BANK OF ENGLAND

Working Paper No. 398

The sterling unsecured loan market during 2006–08: insights from network theory

Anne Wetherilt, Peter Zimmerman and Kimmo Soramäki

July 2010

BANK OF ENGLAND

Working Paper No. 398

The sterling unsecured loan market during 2006–08: insights from network theory

Anne Wetherilt,[1] Peter Zimmerman[2] and Kimmo Soramäki[3]

Abstract

We model the unsecured overnight market in the United Kingdom as a network of relationships and examine how the structure has changed over the recent period of crisis. Using established network techniques, we find strong evidence of the existence of a core of highly connected banks alongside a periphery. We find that membership of this core expanded during the crisis and suggest that this is due to a few intermediate banks becoming more connected. The widened reserve target bands may have also had an effect, by partially alleviating the need to manage reserve accounts close to a target and therefore allowing banks to exercise more discretion in forming relationships. However, there is an asymmetry between borrowers and lenders in the overnight market, with borrowers more reliant on the most established of the core banks during the crisis.

Key words: Network, topology, interbank, unsecured loan, systemic risk, financial stability.

JEL classification: G21, D85, E58.

---

(1) Bank of England. Email: anne.wetherilt@bankofengland.co.uk
(2) Bank of England. Email: peter.zimmerman@bankofengland.co.uk
(3) Helsinki University of Technology. Email: kimmo@soramaki.net

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We thank Claire Halsall for her help in preparing the data and Michael Jervis and Joanna McLafferty for their research assistance. We are also grateful to Ian Bond, Paul Chilcott, Rod Garratt, Jon Relleen, Gabriel Sterne, Goetz von Peter, seminar participants at the Bank of England and at the Bank of Finland, and an anonymous referee for their feedback on this paper. An earlier version of this work appeared in 'Simulation analyses and stress testing of payment networks', Bank of Finland (2009). This paper was finalised on 28 April 2010.

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

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

Financial markets in general can be viewed as networks, where buyers and sellers engage in repeated interactions. In particular, this analogy can be applied to money markets, as borrowers and lenders rely on each other for their daily funding needs. This paper examines the unsecured sterling overnight money market during a period which covers the crisis of 2007-08. A unique data set of individual trades in the UK CHAPS interbank payment system is used to construct a network of lending relationships between banks in the overnight market.

Network analysis of the overnight money market indicates that the structure of relationships between banks changed as the crisis unfolded. First, the data show that there is a core of a small number of banks which account for a large portion of overnight relationships. But, when concerns about counterparty risk increased, banks in the network diversified their relationships, reducing their dependence on the core. A possible explanation is that banks attempted to reduce funding liquidity risk by establishing more funding relationships.

Second, the analysis indicates that some of the observed changes in the network are asymmetric, in that they affected borrowers more than lenders. The paper argues that this asymmetry may be unique to the overnight market where increased counterparty risk is a concern for borrowers, but perhaps less so for lenders. This may be because many borrowers hope to roll overnight loans for an extended period. Thus borrowers may be keen to establish a relationship with one or more core counterparties, who are more likely to be able to provide this funding on a daily basis.

Third, the paper also suggests that changes to the reserve regime in September 2007 made liquidity management more straightforward, because banks had less strict end-of-day targets to meet. Banks therefore had much more discretion about whether to participate in the overnight market, and who to trade with. The network data show a drop in the probability of forming a relationship at this time.

The paper does not attempt to measure whether the impact of market events was greater or less than the impact of policy events. This question could be important when attempting to gauge the effect of central bank actions.

The analysis is confined to the overnight unsecured market, reflecting data availability. It does not examine to what extent this market was affected by changes in the term markets and in the secured markets. Hence, this research does not permit conclusions about the resilience of liquidity in the money markets in general, or the case for any changes in the underlying infrastructure. These issues are left for future research.

1 Introduction

In their efforts to improve the assessment of vulnerabilities in the financial system, economists have started to view the network of financial markets and institutions as a complex, adaptive system (Haldane (2009)). Using language and statistical methods developed in other sciences, they are aiming to develop a better understanding of how weaknesses can develop, how these networks behave under stress, and what mitigating actions can be taken.

Such network methods have in recent years been applied to high-value payment systems, and have proved a useful tool to model the intraday payment flows between settlement banks. This has allowed researchers to assess the robustness of payment flows, both during normal circumstances and following operational outages (Becher *et al* (2008)). Similar analysis has been used to model intraday flows in money markets (first by Boss *et al* (2004)). Again this has proven a useful method to evaluate the behaviour of money market participants, with recent work (Bech and Atalay (2008)) relating network measures to changes in the federal funds rate.

In this paper, we model the unsecured overnight market in the United Kingdom as a network of relationships and examine how it has changed over the recent period of market turmoil. Using a unique data set of individual trades in the UK CHAPS interbank payment system, we employ established network techniques to assess how the network of lending relationships between individual settlement banks changed during 2006-08.

In this paper, we use methods in network theory to examine whether banks concentrated or diversified their relationships as the crisis unfolded, and which counterparties were affected by these changes. We also investigate whether concerns about counterparty credit risk affected borrowers and lenders in different ways during the crisis.

We argue that this asymmetry may be unique to the overnight market where increased counterparty liquidity risk is a concern for banks seeking lenders, but perhaps less so for those looking for borrowers with whom to place cash. If a lending counterparty develops liquidity problems, it may choose to cut down on the amount of funding rolled. But a borrowing counterparty would need to experience very severe liquidity problems before failing to repay its overnight loan, since that would be a credit event. Therefore we want to see whether banks looking for a lender became more likely to choose a counterparty from the core than those looking for a borrower.

Our results indicate that although total overnight unsecured sterling activity was not lower during the crisis compared to pre-crisis levels, we do find a reduction in the number of bilateral relationships, while dependence on a small number of core banks fell. This core did exist pre-crisis, but we suggest that it expanded during the crisis phases as participants became more aware of counterparty credit risk and sought to diversify.

While we do not attempt to assess whether the impact of market events was greater or less than the impact of policy events, we observe the clearest network changes after the widening of reserves target bands in September 2007. We suggest that this policy change may have given banks more flexibility in their ability to adjust their liquidity management, and hence allowed them to focus on counterparty risk.

The remainder of the paper is organised as follows. After providing an overview of the sterling money market (Section 2), we discuss how to visualise the market as a network and present some statistics (Section 3). Section 4 interprets these statistics in the broader financial context, while Section 5 looks in more detail at the most intense period of turbulence. Section 6 concludes.

2 The UK money market

This section provides an overview of the UK unsecured overnight money market operating framework and presents a statistical overview.

2.1 Overview of the money market turmoil and Bank of England actions

Conditions in global money markets were unusually stressed between Summer 2007 and the end of 2008. Market liquidity fell sharply, particularly at maturities beyond one month, and spreads over policy rates widened. Many banks therefore found it difficult to access longer-term funding on acceptable terms. In the United Kingdom too, term money markets saw a fall in liquidity and term spreads widened compared to pre-August 2007 levels. At shorter dates, however, market activity has been less impaired and rates have stayed closer to policy rates.[1]

During this period central banks continued to provide liquidity via the normal channels (eg open market operations, standing facilities) to keep short-term market rates close to policy rates. In addition, central banks responded to the continued strains in money markets by introducing a range of extraordinary measures such as offering longer-term funding, widening collateral lists and broadening the range of counterparties (Committee on the Global Financial System (2008)).

The UK money market framework in place between May 2006 and March 2009 allowed participating banks to choose their own reserve targets.[2] Each bank's average reserve balance over the maintenance period (which lasts four or five weeks between regular Monetary Policy Committee meetings) would have to be within a certain band around this target in order to be remunerated at Bank Rate. If the balance fell short, the bank would be forced to borrow the shortfall from the Bank of England at a penal rate. If the balance was over the target, the bank would earn no interest on the excess.

In response to the strains in sterling money markets, the Bank of England widened the range around banks' reserve targets in September 2007, thereby giving banks greater flexibility in managing their reserve accounts.[3] In addition, banks in aggregate increased their reserve targets

[1] See Bank of England (2008a), Chart 1.5; Bank of England (2008b), Chart 23; and Bank of England (2008c), Chart 2.
[2] Bank of England (2005) pages 211-20 describes the regime in more detail.
[3] Annex 5 summarises the changes in reserve target bands throughout the period examined in this paper.

significantly after August 2007.[4] At the time of writing, the reserve targeting regime had been suspended since 5 March 2009.[5]

The Bank of England's monetary framework also allows banks to borrow or lend on an overnight basis using the standing facilities. Initially, information regarding use of these overnight standing facilities was published on the following day. During Summer 2007, stigma became attached to the borrowing facilities and banks ceased to use them. In October 2008, the Bank of England announced reforms to the facilities in order to reduce this stigma.[6]

In the remainder of this paper, we examine how these developments in the markets and the monetary framework affected interbank relationships in the sterling overnight market.

2.2 *A first look at the data*

We use the data on payments in the large-value payment system CHAPS Sterling available to the Bank of England in its role as operator of the underlying real-time gross settlement processor. From the raw data, it is difficult to distinguish cash payments from loan payments (either advancement of principal or repayment of principal plus interest). Following Millard and Polenghi (2004), we use a variation of the algorithm developed by Furfine (1999) which identifies pairs of payments on consecutive days that could be interpreted as overnight loan advances and repayments. Annex 1 discusses the construction of the data set in detail.

It should be noted that the CHAPS database includes only payments made between the clearing banks which were CHAPS members during the period examined. It therefore excludes loan payments between two customers of the same settlement bank, which are settled across that bank's books in commercial bank money rather than in CHAPS. Data on these are not available. Furthermore, we cannot distinguish between settlement banks and their customers in the data. Our data indicate that payments relating to overnight unsecured activity (advances and repayments) account for about 20% of CHAPS values; however, this may be an underestimate of the true size of the market due to the internalisation issue.

Our sample period runs from 18 May 2006 to 16 December 2008. Following Borio (2008), we take the start of the crisis as 9 August 2007. On this day, interbank markets in the United Kingdom and other countries came under severe and lasting stress. BNP Paribas blocked withdrawals from three of its investment funds, and there were rumours of losses at other banks. In mid-August problems became apparent at Northern Rock, a UK mortgage bank which had relied on funding in the wholesale markets, leading to it seeking emergency liquidity assistance from the Bank of England in September.

We consider a pre-crisis phase 1 (18 May 2006 to 8 August 2007) and a crisis phase 2

[4] See Bank of England (2008c), Chart 30.
[5] See Bank of England (2009).
[6] See Bank of England (2008d).

(9 August 2007 to 16 December 2008). There are other events during the crisis phase which might be expected to have had an impact on the money markets, the most significant of which are listed below:

- On 13 September 2007, the Bank of England announced an increase in the bands around the target at which reserves would be remunerated at Bank Rate. These bands remained wider throughout the crisis phase (see Annex 5).
- On 21 April 2008, the Bank launched the Special Liquidity Scheme, which allowed banks to exchange illiquid assets for UK Treasury bills.
- On 15 September 2008, Lehman Brothers defaulted, triggering a period of acute stress.
- On 8 October 2008, the UK government announced a recapitalisation scheme for UK banks.

With these key dates in mind, we take a first look at the overall trends in the overnight sterling market. The diamonds in Chart 1 mark the key events described above. Both values and volumes were higher in the initial period of the crisis until mid-September 2007 – around the time of the widening of reserve bands – after which activity drops. It picks up in March 2008, but falls again after the announcement of the Special Liquidity Scheme in April. After the Lehman default in September 2008 activity increases again. Value of trades is more volatile than volume.

Chart 1: Total daily advances in overnight unsecured sterling – 21-day rolling average[7] (value £m on left-hand axis; volume on right-hand axis)

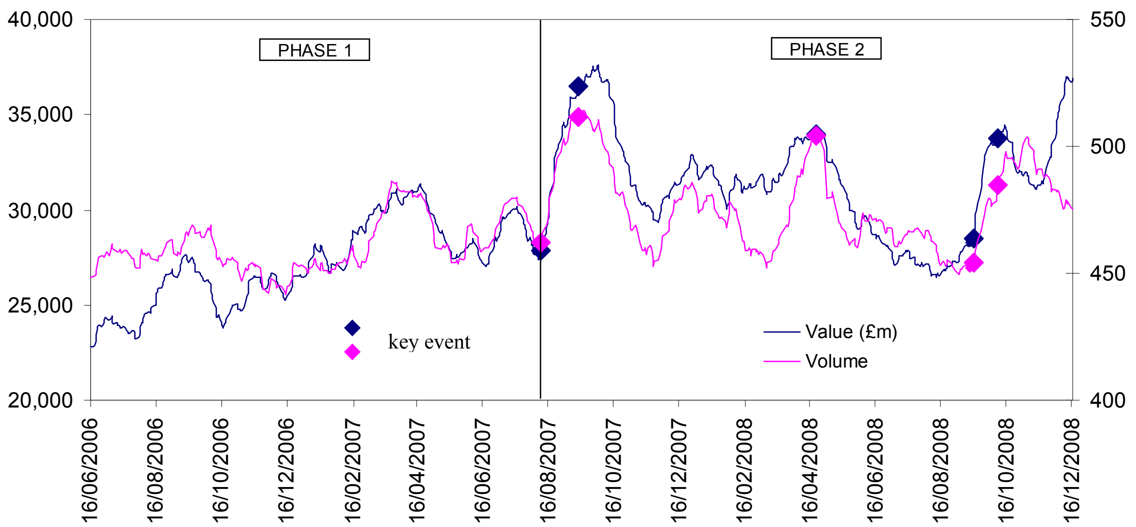

Since a new settlement member joined CHAPS in October 2007 and another ceased direct membership in September 2008, Table 1 includes an average daily figure per bank. It shows that average daily value of loan advances was higher during the crisis phase, increasing from £27.0 billion in phase 1 to £31.8 billion in phase 2. Despite the presence of a thirteenth bank during most of phase 2, average daily value per bank was significantly higher as well. The volume of loan contracts increased too during the crisis phase, but not to the same extent.

[7] We plot a rolling average rather than the actual data as the series is highly volatile with no obvious breakpoints, making it difficult to pick out long-term trends.

During phase 2 the number of contracts per bank per day was actually slightly lower than during phase 1.

Table 1: Summary statistics

	All	Phase 1	Phase 2	t-stat	p-value
Days in phase	656	311	345		
Average number banks per day	12.37	12	12.70		
Daily average value £m	29,522	27,015	31,782	14.62	***0.0%
Daily average value per bank £m	2,387	2,251	2,502	8.76	***0.0%
Daily average volume	469	460	477	6.98	***0.0%
Daily average volume per bank	37.9	38.3	37.5	-3.62	***0.1%

In Table 1 and elsewhere in this paper, the t-statistics are obtained from Welch's two-tailed t-test of the hypothesis that the daily mean figure in phase 2 is significantly different from the mean in phase 1.[8] The results so far tell us that overall overnight loan activity increased post August 2007. Of course, our data allow us only to comment on the unsecured overnight markets and not loan activity in general.

3 The market as a network

In this section, we briefly survey the existing literature and undertake a statistical examination of the network describing the sterling unsecured overnight market. We focus on network-wide measures rather than studying individual nodes, since the phases identified are defined by market-wide and not bank-specific events. We do not examine measures relating to path length since these do not have an obvious interpretation in the money markets, particularly given the small size of the CHAPS settlement bank community.

3.1 Empirical financial networks – a literature review

The study of networks has been applied to a wide range of fields, such as social interactions, epidemiology and the worldwide web. A large amount of work from the physics community has focused on the structure of complex networks – that is, those displaying features that are neither regular nor purely random.[9] As mentioned in Section 1, economists have recently started using these methods to analyse the patterns in payment and loan flows, and to assess the stability of these networks.

Boss *et al* (2004) and Inaoka *et al* (2004) were among the first to use network topology in empirical studies of interbank markets, examining the Austrian and Japanese systems respectively. These papers confirm that topology measures are suited to describe financial networks in general and their resilience to shocks in particular.

[8] Welch's t-test uses the null hypothesis of equality of means. It requires the underlying observations to be normally distributed, but due to the large number of observations in each phase we can appeal to the central limit theorem. In all our tables, one asterisk denotes significance at the 90% level, two asterisks significance at the 95% level, and three asterisks significance at the 99% level. Throughout this paper, we say a result is 'significant' if the null hypothesis is rejected at a 95% confidence level.
[9] See, for example, Dorogovtsev and Mendes (2003), Albert and Barabási (2002) and Newman (2003).

Soramäki *et al* (2007) examine the network topology of interbank payments across the Fedwire Funds Service in the United States. They find that participation in the payment system fell following the attacks of 11 September 2001, both in terms of number of active banks and number of transactions. There is evidence of less co-ordination between banks, likely a result both of the operational problems faced by some participants and the responses by others to the resulting liquidity problems. However, once the operational problems were over, activity rose to above-average levels as banks settled their backlogs of payments.

Becher, Millard and Soramäki (2008) examine interbank payments across CHAPS Sterling. They investigate the impact on the network of an operational outage of a major settlement bank, and find that the network topology did not look significantly different during the outage day. Furthermore they find that non-stricken banks were able to manage their liquidity effectively, so that payment flows between them were much as normal.

Papers by Pröpper *et al* (2008) and Lublóy (2006) use payment data from the Dutch and Hungarian systems, respectively. Pröpper *et al* find that payment activity has been considerably higher since the market turmoil of 2007 began, but network properties have been less affected. Lublóy looks at the permanence of linkages over time and finds that, although there are relatively few bilateral lending relationships which exist every single day, those that do account for the majority of payment orders by value.

Topology measures are also used in a number of recent papers which look at the overnight money market. Iori *et al* (2008) study the Italian money market and find that the network has changed over time. Here, lending has become more concentrated while borrowing has become more diluted. They further document that a few large banks borrow from a large number of small counterparties. Bech and Atalay (2008) and Bech and Rørdam (2008) examine overnight money markets in the United States and Denmark, respectively.

3.2 Defining network terminology

In this subsection we briefly explain the network terminology that we shall use throughout this paper. The definitions that we use are mostly consistent with the literature.

A **network** can be expressed as a set of **nodes** and **links** between those nodes. Links can be **directed** (so that the link from node i to node j is different from the link from j to i) or **undirected**.

Each link can be assigned a value, usually non-negative. A value of 0 means the link does not exist. If each link can only take on values 0 or 1 then we say the network is **unweighted**. If a link can take on more than one strictly positive value then we say the network is **weighted**. For example, in the money markets an unweighted network can be used to determine the presence of bilateral relationships, while a weighted network might be used to analyse value or volumes traded.

Given n nodes, there can be up to $n(n\text{-}1)$ links in a directed network and half this number in an undirected network. The total value of all links in the observed network is called the **degree**. A **complete** network is one where every possible link exists; that is, every link has non-zero value.

A **subnetwork** is a subset of the nodes, with the links between them. A **component** of a network is a subnetwork where no links exist between nodes in the subnetwork and those in its complement.

In what follows, we examine to what extent relationships between counterparties have been affected by the market and policy events over the course of the crisis. To do so, we model the series of unsecured loan payments between banks as an evolving network. Each bank is represented by a node, and each loan advance is a directed link between nodes from the lending bank to the borrowing bank. We take a business day as our unit of time, and do not look at intraday networks. This is because a bank's reserve management – and thus lending and borrowing – behaviour is driven to a large extent by the end-of-day target rather than intraday needs, so a day is a natural unit of time.

3.3 Using an unweighted network

In the literature, practice varies as to whether unweighted networks (where the value of each link is 0 or 1 according to whether a relationship exists or not) or weighted networks (where each link takes on a real number value such as the value or volume lent) are examined. Most of the papers we cite examine the topology of the unweighted network, considering weighted links only for non-topological measures such as total amounts lent by a bank.[10]

In this analysis, we look at unweighted networks – that is, a link exists from bank i to bank j if i lends to j, regardless of the actual value lent (though note we only detect loan advances of £1 million or more – see Annex 1). Therefore we are examining the *existence* of relationships in the overnight market, rather than the weight of these relationships in terms of values or volumes of loan activity. We decide to do this because we are interested in the extent to which credit/liquidity events influence the counterparties than banks choose.

If we were to look at the weighted network, we would need to take account of differences in banks' sizes. For instance, we might expect the links between the biggest banks to have the highest weights regardless of market events, so comparison of absolute values may not tell us much about how relationships change. Furthermore this analysis would be more sensitive to the behaviour of indirect members of the system, which we cannot directly observe. Some indirect members do not have reserve accounts and thus do not have the same end-of-day targets in sterling: these may exhibit different behaviour to settlement banks which do.

On a more practical level, the measures cannot all be calculated easily for a weighted network. For example, connectivity requires some concept of the maximum value of a link. In an unweighted network this is of course 1, but there is no theoretical upper limit if we weight by value or volume. This is sometimes dealt with in the literature by using the empirical maximum; however, this can skew values if the maximum is an outlier. This can be avoided by using a cut-off at a particular value (for example, the 95th percentile), but the choice of this value is somewhat arbitrary and the results may be sensitive to this.

[10] There are some exceptions – for example Bech and Atalay (2008) define and examine a weighted version of reciprocity.

The core can be more easily computed on a weighted network: we simply need to assume a distributional form on the links (just as we do in the unweighted case where we use a Bernoulli distribution). We leave this for future research.

3.4 Network graphics

In this section we present the network graphically, using the software Pajek.[11] As an illustration we display here six pictures for each of the key dates described in Section 3, plus the start and end of the sample period. In these pictures, a link is represented by a black line between two nodes. The arrow shows the direction of the link. Where a link exists in both directions, there is an arrow on both ends of the line. Note that the order of the nodes has been permuted between pictures to preserve the anonymity of individual banks.

At first sight, the interbank network did not change greatly. On most days the network was connected as a single component, indicating that most banks were active in the market and lending to one another. It appears that a small number of banks dominated overnight activity, though we cannot tell from the pictures whether the identity of these 'core' banks remained unchanged. It looks as if some of the peripheral banks may have become better connected in the last two pictures, but we cannot tell whether this was significant. In the rest of this section, we explore more sophisticated techniques for examining these networks.

[11] For more information about Pajek, see http://pajek.imfm.si/doku.php?id=pajek.

Charts 2.1-2.6: Network graphics on selected dates

(2.1) 18 May 2006

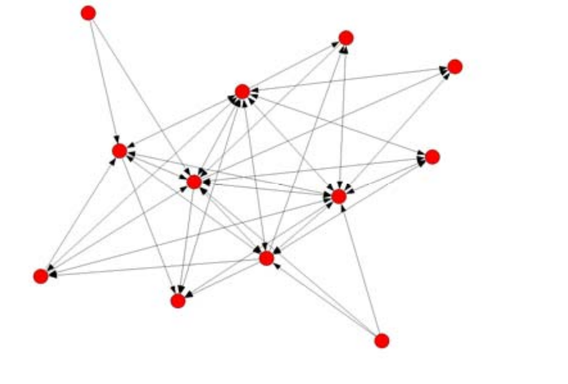

(2.2) 9 August 2007

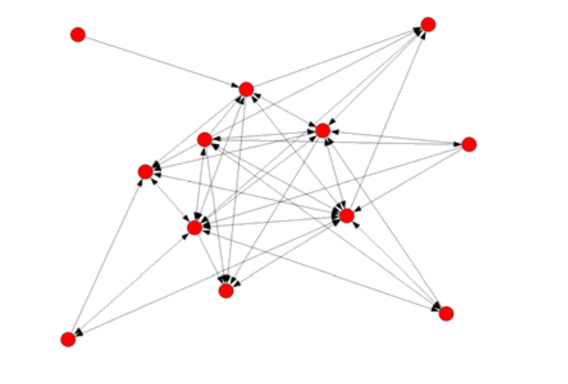

(2.3) 21 April 2008

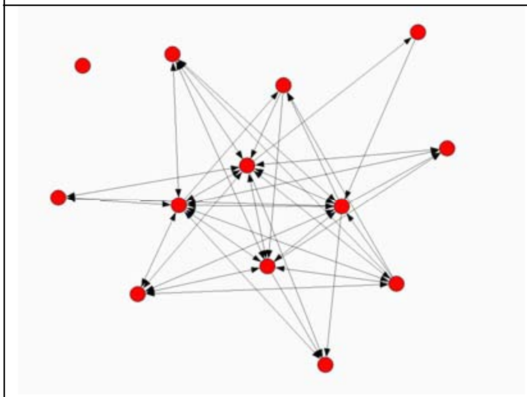

(2.4) 15 September 2008

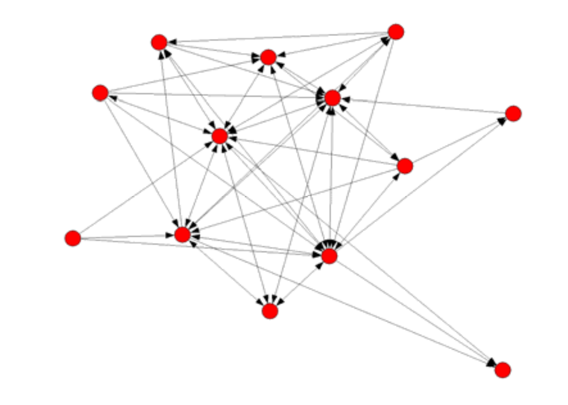

(2.5) 8 October 2008

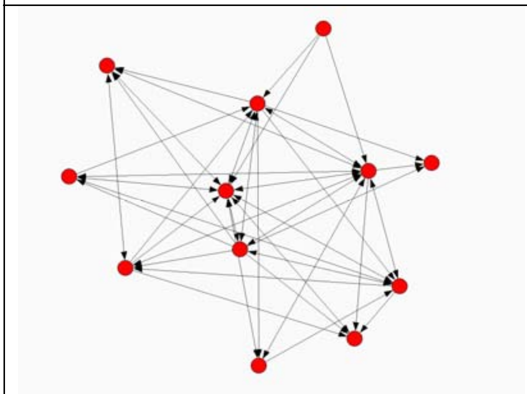

(2.6) 16 December 2008

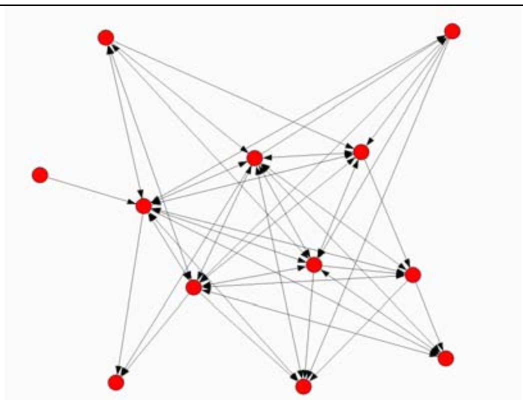

Key

Loan (arrow points to borrower)

Bank

3.5 Importance of individual banks

These pictures indicate that there may have been a few banks accounting for a relatively large proportion of the links. To look at this in more detail, on each day we count the number of links (both in and out) that each bank accounted for, and rank them in order. We then divide by the sum of these to obtain a proportion.

Chart 3 shows how the proportion of links accounted for varies by rank of bank. The vertical axis denotes the average proportion of links across the whole period (or, in the case of the 13th bank, the period in which there are 13 banks in the network). Note that this chart does not consider the identity of each bank – on each day, we take whichever bank is ranked at the appropriate level. The vertical axis shows the average proportion over the entire sample period.

Chart 3: Average proportion of links according to rank of bank

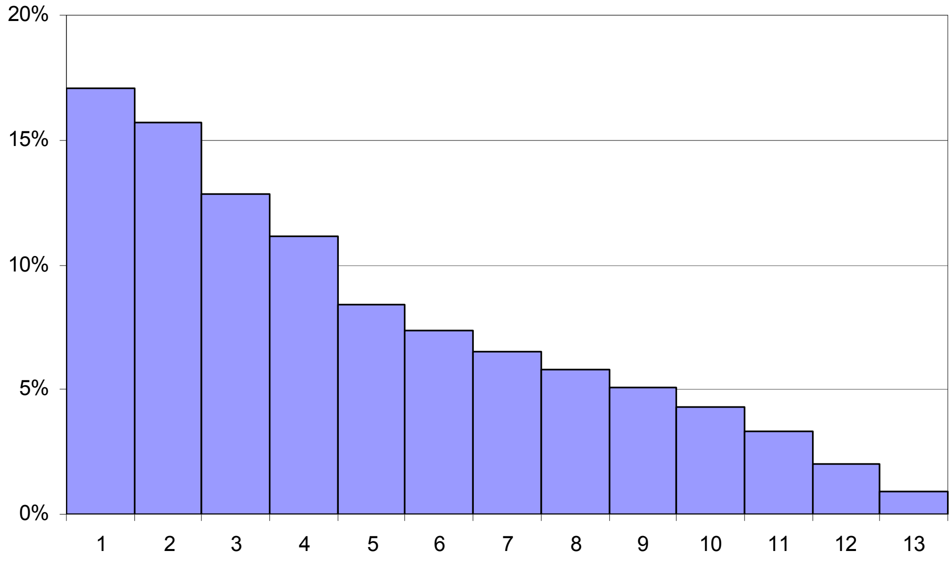

The slope of the chart is fairly gentle, suggesting that there may be no unique 'cut-off' point between the more and less-connected banks. There is a noticeable drop between the second and third banks, and between the fourth and fifth, so it may be that there was a 'core' of size 2 or 4. But we need to examine who the less-connected banks linked to: for example, if they linked to each other but not to the more-connected banks then we would have a network that resembles two clusters rather than a core and periphery.

3.6 Existence of the core

We use a maximum likelihood approach similar to Čopič *et al* (2009) to define the core. Annex 2 provides a full description, but we outline the process briefly here. Consider all possible partitions of our set of nodes into two subsets. Each such partition implies a partition of links into four subsets. Assume that, within each of these four subsets, links are formed independently and with equal probability. We then define the 'core' allocation as the partition which maximises the likelihood of producing the observed network. There will actually be two such allocations, since any given partition will yield the same likelihood as its complement (ie we could switch the labels of 'core' and 'non-core'). To select between these, we impose the restriction that the probability of links between core nodes must be no smaller than the probability of links between non-core nodes.

A maximum likelihood ratio test suggests that the core is significant on each day – that is, the network can be better described by a core/non-core partition than by assuming all nodes are identical. This suggests that a core existed throughout the sample period. This might be thought to be a trivial result given the concentration in the top-four in the UK payments system by value, but Chart 3 suggests that this concentration may not be so apparent in terms of number of relationships.

We first look at the size of the core and changes in underlying membership. Chart 4 below shows the number of members of the core on each day over the period, according to our likelihood-maximising core allocation. The blue diamonds denote the first days of the two phases and the other key dates mentioned on page 7.

Chart 4: Size of the core

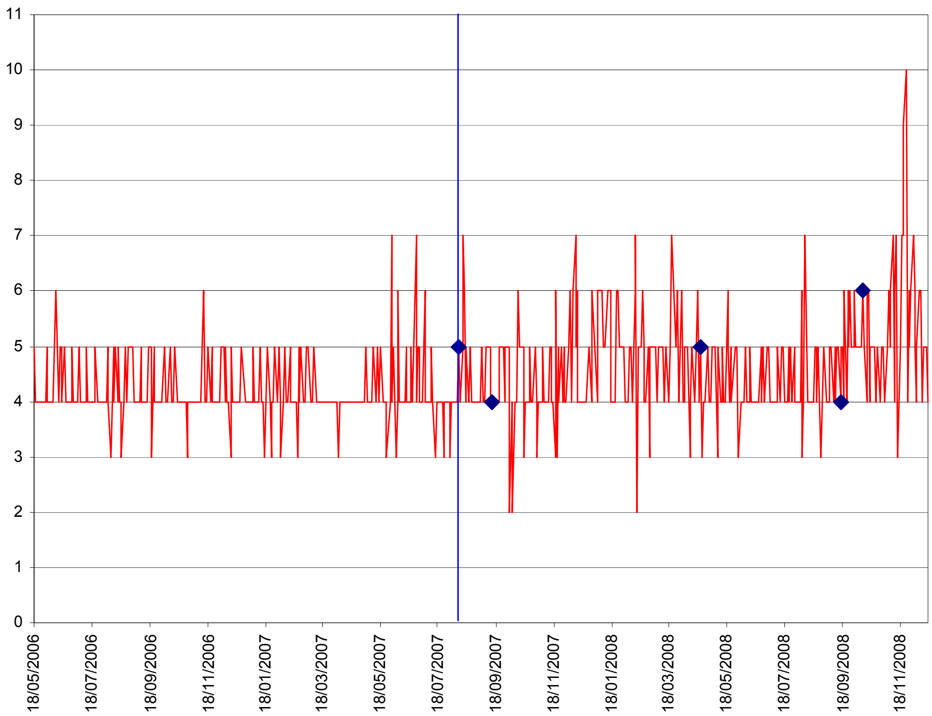

The core was of average size 4.18 during the pre-crisis phase and 4.67 during the crisis, which is a significant increase at the 1% level based on Welch's t-test. The chart shows that there was a sudden leap in the size of the core following Lehman's default in September 2008, reaching 10 (out of 12 banks) on 24 November 2008. This suggests that, not only did overall connectivity increase in the wake of the Lehman collapse (see Section 3.7) but that this increase was fairly well distributed across the system. It was not only the core banks that established more links; the formerly non-core banks also became better connected.

This result is much starker than that suggested by Chart 3, which did not show a big difference between the top four and the rest, even though that chart re-ranked the banks every day. This shows that it is not enough to look at how connected a bank is; one needs to consider who the connections are to.

3.7 *Linkages between the core and non-core banks*

We can further investigate whether this interpretation is correct by looking at the probabilities of links forming among the different types of banks. The partition of the network into two sets (core and non-core) implies a partition of the links into four sets. We define p_{CC} as the probability that a link exists between two core banks, p_{CN} as the probability that a core bank lends to a non-core bank, and similarly for p_{NC} and p_{NN}.

Chart 5 shows the values of p_{CC}, p_{CN}, p_{NC} and p_{NN} each day. First, note that p_{CC} is higher than the other measures and p_{NN} lower. This points to a network characterised by a core and periphery, as described above. If the network was actually composed of two clusters – so that links form with high probability within a cluster but a low probability across clusters – then we would expect p_{NN} to be higher than p_{CN} or p_{NC}.

p_{CC} was on average 0.98, implying that the subnetwork of links between the core banks is almost complete (in fact, it was complete on 79% of days in the sample period). But Table 2 shows that p_{CC} dropped during the crisis phase: the peripheral banks became connected enough to join the core, but they were still less connected than the core banks were previously. Although all four probabilities fall, the drop in p_{NN} is slightly less significant than the other three.

p_{CC} reached a low of 0.53 on 24 November. On this day, p_{CN}, p_{NC} and p_{NN} were all equal to zero. This suggests that, as most banks were in the core on this day, links between core members themselves were sparser than usual. The two banks which remained outside the core had no overnight activity at all that day.

Chart 5: Estimated probability of links between core and non-core nodes

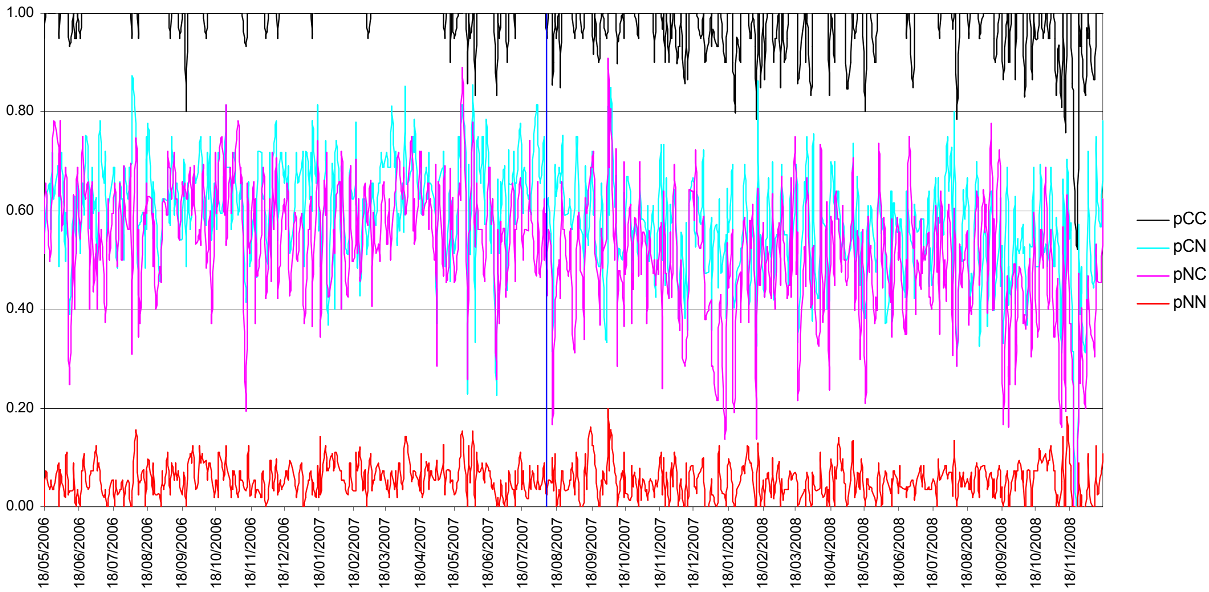

Table 2: Changes in the probability of links during the crisis

		Phase 1	Phase 2	t-stat	p-value
p_{CC}	Mean	0.9937	0.9684	-7.41	***0.0%
	Variance	0.0006	0.0034		
p_{CN}	Mean	0.6298	0.5461	-10.28	***0.0%
	Variance	0.0095	0.0123		
p_{NC}	Mean	0.5781	0.4705	-11.58	***0.0%
	Variance	0.0107	0.0178		
p_{NN}	Mean	0.0581	0.0507	-2.65	***0.8%
	Variance	0.0011	0.0014		

3.8 Probability of core membership for individual banks

Chart 6 below examines the probability of each bank being in the core during each phase. For instance, in phase 2 there were six banks which were in the core on at least one but less than 25% of the 656 days examined. We see that two banks were never members of the core, and another two were members only during the crisis phase. Four banks were members of the core more than 75% of the time. Let us call these four the ‘regular core’ banks.

The histogram moves to the right in phase 2, meaning that membership of the core for individual banks appears to have become slightly more likely during the crisis phase. However, there is still a clear distinction between the four regular core banks and the rest. Note that the heights of the bars sum to 12 during the pre-crisis phase and 13 during the crisis, due to changes in the underlying membership of CHAPS.

Chart 6: Probability of core membership

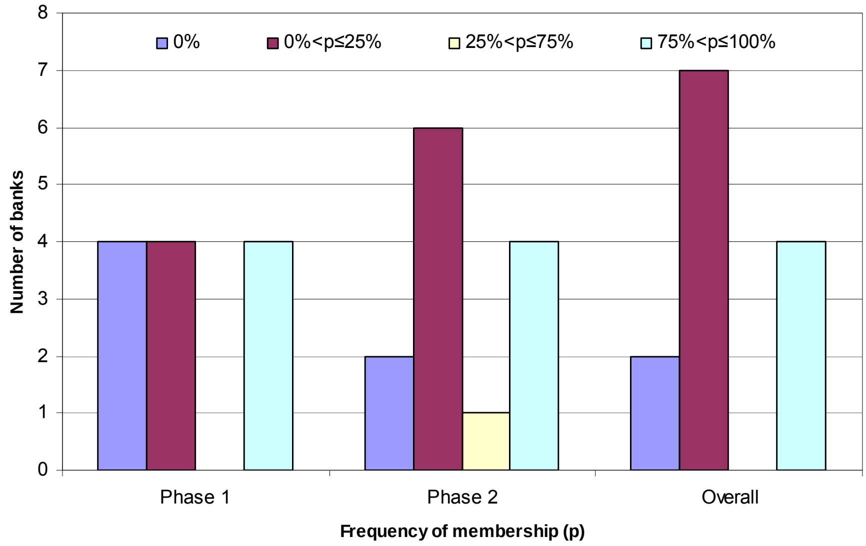

This shows that there is some stability in the core. The same four banks are members the majority of the time, while the others are there only a fraction of the time.

3.9 *Importance of the core banks*

The analysis so far confirms the existence of a core and examines how its membership changed. It would be useful to establish whether dependence on the core changed too. Define $Q_l^{(t)}$ as the proportion of links on day t for which the lender is a core bank. Similarly, $Q_b^{(t)}$ is the proportion of links with a core borrower. They reflect the degree to which lenders and borrowers respectively depend on the core banks (see Annex 3 for a mathematical definition).

Table 3: Dependence on the core

		Phase 1	Phase 2	t-stat	p-value
$Q_l^{(t)}$	Mean	0.5770	0.5976	3.76	***0.0%
	Variance	0.0037	0.0062		
$Q_b^{(t)}$	Mean	0.6080	0.6474	6.32	***0.0%
	Variance	0.0037	0.0093		
$Q_b^{(t)} / Q_l^{(t)}$	Mean	1.0608	1.0867	2.83	***0.5%
	Variance	0.0143	0.0129		

Table 3 shows that dependence on the core increased significantly during the crisis phase 2 for both borrowers and lenders. Furthermore, it increased by more for borrowers than for lenders, indicating that borrowers chose to – or were forced to – choose counterparties from the core.

Another point to note is that borrowers were more reliant than lenders on the core, since in expectation $Q_b^{(t)} > Q_l^{(t)}$. This implies that the number of links from the core to the non-core exceeded those in the other direction – in other words, non-core banks were more likely to borrow from the core than lend to it.

3.10 Importance of the regular core banks

Given the changes in the size of the core illustrated in Chart 4, increased dependence on the core is hardly surprising. For example, it seems extremely likely that the proportion of links connecting to the core would be much higher on 24 November 2008 because almost every bank is in the core. Therefore let us instead examine dependence on the four 'regular core' banks described in Section 3.8. We define $R_l^{(t)}$ and $R_b^{(t)}$ as the proportion of links which have a regular core bank as the lender or borrower respectively. This is analogous to the definitions of $Q_l^{(t)}$ and $Q_b^{(t)}$ in Subsection 3.9; again Annex 3 provides a more rigorous definition.

Table 4: Dependence on the regular core banks

		Phase 1	Phase 2	t-stat	p-value
$R_l^{(t)}$	Mean	0.5594	0.5370	-8.30	***0.0%
	Variance	0.0011	0.0013		
$R_b^{(t)}$	Mean	0.5919	0.5838	-2.79	***0.0%
	Variance	0.0015	0.0013		
$R_b^{(t)} / R_l^{(t)}$	Mean	1.0640	1.0932	3.25	***0.1%
	Variance	0.0138	0.0126		

Table 4 shows that dependence on the regular core banks actually fell during the crisis phase. Taken together with Table 2 and Chart 4, this suggests that banks moved away from these four and formed links with other counterparties instead, which then moved into the core as they become more connected. Note that borrowers were more dependent on the regular core than lenders.

3.11 Connectivity

Connectivity is defined as the proportion of potential links that exist. Thus, for a network with n nodes, connectivity is equal to degree divided by $n(n\text{-}1)$. We need to normalise in this way since the number of nodes changes over our sample period. Connectivity can be thought of as the average probability that any given link is formed.

Chart 7 below shows how the connectivity of the network changed over the sample period. The black line shows the backward-looking 21-day rolling average, while the blue diamonds denote the new phases and key dates. There appears to have been a very large drop around the time when the reserve bands were widened in September 2007, after which connectivity remained at a lower level before a considerable rise starting around the beginning of September 2008. From the start of October it was fairly stable, at around the same level as in phase 1.

Chart 7: Connectivity of the network over the period

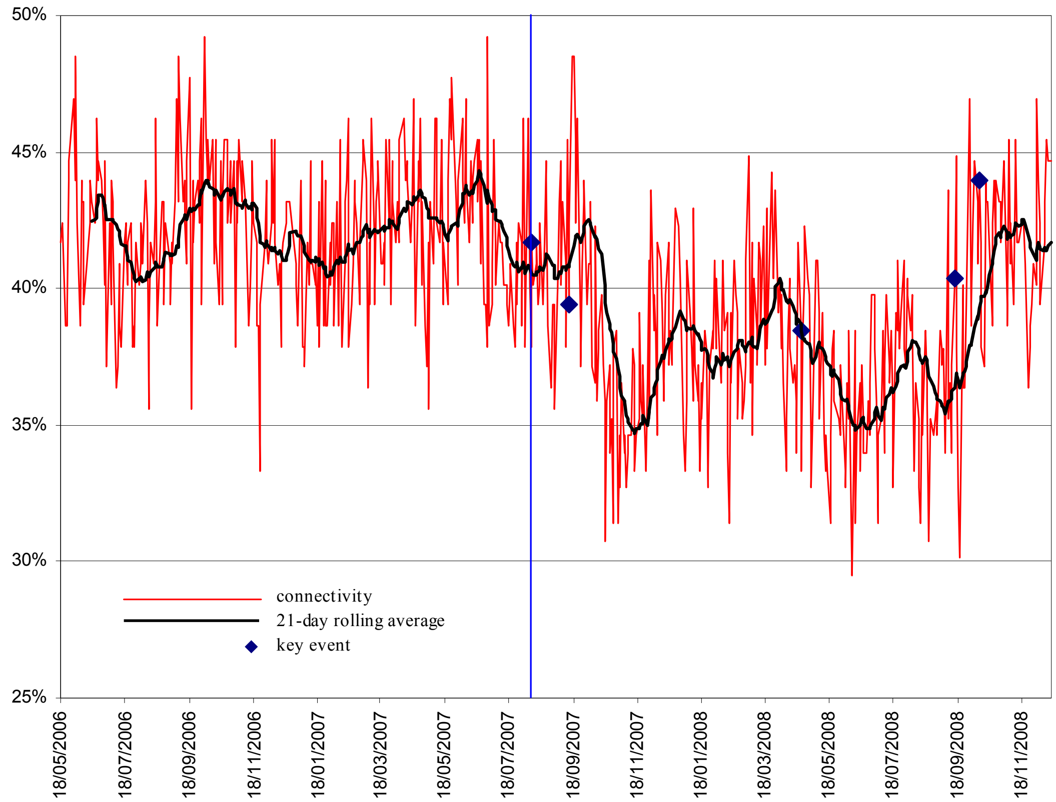

Table 5 below shows that mean connectivity was significantly lower during the crisis phase. Our results in Section 3.10 suggest that a disproportionate number of the links which disappeared involved the four regular core banks.

Table 5: Mean connectivity during each phase

	Phase 1	Phase 2	t-stat	p-value
Mean	0.4210	0.3845	-14.70	***0.0%
Variance	0.0007	0.0013		

Lower connectivity during phase 2 explains why all four of p_{CC}, p_{CN}, p_{NC} and p_{NN} were found to significantly decline. We can factor this out by dividing each of these by connectivity on the day in question: this would give us a measure of how links in the four subnetworks changed relative to the overall probability of forming links.

Table 6: Changes in the relative probability of links during the crisis

		Phase 1	Phase 2	t-stat	p-value
p_{CC}/χ	Mean	2.3704	2.5412	9.35	***0.0%
	Variance	0.0274	0.0846		
p_{CN}/χ	Mean	1.4983	1.4261	-3.54	***0.0%
	Variance	0.0534	0.0842		
p_{NC}/χ	Mean	1.3706	1.2279	-6.40	***0.0%
	Variance	0.0478	0.1184		
p_{NN}/χ	Mean	0.1366	0.1305	-0.93	35.5%
	Variance	0.0057	0.0086		

Table 6 shows that, relative to the entire network, links between core banks actually became more likely. However, links between the core and the non-core banks became significantly less likely. This suggests that the non-core banks became even more peripheral, relative to the rest of the network.

3.12 Reciprocity

Reciprocity is defined only for directed networks, and is the probability that a link from one node to another exists, given that a link exists in the opposite direction between the same two nodes.[12] A mathematical definition is given in Annex 3. In the context of the network of the overnight market, this can be thought of as the strength of bilateral relationships. If reciprocity is high, then banks tend to use the same counterparties for both lending and borrowing.

In fact, reciprocity did not rise during the crisis phase. But this may reflect the overall fall in the probability of any relationship after October 2007 (as shown in Section 3.6), rather than reduced willingness to enter into borrowing and lending relationships with the same counterparties. To separate the two effects and ascertain whether reciprocity rose relative to connectivity, we compute a new measure called normalised reciprocity. This is calculated as reciprocity on day t divided by connectivity on day t.[13]

Chart 8 below shows how this normalised reciprocity measure changed over the period. There was a rising trend through the first half of 2008 before a decline. In October 2008 it increased again. Since normalised reciprocity was always greater than 1 throughout the period under examination, we can surmise that reciprocity was greater than would be expected if the network were Erdős-Rényi.

[12] There is a theoretical issue of how we would define reciprocity in an empty network – that is, one with no links. Of course, the probability of an empty network tends to zero as n tends to infinity so this is less of a concern for larger networks. Over the sample period examined in this paper, the unsecured money market network is never empty.

[13] An explanation of this is as follows. Suppose we have an Erdős-Rényi random network, where each link is formed independently and with probability p. Then connectivity has expected value p. For a network with n nodes, it can be shown that the expected value of reciprocity tends toward p as n tends to infinity. Suppose we observe a decrease in reciprocity: how can we distinguish between the case where p has simply declined and the case where reciprocated links have become less likely (ie the structure of the network has changed)? An obvious solution is to 'normalise' reciprocity by dividing by p. In an Erdős-Rényi network, the expected value of this normalised measure should be equal to 1, and invariant to changes in p. Therefore significant changes in its value suggest that the structure of the network is different, and the probability of reciprocal links has changed. In practice we cannot observe p directly but connectivity is the likelihood-maximising estimate.

Chart 8: Normalised reciprocity of the network over the period

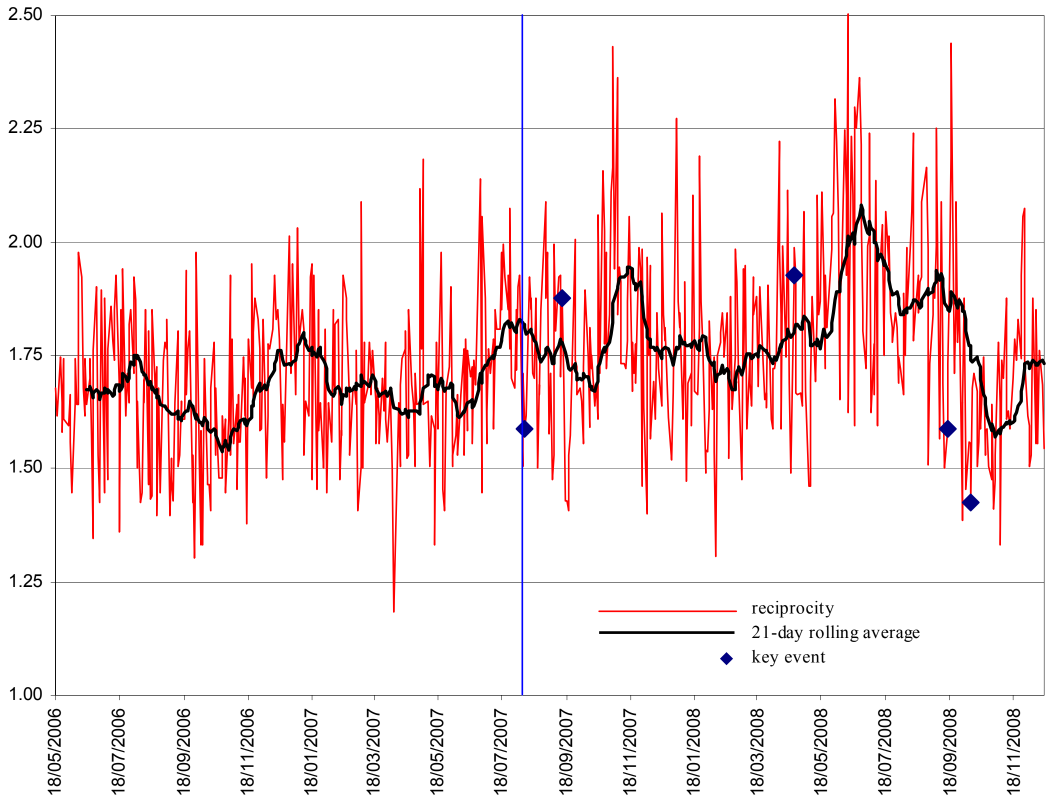

Chart 8 does not suggest any associations between normalised reciprocity and specific events (with the exception of the Lehman default), but Table 7 below confirms that normalised reciprocity was higher during the crisis than before.

Table 7: Mean normalised reciprocity during each phase

	Phase 1	Phase 2	t-stat	p-value
Mean	1.6805	1.7932	7.59	***0.0%
Variance	0.0282	0.0449		

3.13 Persistence of relations

We define persistence on day *t* as the probability that, given a link exists in the network on day *t*, the same link exists in the network on day *t*+1. As with reciprocity, we need to account for the effect of changes in connectivity: we therefore define normalised persistence on day *t* as persistence on day *t* divided by connectivity on day *t*+1.[14]

Chart 9 and Table 8 show that normalised persistence was higher during the crisis than before. And, since normalised persistence is always greater than 1, we can deduce that reciprocity was greater than would be expected if the network were Erdős-Rényi.

[14] Suppose we have Erdős-Rényi random networks on days *t* and *t*+1 which are independent of one another. Clearly, the expected value of persistence on day *t* will be equal to the probability of a link forming on day *t*+1, which can be estimated by connectivity on day *t*+1.

Chart 9: Normalised persistence of the network over the period

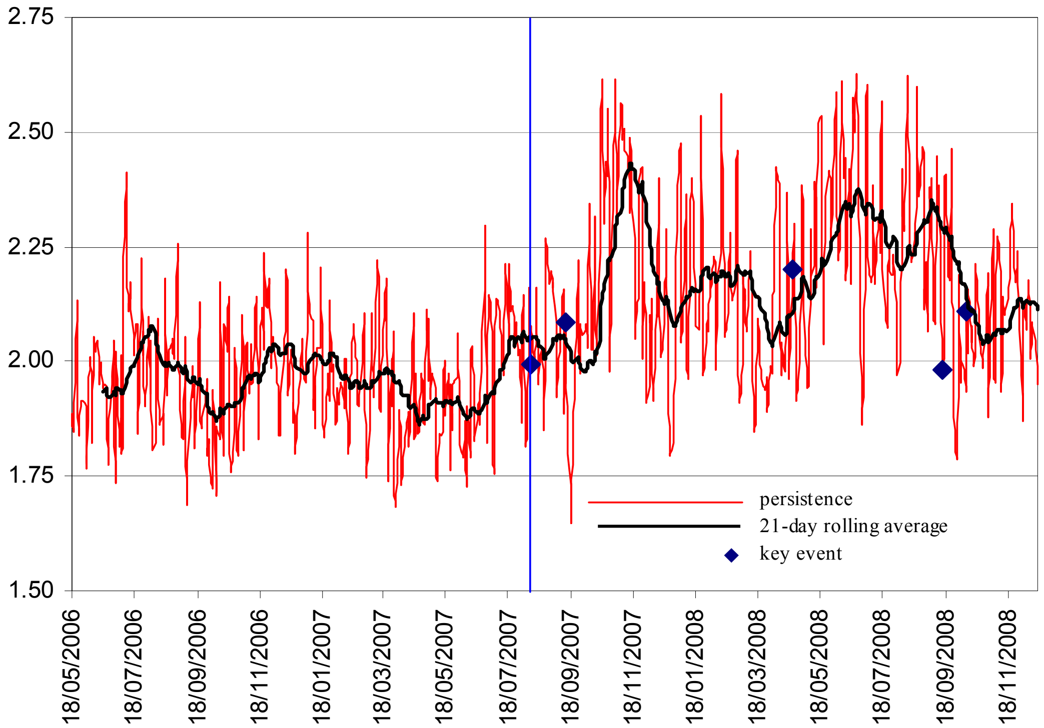

Table 8: Mean normalised persistence during each phase

	Phase 1	Phase 2	t-stat	p-value
Mean	1.9627	2.1792	17.51	***0.0%
Variance	0.0158	0.0352		

3.14 Summary

To conclude this section, we summarise the results from the network analysis, before relating them to the results of the theoretical banking literature and the main events of 2007-08 in Section 4 below. For convenience, Table 9 summarises the various network measures.

Table 9: Summary of results in Section 3

	Mean		
Measure	**Phase 1**	**Phase 2**	**p-value**
Size of core	4.18	4.67	***0.0%
p_{CC}	0.99	0.97	***0.0%
p_{CN}	0.63	0.55	***0.0%
p_{NC}	0.58	0.47	***0.0%
p_{NN}	0.06	0.05	***0.8%
p_{CC}/χ	2.37	2.54	***0.0%
p_{CN}/χ	1.50	1.43	***0.0%
p_{NC}/χ	1.37	1.23	***0.0%
p_{NN}/χ	0.14	0.13	35.5%

	Mean		
Measure	**Phase 1**	**Phase 2**	**p-value**
$Q_l^{(t)}$	0.58	0.60	***0.0%
$Q_b^{(t)}$	0.61	0.65	***0.0%
$Q_b^{(t)} / Q_l^{(t)}$	1.06	1.09	***0.5%
$R_l^{(t)}$	0.56	0.54	***0.0%
$R_b^{(t)}$	0.59	0.58	***0.0%
$R_b^{(t)} / R_l^{(t)}$	1.06	1.09	***0.1%
Connectivity	0.42	0.38	***0.0%
N. reciprocity	1.68	1.79	***0.0%
N. persistence	1.96	2.18	***0.0%

At first glance, Charts 3 and 4 may appear to contradict each other. While Chart 3 suggests that connectedness of banks decreased gradually, Chart 4 implies that the networks could be characterised by a core of 4-5 banks which were much more connected than the others. There is no contradiction. Chart 3 does not examine who was connected to whom; for example it does not tell us whether the less-connected banks linked to each other or to the more-connected banks. It cannot distinguish between a core-periphery network and one with a series of loosely connected clusters. However, Chart 4 tells us that there was a set of peripheral banks which tended not to connect to one another but only to the core. And Chart 5 shows that there was a substantial difference between the core and non-core banks in terms of the likelihood of forming links.

Chart 6 shows that the core tended to be fairly stable, in that there are four 'regular core' banks which were usually in the core, with the others only joining occasionally. This suggests that there was a small group of banks which were dominant in the overnight market, but that others may from time to time have become more linked to this core according to the circumstances. During the crisis phase we do indeed observe that formerly peripheral banks did join the core, but connectivity did not increase in general. It seems then that banks became more connected to the core banks and less linked to those in the periphery. Since the core was generally smaller than the non-core – Chart 4 implies that on most days more than half of banks were outside the core, even during the crisis – this meant links cut (to the non-core) exceeded new links established (to the core), and connectivity therefore fell.

Indeed, banks could even join the core as a result of changes in counterparties' behaviour, rather than their own. Consider two non-core banks A and B which linked to one another. Suppose bank A became more linked to the core, and thus joined the core itself. Bank B now had one more link to the core than before, and one less to the non-core. This may have been enough to move B into the core itself, even though it had not changed its relationships with counterparties at all.

The measures of dependence show that banks became more dependent on the core during the crisis phase, but less dependent on the 'regular core' (that is, the four banks most often in the core). This is consistent with our story: as the core became larger, the influence of the 'regular core' banks within it became less important. A bank could have joined the core simply because some of its hitherto non-core counterparties joined themselves; in principle it did not need to link to the 'regular core' at all. We see that banks formed more links with new entrants to the core, and fewer with both the non-core banks and the traditional 'regular core'. In other words, there were intermediary banks which became more important.

We also see that borrowers linked to the core more frequently than lenders. This is apparent from the measures of dependence, and also from the fact that p_{CN} is larger than p_{NC}, both in absolute terms and relative to overall connectivity. And the tables show that this asymmetry increased significantly during the crisis phase.

Normalised reciprocity was significantly higher during the crisis, which is consistent with the explanation above. We have established that during the crisis phase a greater proportion of links involved the core, and that the core was generally smaller than the non-core. This means that links became more concentrated among a relatively narrow range of counterparties. Therefore there was a greater chance of lending to and borrowing from the same counterparty. A similar argument explains the increase in normalised persistence. As counterparties become more concentrated during the crisis phase, banks were more likely to form links with the same counterparties on consecutive days.

In summary, we find that:

- Banks that were previously peripheral join the core during the crisis phase, but this could be due to only a few intermediary banks changing their behaviour.
- Since there was an overall decline in connectivity, there were more links cut to the non-core than there were new links established to the core.
- Borrowers were more likely to link to the core than lenders.

4 Analysing the network measures

The results discussed so far are consistent with many of the classic features of disintermediation explained by the theoretical banking literature. First, theoretical models tell us that interbank markets may cease to function efficiently when concerns about credit worthiness increase and banks are hit by aggregate liquidity shocks. The result is an overall reduction in interbank activity, often accompanied by a reallocation of flows away from weaker banks.

Freixas and Jorge (2008) model the impact of an aggregate liquidity shock and show how severe liquidity shortages may arise. Together with asymmetric information, this causes liquidity in the interbank market to flow towards the most liquid banks, at the expense of the less liquid ones. Likewise, Acharya *et al* (2008) show that during a liquidity crisis, liquidity-poor banks will find access to the interbank market greatly reduced, as liquidity-rich banks exert their market power and charge higher rates. This forces the former to exit the interbank market and rely on asset sales instead.

There is a second mechanism at work. The banking literature also demonstrates that, when faced with unexpected shocks, banks may need to build up their own liquidity reserves. Freixas *et al* (2000) consider a situation where lenders withdraw from the market because they are uncertain about their own ability to borrow in the future. Allen *et al* (2009) show that banks reduce their interbank lending when there is uncertainty about the overall demand for liquidity in the banking system. In this model too, banks cease to use the interbank market and start hoarding liquidity. Caballero and Krishnamurthy (2008) demonstrate that when Knightean uncertainty (ie uncertainty about future states of the world) increases, financial intermediaries are inclined to assume the worst-case scenario and hoard liquidity. Hence, this second line of research points to an overall reduction of lending activity following a rise in uncertainty, whether that uncertainty relates to banks' counterparties, to their own funding needs or to the general economic outlook.

In a recent paper, Ashcraft *et al* (2009) challenge this view and argue that faced with increased uncertainty about intraday liquidity shocks, banks may hoard liquidity in the early part of the day. Later in the day, as their payment obligations become clearer, they are more willing to lend their excess reserves overnight. This would explain both rises in overnight lending activity and large intraday variations in overnight rates. Furthermore, they show how smaller banks, who typically face greater liquidity constraints, build up larger intraday reserves and tend to be net lenders to the larger banks.

Empirical studies of interbank markets generally support this paper's conclusions, but add some interesting insights. Furfine (2001a) finds no evidence that interbank activity (in the federal funds market) declined following the 1998 events (Russian default, near-collapse of LTCM). Instead, he shows that, apart from the days surrounding the LTCM rescue, volumes in the second half of 1998 were higher than in the first half. He attributes this to increased dealer activity. At the same time, he finds some evidence of reduced borrowing by the most active and risky banks, which is consistent with the conclusions of Freixas and Jorge (2008) and Acharya *et al* (2008).

A third group of papers, relevant for our study, look at the importance of relationships in interbank markets. In a seminal paper, Rochet and Tirole (1996) model the monitoring role of banks in these markets. They show that banks have strong incentives to monitor each other when interbank loans are large and unsecured. But these incentives can be undermined if banks believe that large financial institutions would never be allowed to fail. Furfine (2001b) confirms this risk monitoring role showing that US federal funds rates do indeed differentiate between banks in ways which plausibly reflect counterparty credit risk. At the same time, he finds that access to this market can rapidly dry up, partly as a result of banks' reluctance to signal to the market that they need funds, and partly because other banks wish to limit their credit risk exposure. In other words, Furfine finds evidence of credit rationing rather than an increase in the rates charged to individual banks when their condition deteriorates. Cocco *et al* (2003) highlight the importance of relationships in the interbank market in providing banks with insurance against liquidity risk, in the form of both unexpected shortages and surpluses of funds. Using data for the Portuguese interbank market, they find evidence of riskier borrowers relying on established relationships. Furthermore, they find that during the 1998 crisis, when overall liquidity fell, borrowers relied more than usual on banks with which they had an existing lending relationship.

In summary, theory suggests that when banks become more concerned about funding liquidity risk across the market, they may attempt to reduce risk exposure either by reducing their lending activity or by changing their borrowing and lending relationships. Our paper finds evidence of the latter. Indeed, the data suggest that banks relied less on established relationships and instead formed relationships with medium-sized intermediaries, which then themselves joined the core. This may be because borrowers were concerned about specific names, or more generally that they wished to diversify funding sources and have more rather than fewer relationships.

Alternatively, it may be the lenders who drove this change as revised credit criteria may have limited the amount and terms of funding that they were prepared to provide.[15] It is difficult to assess whether the changes were due to lender or borrower preferences, or a combination of the two. Since the total amount of overnight activity increased during the crisis phase, a right-shift in the demand curve could be considered more likely rather than a left-shift in supply. But there

[15] See Bank of England (2008a), Chart 1.5.

may be other factors to consider here: for example supply may have fallen by even more in the term markets, leading to borrowers to substitute term funding with overnight money.

Furthermore, our results suggest an asymmetry between borrowers and lenders. This means that provision of funding was concentrated at the more-connected banks, and this concentration increased during the crisis phase. Since the core is small compared to the periphery, this implies that providers of funding may have had some scarcity power, which suggests that the change in relationships during the crisis was caused by a shift in supply rather than demand.

There is another explanation of this asymmetry. If a lender has an adverse liquidity shock, it can choose to cut down on the amount of funding rolled in overnight markets and use the money instead to improve its reserve account balance. The lender has discretion (unless, of course, it is a committed line). But a borrower has less discretion: refusal to pay triggers a credit event. Therefore borrowers – which seek lending counterparties – may be more concerned about liquidity risk at their counterparties than lenders are. If we believe that borrowers had some market power, during the crisis phase they may have sought out the more secure counterparties and relied more on the core than lenders did. And we might expect that this effect was magnified as term markets closed and borrowers were forced to roll overnight funding instead. With term markets resuming to a limited extent toward the end of the phase, we might expect this effect to be less pronounced, but not enough to counterbalance the impact of the worst period of the crisis.[16]

Changes in the monetary framework may also have had an effect. One reason for banks to participate in the overnight market was to manage their reserve account balances close to the target, as described in Subsection 2.1. Although the target only related to average end-of-day balances over the maintenance period, banks are known to actively manage their positions and set themselves targets on a daily basis. If a bank is long or short near the end of the day, it may choose to lend or borrow overnight in the market rather than miss its internal target.

Since this reserve management activity behaviour is largely discretionary, we might expect it to be strongly affected by the heightened counterparty concern during the crisis. But the bands remained relatively tight at ±1% until September 2007, so banks did not have much room to manoeuvre. To the extent that counterparty risk limits were a constraint, they would have had to diversify this business across counterparties. Once the bands were widened, banks had much more discretion about whether to participate in the overnight market, and so could have been more selective about their relationships.

It seems then that both market and policy events had an impact on the network. Although we do not attempt to measure the separate impact of these, we argue that combined they would have led to changes to the core during the crisis, by affecting the actions of both lenders and borrowers.

[16] Bank of England (2008a), page 34.

5 A closer examination of sub-phases of the market turbulence

The crisis is widely recognised to have worsened considerably following the bankruptcy of Lehman Brothers on 15 September 2008. We now take a closer look at the two sub-phases defined by the Lehman default and the announcement of a bank recapitalisation scheme by the UK government on 8 October. Do these sub-phases appear different to the rest of the crisis phase?

5.1 Comparing the measures

We define the following divisions of phase 2: sub-phase 2a from 9 August 2007 to 12 September 2008; sub-phase 2b from 15 September to 7 October 2008; and sub-phase 2c from 8 October to 16 December 2008. It should be noted that, as sub-phase 2b is only 17 days long, the underlying distributions of the means of our statistics are less likely to be normal so the accuracy of the Welch test may be compromised. Furthermore, results are less likely to be significant due to the short sample period.

Table 10 shows that the core continued to increase in size during sub-phases 2b and 2c, and the probability of forming links to the core fell. The probability of forming core-core links fell relative to connectivity, to below pre-crisis levels after the recapitalisation plan was announced. Lenders used the regular core less in sub-phase 2c, but not borrowers. These indicate that the changes to the core became more significant as the crisis phase went on. Despite this, connectivity returned to pre-crisis levels in sub-phase 2c, and normalised reciprocity and persistence were a little lower.

Table 10: Comparison of the post-Lehman sub-phases

	2a mean	2b mean	2c mean	p-value 2b vs 2a	p-value 2c vs 2a
Size of core	4.53	5.06	5.34	***0.5%	***0.0%
p_{CC}/χ	2.60	2.48	2.22	*5.4%	***0.0%
p_{CN}/χ	1.47	1.28	1.22	***0.0%	***0.0%
p_{NC}/χ	1.29	1.00	0.93	***0.3%	***0.0%
p_{NN}/χ	0.13	0.11	0.13	30.4%	95.1%
$R_l^{(t)}$	0.54	0.53	0.52	42.0%	***0.0%
$R_b^{(t)}$	0.58	0.60	0.58	26.7%	77.6%
connectivity	0.38	0.40	0.42	*9.0%	***0.0%
normalised reciprocity	1.80	1.75	1.64	51.9%	***0.0%
normalised persistence	2.20	2.11	2.09	*8.5%	***0.0%

There are mixed findings here. It appears that, as the core continued to grow in size during sub-phases 2b and 2c, the probability of core banks connecting fell. And dependence on the regular core continued to drop slightly too. But overall banks became more connected, and normalised reciprocity and persistence fell, reversing the trend seen earlier in the crisis phase.

It seems then that we see a different dynamic in the later sub-phases of the crisis. The core became 'diluted', as less-connected banks joined it. Lenders continued to move away from the regular core but borrowers did not. This could be because these banks – which had hitherto been viewed as safe, reliable counterparties – began to be seen as risky credit and lenders preferred to move their money elsewhere. The regular core banks responded to this by lending money out, trying to send a message to the money markets that they were not desperate for funds. This is consistent with the drop in normalised reciprocity.

It may be surprising perhaps that these effects should have been more significant in 2c than 2b, which is generally viewed as the worst period of the turmoil. This may be partly due to the smaller sample for 2b, but it could also be because counterparties needed time to adjust their activity. Lenders might have preferred to bring rolled overnight funding to the end of its term naturally rather than terminate it suddenly, for fear that the counterparty could view this negatively. Therefore there may be a lag here, with the shock in sub-phase 2b having its full impact in 2c. The decline in normalised persistence throughout 2b and 2c suggests that rolled funding may have been moved gradually and not at the start of sub-phase 2b.

5.2 Changes in the size of the network

Table 7 and Charts 7, 8 and 9 show that connectivity and normalised reciprocity and persistence changed the most from pre-crisis levels between Autumn 2007 and Autumn 2008. These phases coincide with changes in the underlying membership of CHAPS. One bank became a direct member in October 2007 and another ceased its direct membership in September 2008. It is possible that these changes have had an impact on the structure of the network.

We could attempt to control for these changes by considering a scenario where neither of the 'floating' banks was ever a direct CHAPS member, and instead all of their overnight activity was done through their former/later correspondent bank during the entire period. This would give us a network with 11 nodes. In this network we find that there is no longer a significant drop in connectivity in Autumn 2007 or a rise in Autumn 2008.

However, there are limitations to this analysis. We do not know that banks would have kept their lending and borrowing behaviour the same if the floating banks had been second-tier members throughout the entire sample period. For example, for some customers it may be operationally easier to use the correspondent for both intraday and overnight credit. For others, they may prefer to choose a different counterparty for overnight funding so as not to impair the intraday credit line granted by the correspondent.

In summary, although there is evidence that changes in CHAPS membership may have contributed to our results, we cannot conclude that they alone explain the observed changes in the network. We do not know what the counterfactual situation would have been.

6 Concluding remarks

Our network results indicate the existence of a core of a small number of banks which accounted for a large portion of overnight relationships. When concerns about counterparty risk increased, banks in the network preferred to diversify their relationships rather than rely on the core, leading to some previously non-core banks becoming connected enough to join the core themselves. Often these banks acted as intermediaries: by joining the core, their counterparties became more connected to the core too and thus joined themselves.

This may have been due to changes in borrower or lender preferences, or a combination of the two. Lenders may have cut counterparty limits, forcing borrowers to acquire funding from a wider range of counterparties. And borrowers may have sought to diversify funding sources as concerns grew about the availability of liquidity from established counterparties. Changes to the monetary framework may have had an effect too, as wider reserve target bands allowed banks to exercise more discretion about when to enter the overnight market.

This asymmetry between borrowers and lenders – and its increase during the crisis – may be a result of reliance on the core for rolled overnight funding. As the availability of term liquidity dried up, borrowers obtained more funding from the overnight market and rolled it on a daily basis. This may have made borrowers more concerned about counterparty risk, since this funding could be cut off at any time, and so they chose to borrow from the core. Alternatively, it may be that these banks – or their clients – were the only ones willing to provide overnight funding.

The network indicators do not improve after the recapitalisation plan was announced, though this may be because of a lag between events occurring and being able to find new relationships. But the short sample size of the deepest part of the crisis makes drawing strong conclusions rather difficult.

We do not attempt to measure whether the impact of market events was greater or less than the impact of policy events. This question could be important when attempting to gauge the effect of central bank actions. For example, increased access to central bank liquidity as a result of policy changes during the crisis may have reduced the need for private provision, affecting the relationships that commercial banks have with each other.

There are some limitations to our analysis. Changes in the underlying membership of the system mean that it is difficult to compare the network over time. A similar, but unobserved, factor is the activity of second-tier banks: as their sterling funding and lending needs change – or they change the settlement bank used – there could be a confounding effect on our analysis. Many of these second-tier banks do not have reserve accounts, so are not subject to the same end-of-day constraints as settlement banks. And the regular core members may be significant providers of correspondent services for these second-tier banks, which may affect their observed behaviour.

The notion of intermediaries drawing new members into the core could be explored by expanding the framework. Instead of looking at partitions of the nodes into two sets, we could look at partitions into three: the core, an intermediate level, and the periphery. The theoretical model would be similar, but the limiting factor here is the amount of computation power required. To look at all possible partitions for a network of thirteen nodes currently takes 8,192 calculations per day examined; with this extension it would take nearly 200 times as many. This may be more feasible with an optimising algorithm. We leave this for future research.

In conclusion, our network analysis of the overnight money market indicates that the structure of relationships between banks changed as the crisis unfolded. But the analysis also suggests that the observed changes were small, and that underlying trading activity was largely unaffected. More work is needed to understand how activity in the overnight unsecured market was affected by changes in the term markets and in the secured markets. Further research could provide indications of the resilience of liquidity in the money markets and establish whether changes in the underlying infrastructure could be beneficial.

Annex 1: Finding overnight loans from payment data

The Furfine method for finding overnight loans employed in this paper can be summarised as follows:

1. Generally loan principal amounts are in fairly round numbers. On day *t*, find all payments in round numbers and label them as possible overnight loan advances. For overnight loans, we define a 'round number' as being of value £1 million or above and divisible by £100,000. Thus we only consider loans of value £1 million or more.
2. On day *t*+1 label all payments of value £1 million or above and not a 'round number' as possible overnight loan repayments. For each of these, calculate the implied principal amount by rounding down to the nearest £100,000, or round down to the nearest £1 million if the repayment amount is greater than £250 million.[17]
3. Match possible advances on day *t* with implied principal amounts from repayments on day *t*+1. For each potential matched pair, check:
   - The advance has not already been matched with another repayment, and *vice versa*;
   - The payer of the advance matches the payee of the repayment, and *vice versa*;
   - The implied interest rate falls within ±2% of Bank Rate;
   - The implied interest rate is plausible, meaning that is an exact number of basis points or half-points. To allow for rounding errors, we accept interest rates within 0.01 basis points of an exact number.

This algorithm is relatively accurate for overnight loans. It is unlikely to be 100% accurate since it could happen that opposing payments on consecutive days resemble a loan advance and repayment without actually being part of loan. Also there is a problem of ambiguity if there are payments on three or more consecutive days that could be advances and repayments (in this case the algorithm assumes an advance on the first day and repayment on the second). But the likelihood of these coincidences is low enough to justify use of the method. We also lose any overnight loans that are less than £1 million in value, but again this should not have a major impact. Millard and Polenghi (2004) conduct robustness checks on this data set and confirm that the data appear representative of the sterling unsecured overnight market.

In the case of the UK market, we exclude any matches involving the Bank of England or CLS. Since these institutions do not participate in the unsecured overnight market, any matches must be erroneous. We also treat RBS and NatWest as a single entity. Since these institutions are part of the same banking group, counterparties will have identical perceptions of them. This leaves us with twelve or thirteen CHAPS settlement banks (depending on the date) participating in the unsecured overnight sterling market over our sample period.

[17] There may appear to be a danger here of making an error if the repayment amount is slightly more than a multiple of £100,000 or £1 million. But if a loan is of size less than £250 million and the interest is greater than £100,000, then the implied annualised interest rate is over 14%, which is unreasonable considering prevailing interbank rates in the United Kingdom during the period studied. And if a loan commands interest of more than £1 million, then even at a rate of 8% (very high compared to three-month Libor over the period), the principal amount would be over £4.5 billion. Payments of this magnitude are very rare, and there have been none on consecutive days between the same pairs of banks over the period we are examining.

Annex 2: Algebraic definition of the core

Formulation

Given n nodes, we label them $\{1,2,\ldots,n\}$. The topology measures are invariant under permutations of this labelling. If there is a link from node i to node j, then $l_{ij}=1$. If there is no such link, then $l_{ij} = 0$. As no node can have a link to itself, $l_{ii}=0$ for all i. Then L is an nxn matrix with 0s along the leading diagonal.

Now consider a partition of the n banks into two sets, the core (denoted by C) and the non-core (denoted by N). We can describe this partition by an n-vector v such that $v_i = 1$ if $i \in \mathrm{C}$ and $v_i = 0$ otherwise. This partition of the nodes implies a partition of the links into four separate sets, determined by whether they link to/from the core/non-core. We use CC to denote the set of links connecting core banks to other core banks, CN the links from core to non-core banks, NC from non-core to core banks, and NN non-core to non-core banks.

In order to pick an optimal core, we need to make an assumption about the structure of the network. We assume that links in each of the four sets are formed independently and with equal probability. In other words, our network can be assumed to be four Erdős-Rényi networks overlaid on each other. Denote p_{CC} as the probability that one core banks lends to another, p_{CN} the probability that a core bank lends to a non-core bank, and similarly for p_{NC} and p_{NN}.

Let u be the n-vector consisting entirely of 1s. Then $m = u^{\mathrm{T}}v$ is the size of the core, and $s_{CC} = m(m\text{-}1)$ is the number of possible core-core links. Similarly, we can define $s_{CN} = m(n\text{-}m) = s_{NC}$ and $s_{NN} = (n\text{-}m)(n\text{-}m\text{-}1)$. We assert $s_{CC} = 0$ if $m \leq 1$, $s_{CN} = s_{NC} = 0$ if $m=0$ or $m=n$, and $s_{NN} = 0$ if $m \geq n\text{-}1$.

Define $\lambda_{CC} = v^{\mathrm{T}}Lv$, the number of core-core links which exist in network L.
Similarly, $\lambda_{CN} = v^{\mathrm{T}}L(u\text{-}v)$, $\lambda_{NC} = (u\text{-}v)^{\mathrm{T}}Lv$ and $\lambda_{NN} = (u\text{-}v)^{\mathrm{T}}L(u\text{-}v)$.

The likelihood of this particular core allocation v is thus:
$$\mathcal{L}(v) = p_{CC}^{\lambda_{CC}}\left(1-p_{CC}\right)^{s_{CC}-\lambda_{CC}}\, p_{CN}^{\lambda_{CN}}\left(1-p_{CN}\right)^{s_{CN}-\lambda_{CN}}\, p_{NC}^{\lambda_{NC}}\left(1-p_{NC}\right)^{s_{NC}-\lambda_{NC}}\, p_{NN}^{\lambda_{NN}}\left(1-p_{NN}\right)^{s_{NN}-\lambda_{NN}},$$
where we evaluate $0^0 = 1$.

The value of p_{CC} which maximises the likelihood is $\frac{\lambda_{CC}}{s_{CC}}$, and similarly for p_{CN}, p_{NC} and p_{NN}.

Having found these in terms of the other unknowns, which are themselves functions of L and v, we need only find the likelihood-maximising core allocation v^*.

We examine all possible $v \in \{0,1\}^n$ to find that which maximises the likelihood.[18] We then have the optimal core allocation, along with the associated probabilities.

[18] Unfortunately, this is not a problem that lends itself to a greedy algorithm. Consider an algorithm which finds the best v for m=1, then finds the best node to add to the core to give an optimum for m=2, and so on. We cannot be sure that the member of the singleton core would be a member of a core with two elements. There are ways to test this, but given the small value of n in our network it is not too computationally expensive to consider all possible members of $\{0,1\}^n$.

Clearly, $\mathcal{L}(v) = \mathcal{L}(u\text{-}v)$ for any $v \in \{0,1\}^n$ – that is, we could switch the core and non-core nodes without changing the likelihood, since this would represent the same partition. We choose between two likelihood-maximising allocations by imposing $p_{CC} \geq p_{NN}$ – in other words, we choose the allocation which gives a higher probability of links between core banks than links between non-core banks.

Since we have a different L for each day, we produce a different optimal core each day.[19] We can thus examine any changes in the membership of the core, and of the probabilities of forming links.

On each day, we find a unique likelihood-maximising core allocation. Furthermore, we find that the following inequalities are true on each of the 656 days:

$$p_{CC} > p_{CN}\,, \qquad p_{CC} > p_{NC}\,, \qquad p_{CN} \geq p_{NN}\,, \qquad p_{NC} \geq p_{NN}$$

with the last two holding with equality only when the probabilities are zero. This is a robustness check: since banks should form links with the core with higher probabilities than they do with the non-core, we would expect these four inequalities to hold. There is nothing in the way the model is formulated to require them to be true, but they imply that the model is self-consistent.

Likelihood ratio testing for existence of the core

To test for significance of the core, we form a null hypothesis that the partition is trivial and that all links have the same probability of being formed. We write our null hypothesis $v = u$ and alternative hypothesis $v = v^*$, where v^* is the n-vector which maximises the likelihood $\mathcal{L}(v)$. Define the ratio $\Lambda = \mathcal{L}(u)/\mathcal{L}(v^*)$. Then the test statistic $-2\log(\Lambda)$ is approximately chi-squared with n degrees of freedom.

The approximation improves as $n \rightarrow \infty$, though it is usually considered adequate for smaller values. As n is no more than 13 in our network, we can compensate for a possible poor fit to the distribution by increasing the significance level of the test. As it happens, the null hypothesis is rejected every day, even at extremely low significance levels. Prior to the Lehman episode, the null hypothesis is rejected every day with a level of 0.0006%. The required levels are considerably higher after Lehman; this is because more banks join the core (as explained above) and the structure can be better specified by a single tier. However, the change is relative; the actual p-values remain very low, peaking at 0.0024%.

Possible extensions

There are several ways in which this work could be extended.

<u>Probabilistic core:</u> Instead of calculating a core each day, we could take an average network over each phase – in other words, two matrices where the $(i,j)^{\text{th}}$ entry is the average of all l_{ij} over the phase. Since these matrices would have entries in the interval [0,1] instead of merely {0,1}, we may choose our core from members of $[0,1]^n$. Thus our vector v would tell us the probability of each node being in the core in each phase, and would allow direct comparison of the two

[19] For some L we have n=12 and others n=13.

phases. Preliminary work on this suggests that there are very many local maxima, meaning that a sophisticated algorithm is needed to find the globally optimal core allocation.

Twin core: Another possible extension is to have two cores: one of the most connected borrowers and another of the most connected lenders. This would enable us to determine whether the important borrowers differ from the important lenders, and if there is any change during the crisis. If we call the in-core v and the out-core w then we can simply re-define $\lambda_{CC} = v^T Lw$, $\lambda_{CN} = v^T L(u\text{-}w)$ etc and similarly for s_{CC}, s_{CN} etc. The formulation of the problem does not change much but, as the solution now lies in $\{0,1\}^{2n}$ the computation required is much greater.

Multi-layered core: As discussed in our conclusion, we could consider a model where we partition the nodes into three (or, more generally, any $k \leq n$) different sets rather than the two considered here. Again, the formulation would not change much but the size of the set of possible core allocations would be k^n instead of 2^n, making for a more computationally expensive problem.

Weighted network: Finally, this model of the core can be easily extended to a weighted network, using for example the values or volumes of daily overnight activity between pairs of banks. An assumption needs to be made as to the distribution of the values of the links: in the unweighted version, the Bernoulli distribution was the obvious choice. The shape of the distribution can be determined empirically from observing the data: for instance in our network the values between banks appear to be gamma distributed. The likelihood function can then be written down and the parameters estimated in the usual way.

Annex 3: Algebraic definitions of topology measures

Let i and j denote nodes in the network. Then $l_{ij}^{(t)}$ equals 1 if a link exists between banks i and j on day t, and 0 otherwise.

Let C denote the core banks. The degree to which lenders and borrowers depend on the core are defined as $Q_l^{(t)} = \frac{\sum_i \sum_{j \in C} l_{ij}^{(t)}}{\sum_i \sum_j l_{ij}^{(t)}}$ and $Q_b^{(t)} = \frac{\sum_{i \in C} \sum_j l_{ij}^{(t)}}{\sum_i \sum_j l_{ij}^{(t)}}$ respectively.

Similarly, let I denote the set of the four regular core banks. Then the degree to which lenders and borrowers depend on the regular core are $R_l^{(t)} = \frac{\sum_i \sum_{j \in I} l_{ij}^{(t)}}{\sum_i \sum_j l_{ij}^{(t)}}$ and $R_b^{(t)} = \frac{\sum_{i \in I} \sum_j l_{ij}^{(t)}}{\sum_i \sum_j l_{ij}^{(t)}}$ respectively.

Connectivity $\chi^{(t)} = \frac{\sum_i \sum_j l_{ij}^{(t)}}{n(n-1)}$.

Reciprocity $r^{(t)} = \frac{\sum_i \sum_j l_{ij}^{(t)} l_{ji}^{(t)}}{\sum_i \sum_j l_{ij}^{(t)}}$.

Normalised reciprocity $\hat{r}^{(t)} = \frac{r^{(t)}}{\chi^{(t)}}$.

Persistence $\pi^{(t)} = \frac{\sum_i \sum_j l_{ij}^{(t)} l_{ij}^{(t+1)}}{\sum_i \sum_j l_{ij}^{(t)}}$.

Normalised persistence $\hat{\pi}^{(t)} = \frac{\pi^{(t)}}{\chi^{(t+1)}}$.

Annex 4: Major events during the 2007-08 UK market turmoil

Date	Event
July 07	First real signs of a crisis. Bear Stearns announces that two sub-prime hedge funds have rapidly declined in value.
9 Aug 07	Generally accepted start of crisis. ECB and Federal Reserve inject around £45 billion of funds into the financial markets.
20 Aug 07	First use of standing facilities during the crisis; clear signs that borrowing on facilities carries stigma.
13 Sep 07	Bank of England announces a widening of reserve bands for the current maintenance period.
14 Sep 07	Northern Rock granted liquidity support facility from Bank of England.
8 Oct 07	UBS joins CHAPS as a direct member.
17 Mar 08	JP Morgan Chase agrees to buy Bear Stearns.
21 Apr 08	Bank of England launches Special Liquidity Scheme, which allows commercial banks to borrow Treasury bills in exchange for less liquid collateral.
15 Sep 08	Bankruptcy of Lehman Brothers.
17 Sep 08	Disclosure of merger talks between HBOS and Lloyds TSB.
22 Sep 08	ABN Amro ceases its direct membership of CHAPS.
29 Sep 08	Bradford & Bingley, a UK mortgage bank, part-nationalised.
8 Oct 08	Announcement of UK bank recapitalisation plan.
20 Oct 08	Reform of standing facilities.

Annex 5: Range around reserve targets within which reserves are remunerated at Bank Rate

Announcement	Band	Effective	Announcement	Band	Effective
18 May 06	±1%	18 May 06	18 Sep 08	±40%	4 Sep 08
13 Sep 07	±37.5%	6 Sep 07	1 Oct 08	±60%	4 Sep 08
20 Sep 07	±60%	6 Sep 07	6 Oct 08	±40%	9 Oct 08
2 Oct 07	±30%	4 Oct 07	3 Nov 08	±20%	6 Nov 08
7 Jul 08	±20%	10 Jul 08	1 Dec 08	±10%	4 Dec 08

Note that some of the band changes were applied retrospectively – in these cases the 'effective' date precedes the announcement.

References

Acharya, V V, Gromb, D and Yorulmazer, T (2008), 'Imperfect competition in the interbank market for liquidity as a rationale for central banking', *CEPR Discussion Paper No. 6984*.

Albert, R and Barabási, A L (2002), 'Statistical mechanics of complex networks', *Review of Modern Physics*, Vol. 74, pages 47-97.

Allen, F, Carletti, E and Gale, D (2009), 'Interbank market liquidity and central bank intervention', *Journal of Monetary Economics,* Vol. 56(5), pages 639-52.

Ashcraft, A B, McAndrews, J and Skeie, D R (2009), 'Precautionary reserves and the interbank market', *Federal Reserve Bank of New York Staff Report No. 370*.

Bank of England (2005), *Quarterly Bulletin,* Vol. 45, No. 2.

Bank of England (2008a), *Financial Stability Report*, October.

Bank of England (2008b), *Quarterly Bulletin*, Vol. 48, No. 3.

Bank of England (2008c), *Quarterly Bulletin*, Vol. 48, No. 4.

Bank of England (2008d), *The Development of the Bank of England's Market Operations*.

Bank of England (2009), *Market Notice: Sterling Monetary Framework; Asset Purchases*.

Bech, M L and Atalay, E (2008), 'The topology of the federal funds market', *Federal Reserve Bank of New York Staff Report No. 354*.

Bech, M L and Rørdam, K B (2008), 'The topology of Danish interbank money flows', *Danmark Nationalbank Working Paper No. 59*.

Becher, C, Millard, S and Soramäki, K (2008), 'The network topology of CHAPS Sterling', *Bank of England Working Paper No. 355*.

Borio, C (2008), 'The financial turmoil of 2007 - ?: a preliminary assessment and some policy considerations', *BIS Working Paper No. 251*, March.

Boss, M, Elsinger, H, Summer M and Thurner S (2004), 'The network topology of the interbank market', *Quantitative Finance*, Vol. 4(6), pages 677-84.

Caballero, R J and Krishnamurthy, A (2008), 'Collective risk management in a flight to quality episode', *Journal of Finance,* Vol. 63(5).

Cocco, J F, Gomes, F J and Martins, N C (2003), 'Lending relationships in the interbank market', *LBS IFA Working Paper No. 384*.

Committee on the Global Financial System (2008), 'Central bank operations in response to the financial turmoil', *CGFS Publications No. 31*.

Čopič, J, Jackson, M O and Kirman, A (2009), 'Identifying community structures from network data via maximum likelihood methods', *The B.E. Press Journal of Theoretical Economics,* Vol. 9(1), Article 30.

Dorogovtsev, S N and Mendes, J F F (2003), 'Evolution of networks', *Oxford University Press*.

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

Freixas, X and Jorge, J (2008), 'The role of interbank markets in monetary policy: a model with rationing', *UPF Working Paper No. 1027*.

Furfine, C (1999), 'The microstructure of the federal funds market', *Financial Markets, Institutions and Instruments,* Vol. 8, No. 5, pages 24-44.

Furfine, C (2001a), 'The interbank market during a crisis,' *BIS Working Paper No. 99*.

Furfine, C (2001b), 'Banks monitoring banks: evidence from the overnight federal funds market,' *The Journal of Business*, Vol. 74, No. 1, pages 33-58.

Haldane, A G (2009), 'Rethinking the financial network', speech to Financial Student Association, Amsterdam.

Inaoka, H, Ninomiya, T, Shimizu, T, Takayasu, H and Taniguchi, K (2004), 'Fractal network derived from banking transaction – an analysis of network structures formed by financial institutions', *Bank of Japan Working Paper No. 04-E-04*.

Iori, G, de Masi, G, Precup, O V, Gabbi, G and Caldarelli, G (2008), 'A network analysis of the Italian overnight money market', *Journal of Economic Dynamics and Control*, Vol. 32(1), pages 259-78.

Lublóy, A (2006), 'Topology of the Hungarian large-value transfer system', *Magyar Nemzeti Bank Occasional Paper No. 57*.

Millard, S and Polenghi, M (2004), 'The relationship between the overnight interbank unsecured loan market and the CHAPS Sterling system', *Bank of England Quarterly Bulletin*, Vol. 44, No. 1, pages 42-47.

Newman, M E J (2003), 'The structure and function of complex networks', *SIAM Review*, Vol. 45, pages 167-256.

Pröpper, M, van Lelyveld, I and Heijmans, R (2008), 'Towards a network description of interbank payment flows', *De Nederlandsche Bank Working Paper No. 177*.

Rochet, J-C and Tirole, J (1996), 'Interbank lending and systemic risk,' *Journal of Money, Credit and Banking,* Vol. 28, No. 4, pages 733-62.

Soramäki, K, Bech, M L, Arnold, J, Glass, R J and Beyeler, W E (2007), 'The topology of interbank payment flows', *Physica A*, Vol. 379, pages 317-33.