry out all interbank activity with the major banks. The major banks, on the other hand, are fully connected with all banks, and with each other. In practice, the money-centre model is estimated by placing additional zero entries into the initial matrix, X^0 . Therefore, subject to the additional restrictions, all banks are assumed to maximise the dispersion of interbank borrowing and lending.

There are two main reasons for studying this type of structure. First, it reflects the fact that only large clearing banks are direct members of the UK payments system: smaller banks, which are not direct members, must keep balances at, and make payments through, larger banks. Second, the model of Allen and Gale (2000) suggests that disconnected structures can be associated with increased risk of contagion; hence, it is interesting to see whether or not restricting the interlinkages between banks in this way does in fact provide support for this theory.

Figure 2: Money-centre structure




Table D shows the number of exposures exceeding 50% of the creditor bank's Tier 1 capital in Model III. Because foreign banks are not able to borrow from each other they borrow more from

large UK-owned banks, relative to the benchmark model. This, in turn, means that large UK banks lend less to each other. For lower levels of loss given default, these changes have little impact on the amount of contagion – the effects are similar to those implied by Models I and II (Table F). For a loss given default rate above 80%, however, the worst case of contagion is more severe than in the earlier models: for 100% loss given default, banks accounting for 42% of balance sheet assets fail as a result of domino effects. And this is accompanied by an increase in the amount of weakening. Charts 6 and 7 shows that, in the worst case, all sizable UK banks lose at least 10% of their Tier 1 capital. In extreme cases, therefore, the disconnected structure appears more vulnerable to contagion.

Table F: Model III: Cases of contagion in money-centre model

| Loss given default (%) | Cases of contagion (out of 32) | Balance sheet assets affected (%) | |
|------------------------|--------------------------------|-----------------------------------|------------|
| | | Median case | Worst case |
| 100 | 7 | 0.99 | 42.22 |
| 90 | 6 | 0.05 | 25.20 |
| 80 | 4 | 0.05 | 6.65 |
| 70 | 4 | 0.05 | 6.65 |
| 60 | 2 | 0.06 | 0.06 |
| 50 | 1 | 0.03 | 0.03 |
| 40 | 1 | 0.03 | 0.03 |
| 30 | 1 | 0.03 | 0.03 |
| 20 | 0 | 0 | 0 |

Chart 6: No. of banks affected in ‘worst case’ in money-centre model





Chart 7: Proportion of total assets affected in ‘worst case’ in money-centre model



5.5 Interbank contagion following a system-wide shock


Historically, multiple bank failures have often been observed during periods of large macroeconomic fluctuations (see Gorton (1988)). This is often attributed to increased volatility in banks' assets relative to liabilities. Although this study focuses on narrow, idiosyncratic insolvency shocks, it is also possible to capture the effect of a narrow shock hitting a single bank during a period of distress for the entire banking system. To model this, we assume that the idiosyncratic shock hits when the capital of all banks has been reduced by a fixed proportion. Such a situation may arise if, in the face of a common macroeconomic shock, all banks raise provisions against non-performing loans by a similar proportion of their total assets. A study by Pain (2002) suggests that, *ceteris paribus*, a 1% rise in the real effective sterling exchange rates has typically been accompanied by a 10% rise in new provisions by UK commercial banks. Using end-2000 data, and assuming that the provisions are taken from capital, this suggests that a 5% rise in the real exchange rate could reduce the Tier 1 capital holdings of the five largest UK-owned banks by around 6%. Similarly, Pain's (2002) work suggests that during a severe recession, sufficient to reduce GDP growth by 4 percentage points, a rise in provisions could lower Tier 1 capital holdings by around 3%.

Of course, these figures are purely illustrative. So rather than hypothesise the extent of capital depletion in periods of stress, we simply report the percentage of total banking assets affected by contagion in the worst case for a range of assumptions about the size of the system-wide shock. Results are presented in Chart 8 for all three structures examined above, where loss given default is assumed to be 100%. At one extreme, if the system-wide shock has reduced all banks' capital by 95% the entire system collapses. Note also that if the system-wide reduction is greater than 40% of Tier 1, the extent of the 'worst case' is similar for all three structures. But for system-wide shocks that reduce capital by less than 30%, the extent of contagion for Model II falls quite sharply, whereas the assets affected in the money-centre model remains around 40% of the total.

When we add the effect of a systemic weakening before the idiosyncratic shock, our results are somewhat consistent with the predictions of Allen and Gale (2000): there is no monotonic relationship between the degree of connection and the severity of domino effects. We do find that the money-centre structure tends to promote contagion. But following a systemic event that reduces Tier 1 capital by up to 40%, Model II (the model that incorporates the pattern of large

exposures) leads to less contagion than the benchmark model, which is closest to a complete structure.

Chart 8: Banking system assets affected by contagion in ‘worst case’ for various levels of systemic weakening



5.6 Stability through time

The preceding analysis used data from 2000 Q4. To assess whether this period is representative, we estimate the benchmark model biannually from 1999 H1-2001 H2. Chart 9 shows the number of banks that would be affected by contagion in our model at each point in time, and the number of banks that suffer a significant capital loss. Chart 10 is similar, but shows the assets affected as a percentage of total system assets. Both charts focus on the worst possible case and assume 100% loss given default. It is interesting to note that, in terms of assets affected, the period used for the detailed analysis above (2000 Q4) results in a relatively large amount of contagion but a relatively small effect in terms of other losses. In general however, the results are similar across time and contagion rarely occurs. When it does, it typically affects just a few small banks. Only in a few extreme cases, with high loss given default ratios, is the effect of contagion large. On average, the worst case triggers 5 additional bank failures, accounting for 19% of total assets, and a further 13 banks lose more than 10% of capital, where these banks account for a further 60% of total banking system assets.

Chart 9: No. of failures and weakness in ‘worst’ case





Chart 10: Percentage of total assets affected by failures and weakness in ‘worst’ case



At each period considered, the results are greatly dependent on the assumed loss given default ratio. Reducing loss given default to 60% means that contagion never affects more than 10% of total assets. Assuming a 40% loss given default (that considered by Furfine (1999)), never results in contagion affecting more than 1% of total assets.

6. Conclusions

The interbank market, while essential for transferring funds between banks, is a channel through which problems in one bank could be transmitted to other institutions. Analysing the potential for this channel of direct contagion is difficult owing to limited data on the network of bilateral interbank exposures. Within the constraints of the available data, this paper constructs three stylised estimates of exposures between banks operating in the United Kingdom in order to gauge the potential for direct contagion.

We study the effect of a narrow and extreme shock – the sudden failure of an individual bank. Our results suggest that, following such a shock, knock-on bank failures are rare. Where they occur, the average effects are typically quite small. For loss given default rates less than 50%, ‘domino’ failures do not affect more than 1% of aggregate banking system assets. But if loss given default is assumed to be 100%, in a few extreme cases the sudden failure of a single bank could trigger domino effects that cause the failure of banks accounting for more than 25% of total banking system assets.

In the face of a given shock, different assumptions about the interbank structure can imply different levels of contagion. In particular, incorporating partial information on large bilateral exposures into the model tends to reduce the risk of contagion between major UK-owned banks. At the same time, it increases the potential for contagion following shocks to overseas banks. On the other hand, assuming that large UK-owned banks act as a money centre for smaller banks and UK branches of foreign banks suggests the potential for spillover is higher.

Although knock-on bank failures are rare, an insolvency shock to a single bank can cause widespread weakening of the UK banking system. In many cases, a single shock can cause banks accounting for over half of total banking system assets to suffer losses exceeding 10% of their Tier 1 capital. Further, and unsurprisingly, the effect of contagion can be much larger if the idiosyncratic shock hits during a period when the banking system is already somewhat weakened, say during a period of unusually high market volatility.

The analysis is subject to important caveats. We do not have full data on the exposure of UK banks to banks located overseas. Given London's position as an important international financial centre, this means that a potentially important channel of contagion cannot be captured. In addition, this study does not include all categories of interbank exposure and makes no allowance for any netting agreements. But, this paper does capture the majority of unsecured exposures between banks operating within the UK system.

Appendix

Estimation of the benchmark exposure matrix proceeds in two stages. First, we choose a distribution for the interbank exposure that has maximum entropy. We then restrict this matrix to have zeros placed on the leading diagonal so that banks do not have exposures to themselves. As this matrix violates the adding-up constraints, we then find another matrix that gets as ‘close’ as possible to restricted matrix, but does satisfy the adding-up constraints. This amounts to minimising the cross-entropy between the two matrices, subject to the zero restrictions.

The concept of entropy, first used in the context of the interbank market by Sheldon and Maurer (1998), is common in information theory. When selecting a distribution for some event, the concept of entropy provides a means of discriminating between feasible alternatives. In particular, in the absence of prior information, one should select the distribution with the maximum entropy. As noted in the main text, entropy maximisation can be understood using the example of selecting a probability distribution for the outcome of rolling a six-sided dice. Without any prior information that the dice is loaded in some way, the most sensible distribution to choose is one that assigns an equal probability to each of the six possible outcomes. This provides a probability distribution for the outcome that maximises its uncertainty, ie the *entropy*, given available information (namely that the dice has six sides). So entropy maximisation allows us to select a unique distribution making full use of available information, without making any assumption about information that is not available.

To simplify the problem of maximising the entropy of the distribution of interbank exposures, we normalise the stock of interbank assets and liabilities to unity ($\sum_i a_i = \sum_j l_j = 1$) and express the problem as follows:

$$\begin{aligned} \min \sum_{i=1}^N \sum_{j=1}^N x_{ij} \ln x_{ij} \\ \text{subject to } \sum_{j=1}^N x_{ij} &= a_i \\ \sum_{i=1}^N x_{ij} &= l_j \\ x_{ij} &\geq 0 \end{aligned} \tag{A1}$$

The Lagrangian to this problem is given by

$$L(x, \lambda, \mu) = \sum_{i=1}^N \sum_{j=1}^N x_{ij} \ln x_{ij} - \sum_{i=1}^N \lambda_i \left(\sum_{j=1}^N x_{ij} - a_i \right) - \sum_{j=1}^N \mu_j \left(\sum_{i=1}^N x_{ij} - l_j \right) \quad (\mathbf{A2})$$

The first-order conditions show that the solution is given by

$$\begin{aligned} x_{ij} &= \exp\{\lambda_i + \mu_j - 1\} \\ &= \exp\left\{\lambda_i - \frac{1}{2}\right\} \cdot \exp\left\{\mu_j - \frac{1}{2}\right\} \end{aligned} \quad (\mathbf{A3})$$

Substituting this into the adding-up constraints from **(A1)** and re-arranging gives

$$\begin{aligned} \exp\left\{\lambda_i - \frac{1}{2}\right\} \sum_{j=1}^N \exp\left\{\mu_j - \frac{1}{2}\right\} &= a_i, \text{ and} \\ \exp\left\{\mu_j - \frac{1}{2}\right\} \sum_{i=1}^N \exp\left\{\lambda_i - \frac{1}{2}\right\} &= l_j \end{aligned} \quad (\mathbf{A4})$$

Note that because of the normalisation, it must be that

$$\sum_{i=1}^N \sum_{j=1}^N x_{ij} = \sum_{i=1}^N \exp\left\{\lambda_i - \frac{1}{2}\right\} \sum_{j=1}^N \exp\left\{\mu_j - \frac{1}{2}\right\} = 1 \quad (\mathbf{A5})$$

Combining this with the expressions in **(A4)** gives

$$\begin{aligned} \exp\left\{\lambda_i - \frac{1}{2}\right\} &= \sum_{i=1}^N \exp\left\{\lambda_i - \frac{1}{2}\right\} a_i, \text{ and} \\ \exp\left\{\mu_j - \frac{1}{2}\right\} &= \sum_{j=1}^N \exp\left\{\mu_j - \frac{1}{2}\right\} l_j \end{aligned} \quad (\mathbf{A6})$$

and substituting these into the solution **(A3)** provides the result

$$\begin{aligned}
x_{ij} &= a_i l_j \cdot \sum_{i=1}^N \exp\left\{\lambda_i - \frac{1}{2}\right\} \sum_{j=1}^N \exp\left\{\mu_j - \frac{1}{2}\right\} \\
&= a_i l_j
\end{aligned} \tag{A7}$$

As mentioned in the main text, this simple solution implies that a bank may have an exposure to itself. To overcome this problem, we construct a new matrix, X^0 , with elements⁽¹⁶⁾

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

Since this matrix may violate the adding-up constraints, the next stage is to find a feasible set of interbank exposures that gets close as possible to X^0 . This amounts to minimising the cross-entropy between the two. Formally,

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

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

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

$$x_{ij} \geq 0$$

Problems of this type can be solved using the RAS algorithm (see Censor and Zenios (1997)).

Given our estimate X^0 , the algorithm works as follows:

$$\text{Step 1 (row scaling): } x_{ij}^U \leftarrow x_{ij}^U \rho_{ij}^U, \text{ where } \rho_{ij}^U = \frac{a_i}{\sum_{\forall j | x_{ij}^0 > 0} x_{ij}^U}$$

⁽¹⁶⁾ Note that the zero restrictions only apply to individual banks. The aggregated groups are allowed to borrow from each other, since foreign-resident banks can borrow from each other.

Step 2 (column scaling): $x_{ij}^{U+1} \leftarrow x_{ij}^U \sigma_{ij}^U$, where $\sigma_{ij}^U = \frac{l_j}{\sum_{\forall i | x_{ij}^0 > 0} x_{ij}^U}$

Step 3: $U \leftarrow U + 1$, and return to step 1.

As mentioned in the main text, the alternative structures are estimated in a similar manner, but by adjusting the initial estimate X^0 .

References

Allen, F and Gale, D (2000), 'Financial contagion', *Journal of Political Economy*, Vol. 108, No. 1, pages 1-33.

Blåvarg, M and Nimander, P (2002), 'Interbank exposures and systemic risk', *Sveriges Riksbank Economic Review*, 2/2002.

Censor, Y and Zenios, S A (1997), *Parallel optimization*, Oxford University Press.

Elsinger, H, Lehar, A and Summer, M (2002), 'Risk assessment for banking systems', *Oesterreichische Nationalbank Working Paper no. 79*.

Fang, S C, Rajasekera, J R and Tsao, H S J (1997), *Entropy optimization and mathematical programming*, Kluwer Academic Publishers.

Furfine, C (1999), 'Interbank exposures: quantifying the risk of contagion', *BIS Working Papers no. 70*.

Gorton, G (1988), 'Banking panics and business cycles', *Oxford Economic Papers*, Vol. 40, pages 751-81.

Pain, D (2002), 'The provisioning experience of the major UK banks: a small panel investigation', *Bank of England Working Paper no. 177*.

Sheldon, G and Maurer, M (1998), 'Interbank lending and systemic risk: an empirical analysis for Switzerland', *Swiss Journal of Economics and Statistics*, Vol. 134, pages 685-704.

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*.

Wells, S (2002), 'UK interbank exposures: systemic risk implications', *Bank of England Financial Stability Review*, December, pages 175-82.