

Liquidity Hoarding, Network Externalities, and Interbank Market Collapse

Prasanna Gai* and Sujit Kapadia†

September 2009

Preliminary – Not for Further Circulation

Abstract

The interbank market collapse has been a central feature of the global financial crisis. In response to concerns about the viability of some off-balance sheet vehicles, banks stopped lending to each other at all but the shortest maturities resulting in a funding drought across the entire global financial system during 2007 and 2008. Taking a network approach, our paper argues that such systemic breakdowns of the interbank market are rare events that can be explained by precautionary hoarding amongst some banks due to concerns about their own future liquidity needs, and endogenous responses by banks to the liquidity hoarding of others. The failure of banks to internalise the potential consequences of their actions for the system as a whole points to a role for prudential policy. Our findings highlight the potential contribution that tougher liquidity requirements for systemically important institutions can make in averting contagion through funding markets.

Keywords: Interbank markets; Network models; Contagion; Systemic risk; Financial crises; Liquidity risk

JEL classification: D85; G01; G21

Disclaimer: This paper represents the views of the authors and should not be thought to represent those of the Bank of England or Monetary Policy Committee members.

*Crawford School of Economics and Government, Australian National University, Canberra, ACT 0200, Australia. *E-mail address:* prasanna.gai@anu.edu.au.

†Bank of England, Threadneedle Street, London, EC2R 8AH, United Kingdom. *E-mail address:* sujit.kapadia@bankofengland.co.uk

1. Introduction

The collapse of lending in interbank markets has been a central feature of the recent global financial crisis. As Figure 1 shows, spreads between interbank rates for term lending and expected policy rates in the major funding markets rose sharply in August 2007, before spiking in September 2008 following the collapse of Lehman Brothers. Throughout this period, banks stopped lending to each other at all but the shortest maturities and the quantity of interbank lending declined dramatically. There was also an unprecedented increase in the amounts placed by banks as reserves at the major central banks. The breakdown of global funding markets and the hoarding of liquidity by banks has profoundly disrupted financial intermediation in most advanced economies, prompting wide-scale reassessment of financial stability policies.

In most accounts of the crisis, the behaviour of the premium charged for interbank loans and the quantity of lending on offer is ascribed to some combination of greater preference for liquidity and concerns about counterparty risk (see, for example, BIS Annual Report, 2009, p20). Figure 2 (a,b,c) decomposes the interbank spread in the sterling, US dollar, and euro markets into contributions from credit premia and non-credit premia. While it is hard to disentangle empirically the relative importance of counterparty credit risk and banks' concerns about their own future liquidity needs, Figure 2 implies that precaution may have been a key factor driving funding costs in both August 2007 and September 2008. Indeed, the gradual build-up in credit risk over the period may have reflected the weak funding position of many banks as a consequence of the precautionary hoarding of others. Anecdotal evidence from the banks themselves, as well as recent empirical work by Acharya and Merrouche (2009) and Christensen et al (2009), lends credence to this view.

The events in the interbank markets raise some important questions that, as yet, do not have satisfactory answers. First, why was the seizure so sudden and why had it not happened before? Second, why did interbank markets dry up across the global financial system, rather than in localised areas where asset quality problems were clearly evident? Third, how crucial were the big, highly connected, interbank lenders in spreading contagion? And, finally, what might policymakers be able to do to avert such funding contagion in the future?

In this paper, we draw on network theory to show how system-wide interbank breakdowns are rare events that can be driven by precautionary motives alone. The interbank market is modelled as a large network of interlocking balance sheets with nodes in the network representing financial institutions, while links between nodes reflect lending relationships. We perturb the network by supposing that a small subset of banks is exogenously forced to raise resources at short notice,

perhaps to finance the off-balance sheet vehicles they are committed to rescuing. Because of liquidity needs, they become “hoarders” and stop lending to other banks (except, perhaps, at extremely short maturities).

The counterparties of these hoarders are faced with several options. Most obviously, they could expand unsecured funding in wholesale markets from other banks that do not hoard. In most circumstances, this is likely to be sufficient and we attach a high probability to this event, though it decreases with the overall fraction of ‘hoarders’ in the system. With a small probability, however, counterparties are unable to raise sufficient resources in this fashion. If they hold large stocks of liquid assets, this is unlikely to be a problem as they can sell or repo these assets to make good their funding shortfall. But if their liquid assets are insufficient, they too may become “hoarders”, withdrawing lending from their own counterparties.¹ Precautionary hoarding is thus sufficient to lead to a collapse of the interbank market, with liquidity evaporating even in the absence of any increase in initial counterparty credit risk. Banks do not internalise the consequences of their hoarding on others in the network, with potentially devastating system-wide repercussions.

The analysis maps how an initial outbreak of liquidity hoarding spreads contagiously through the financial system. Once a critical threshold of banks become “hoarders”, a sudden illiquidity cascade through the entire system is possible, with all banks refusing to lend in interbank markets. Although the probability of this eventuality is very low, the spread of contagion envelopes the entire financial system. *Phase transitions* such as these can come into play if banks – as has been the case in recent years – become highly reliant on wholesale funding and have low stocks of liquid assets. Intuitively, if banks are less reliant on wholesale funding overall, or have large liquid asset buffers, they are less affected by the loss of some of their funding sources and are likely to continue lending in interbank markets.

We also show how some banks are more important to the system than others. In particular, highly connected banks (in terms of their reach as interbank lenders) and banks with large volumes of interbank loans are instrumental in the spread of funding contagion; a result that resonates with the seminal work of Anderson and May (1991) and Albert et al (2000) on the role of critical nodes in network resilience. Haldane (2009) points to the lack of any meaningful relationship between banks’ systemic importance and their Basel capital buffers as evidence that regulators paid scant regard to financial stability considerations when designing prudential policy

¹In extremis, banks may also be called upon to sell their illiquid assets in a fire sale, depressing the price of these assets and forcing others to make mark-to-market losses as a result (see Cifuentes et.al. 2005). Though not the focus of our paper, direct funding contagion due to interbank relationships can also be amplified by price effects in illiquid asset markets.

prior to the crisis. Our findings, thus, offer support for the view that policy measures should target key institutions or “*super-spreaders*”. As Jenkinson (2009) has also argued, systemic liquidity requirements may thus contribute significantly towards financial stability.² In addition, there may also be a case for promoting greater liquidity self-insurance and more diversified funding by banks in future.

In exploring the probability and spread of funding contagion, we model the interbank network using random graph methods.³ A standard criticism of random graph models is that they are context free and fail to capture the economic incentives central to network formation. We therefore use novel techniques that permit the structure of links implied by interbank claims to be entirely arbitrary (Callaway et al, 2000; Newman et al, 2001). Crucially, it means that our results encompass *any possible network* – including those that represent the optimal outcome of network formation games. So we are able to side-step the important issue of the (optimal) formation of interbank markets, without loss of generality, in order to concentrate on the probability and spread of contagion. This is counter to existing work on networks in finance, where strong assumptions about network structure are needed to generate analytical results (e.g. Allen and Gale, 2000; Freixas et al, 2000; Leitner, 2005; Castiglionesi and Navarro, 2007).⁴ As Jackson (2008) emphasises, it also means that our framework is able to provide a natural reference point from which to assess models of the interbank market that are explicitly based on behavioural assumptions.

In what follows, we model the contagion process – analytically and numerically – by describing the behaviour of connected groups in a network and predicting the size of a susceptible cluster. In other words, we identify the number of vulnerable nodes reached via the transmission of shocks along the links of the network. This requires specifying all possible patterns of future transmission. Following Newman et al (2001), we use probability generating function techniques to obtain the number of a randomly selected node’s first neighbours, second neighbours and so on. Recursive equations are constructed to consider all possible outcomes and derive the total number of nodes that the original node is connected to – directly and indirectly. A phase transition or ‘tipping point’ which marks the threshold for an extensive contagious outbreak can then be identified. If, however, no phase transition exists, widespread contagion is impossible because distressed banks cannot spread the shock sufficiently frequently.⁵

²Morris and Shin (2008) and BIS Annual Report (2009) make a related argument for systemic capital charges.

³See Jackson (2008, Chapter 4) for a comprehensive review of network models of this type and their applications to economics.

⁴For a survey of network models in finance, see Allen and Babus (2008)

⁵Gai and Kapadia (2008) show how these techniques can be used to explore asset-side contagion and firesale effects. There is also a large empirical literature exploring contagion through interbank linkages (see Upper, 2007 for a comprehensive survey). But this literature has not focused on liability-side or “funding” contagion.

Our contribution is, thus, also a methodological one, showing how generalised random graph models can be applied to problems in economics where diffusion plays an important part. Although reviewed in Jackson (2008), use of these methods largely resides in fields outside economics, such as statistical physics (Watts, 2002) and epidemiology (Anderson and May, 1991; Newman, 2002; Meyers, 2007; Jackson and Rogers, 2007).⁶ But our paper differs from these contributions in three key respects. First, unlike epidemiology where links are undirected, our model explicitly characterises the direction of claim and obligations on banks' balance sheets. Second, in epidemiological models, the susceptibility of an individual to contagion from a particular 'infected' neighbour does not depend on the health of their other neighbours. By contrast, in our set-up, funding problems at a particular institution following hoarding by one counterparty is more likely if other counterparties are also hoarding. Finally, higher connectivity in epidemiology simply creates more channels of contact through which infection can spread, increasing the potential for contagion. In our model, greater connectivity provides counteracting risk-sharing benefits as funding sources are diversified across a wider set of institutions.

The structure of the paper is as follows. Section 2 reviews some of the recent literature on the interbank freeze and notes our contribution in relation to this. Section 3 describes the model, the transmission process for funding contagion, and presents results on the probability and spread of precautionary hoarding. Section 4 considers why interbank market seizures might be rare events and examines the consequences of '*super-spreaders*' in the network. Section 5 discusses the policy implications of our analysis, and a final section concludes with suggestions for future work.

2. Recent Literature on the Interbank 'Freeze'

A number of complementary papers have also sought to explain the breakdown in the interbank market.⁷ In general, these papers are explicitly behavioural and emphasise imperfections, other than network effects, that influence banks' decisions to form credit relationships and to generate potential collapses in interbank lending.

Allen et al (2009) present a model of interbank liquidity without counterparty risk or information asymmetry in which banks stop lending and hoard liquidity because incomplete financial markets do not allow them to hedge aggregate liquidity shocks. In Diamond and Rajan (2009), banks can only sell assets to realise cash but face a limited set of potential buyers. Faced with

⁶Encouragingly, prominent authors (e.g. May et al 2008) and policymakers (e.g. Haldane, 2009) are beginning to highlight the parallels between financial systems and complex systems in other fields and have called on economists to draw on these techniques when analysing systemic risk.

⁷Related papers that study the functioning of the interbank market also include Bhattacharya and Gale (1987), who show how banks underinvest in liquid reserves in the face of moral hazard, and Repullo (2005) who highlights banks' incentives to free-ride on central bank liquidity.

a liquidity shock, banks are deterred from asset sales because the alternative of holding them is more beneficial. The induced illiquidity in the market for these assets acts to depress lending to others in the interbank market, prompting a credit freeze.

Other papers invoke asymmetric information to generate a market breakdown. Heider et al (2009) place counterparty risk at the centre of their analysis of liquidity hoarding. Banks become privately informed about the risk of their illiquid assets after they choose their asset portfolios. They show how the interbank market becomes impaired by adverse selection as suppliers of liquidity attempt to protect themselves from lending to ‘lemons’. In Bolton et al (2008), long-term investors cannot tell whether short-term investors sell because of adverse asset quality or genuine liquidity needs. The price discount is exacerbated as potential sellers learn about the asset. When confronted with a decision to sell now to meet a liquidity need or risk future sales at a large discount, late trading equilibria that resemble a freeze become possible.

Exogenous shifts in information structure and Knightian uncertainty have also been emphasised as important factors in recent work. Caballero and Krishnamurthy (2008) suggest that the recent bout of financial innovation in credit markets heightened Knightian uncertainty in financial markets and use this to justify flights to quality and hoarding. In Acharya et.al (2009) runs in short-term funding markets arise when the information structure in the market shifts exogenously from being optimistic about asset quality to pessimistic. In their paper, “no news is good news” becomes “no news is bad news”, and they show that if good news arrives at a slower rate than that at which debt is rolled over, interbank lending dries up.

Finally, in a contribution that is most closely related to ours, Caballero and Simsek (2009) perturb a financial network in order to model the spread of precautionary hoarding. But whereas our paper uses a generalised random graph as a metaphor for the complexity of real world financial networks, Caballero and Simsek appeal to the rising costs of understanding the structure of the network as the basis for complexity and, ultimately, the reason for hoarding. Their model adopts a “ring-like” network structure of identical banks with only one neighbour serving as a funding counterparty. Banks must learn about the health of the trading partners of the trading partners of their trading partners, and so on. If information about the network structure is costless, there is no hoarding. But if, following a liquidity shock, these information costs rise sharply, banks’ inability to understand the structure of the network to which they belong leads them to withdraw from their loan commitments.

In contrast to the existing literature, this paper places a complex network at the centre of the analysis and puts particular emphasis on the precautionary hoarding behaviour which, as Figure 2

demonstrates, was central to the systemic liquidity crisis. Our paper thus offers a new perspective on recent events which can help to explain several features of the interbank market collapse. In particular, it identifies a clear network externality which can cause widespread liquidity hoarding to occur even in the complete absence of any counterparty risk; demonstrates how a system-wide collapse may be possible even if only a small subset of the system is affected initially; and highlights how some banks may be considerably more important in triggering an interbank market collapse than others, depending on their position and importance in the network.

3. The Model

3.1. Model Structure

Consider a financial network in which n financial intermediaries, ‘banks’ for short, are randomly linked together by their claims on each other. Each bank is represented by a *node* on the network, and the interbank exposures of bank i define the *links* with other banks. These links are *directed*, reflecting the fact that interbank connections comprise both assets and liabilities. Figure 3 illustrates a directed financial network with five banks.

Our paper characterises the global interbank market by an arbitrarily large network of this kind. The model is flexible enough to allow for heterogeneity in bank balance sheets and, as such, is capable of capturing the role of key players such as large investment banks and money market mutual funds. Importantly, the generality of the assumed network structure implies that our analytical results hold for *any possible network* including, for example, an interbank network that arises as the optimal outcome of a prior network formation game. They also hold for real-world networks, such as the UK interbank market depicted in Figure 4. In the analysis that follows, we illustrate the validity of our results with a stylised (Poisson) distribution of interbank exposures, as well with more realistic (exponential) distributions.

More formally, let the links that point into a bank (or node) represent the interbank assets of that bank (i.e. money owed to the bank by a counterparty) and the outgoing links be the interbank liabilities (i.e. money borrowed by the bank from a counterparty). Denoting the number of incoming links, or *in-degree*, to bank i by j_i , and the number of outgoing links, or *out-degree*, by k_i , we define the *joint degree distribution* of in- and out-degree, p_{jk} , to be the probability that a randomly chosen bank simultaneously has in-degree j and out-degree k .

The joint distribution of in- and out-degree and, hence, the structure of interbank exposures plays a key role in determining how shocks spread through the network. In our model, we treat this joint degree distribution as being entirely arbitrary. This implies that the network is entirely

random in all respects other than its degree distribution. In particular, there is no statistical correlation between nodes and mixing between nodes is proportionate (i.e. there is no statistical tendency for highly connected nodes to be particularly connected with other highly connected nodes or with poorly connected nodes).

Since every interbank asset of a bank is an interbank liability of another, every outgoing link for one node is an incoming link for another node. This means that the average in-degree in the network, $\frac{1}{n} \sum_i j_i = \sum_{j,k} j p_{jk}$, must equal the average out-degree, $\frac{1}{n} \sum_i k_i = \sum_{j,k} k p_{jk}$. We refer to this quantity as the *average degree* and denote it by

$$z = \sum_{j,k} j p_{jk} = \sum_{j,k} k p_{jk}. \quad (1)$$

Figure 5 presents the composition of individual bank balance sheets in the model. The total liabilities of each bank are comprised of interbank liabilities, L_i^{IB} , retail deposits, D_i , and capital, K_i . We assume that the total interbank liability position of every bank is evenly distributed over each of its outgoing links and is independent of the number of links the bank has (if a bank has no outgoing links, $L_i^{IB} = 0$ for that bank). These assumptions do not affect any of our main results. However, they serve to maximise diversification for a given number of links and the extent to which banks can diversify themselves by forming new outgoing links.

Since every interbank liability is another bank's asset, interbank assets, A_i^{IB} , are endogenously determined by the network links. So, although total interbank assets equal total interbank liabilities in aggregate across the entire system, each individual bank can have a surplus or deficit in their individual interbank position. Apart from interbank assets, banks also hold retail assets, such as mortgages, A_i^R and liquid assets, such as cash or government bonds, A_i^L .⁸

Now consider the potential impact of liquidity hoarding in the network. In particular, suppose a fraction μ of banks 'hoard' liquidity from bank i , that is, they withdraw their deposits held at bank i and convert them into liquid assets. In the model presented here, we assume that this represents a genuine drain on the liquidity of the entire banking system which flows outside, for example, ending up as increased reserves at central banks. But hoarding behaviour can also be interpreted as a switch from lending at long maturities (e.g. for three months) to lending at much shorter maturities (e.g. overnight) without changing the results.

Under full withdrawal (i.e. the lending bank withdraws their entire deposit), bank i loses μL_i^{IB} of its liabilities. So bank i will be 'liquid' without selling retail assets or starting to hoard

⁸Formally, A_i^L is taken to be a random variable – the underlying source of its variability may be viewed as being generated by variability in A_i^R , drawn from its appropriate distribution. For notational simplicity, we do not explicitly denote this dependence of A_i^L on A_i^R in subsequent expressions.

itself if

$$A_i^L - \mu L_i^{IB} + L_i^N - \varepsilon_i > 0 \quad (2)$$

where L_i^N represents new interbank borrowing that the bank is able to raise (we assume that banks are unable to raise new retail deposits given the time lags involved) and ε_i represents a liquidity shock. The shock may, for example, be interpreted as stemming from an actual off-balance sheet liquidity demand on bank i , perhaps via committed line to special investment vehicle or conduit, or a fear that such a liquidity line might be drawn. Alternatively, it could be regarded as deriving from an aggregate liquidity shock which has heterogeneous effects across banks. Note that the liquidity condition can also be expressed as

$$\mu < \frac{A_i^L + L_i^N - \varepsilon_i}{L_i^{IB}}, \text{ for } L_i^{IB} \neq 0. \quad (3)$$

If a bank is illiquid according to condition (2), then it needs to take action to avoid defaulting on required payments. One option would be for the bank to liquidate retail assets, A_i^R , in a fire sale. But selling assets cheaply is unlikely to be an attractive option. Diamond and Rajan (2009) suggest that the risk shifting incentives of banks can make them hesitant to enter a firesale because the alternative of holding on to assets are more beneficial. And potential stigma effects of asset sales for the liquidating bank are a further deterrent, placing it at risk of losing funding from its remaining interbank debtholders.

A second option is for the bank to hoard liquidity itself by withdrawing its own lending, A_i^{IB} , from others in the interbank market. This type of hoarding is purely due to the bank's own liquidity needs; concerns over the solvency of its counterparty play no role in its decision. While hoarding may also entail reputational costs, it is likely to be less costly than a firesale of assets. And since interbank transactions are over-the-counter, adverse stigma effects can be kept to a minimum.

During the recent crisis, liquidity hoarding has been frequently and widely observed in advanced country banking systems. Fire sales have been much less in evidence. Therefore, in what follows, we assume that banks which are unable to satisfy (2) choose to raise resources by hoarding liquidity and withdrawing their interbank assets held against other banks in the system. Clearly, it is possible that hoarding alone may not allow banks to raise sufficient resources to meet their obligations, at which point fire sales become an option. But since hoarding behaviour alone is sufficient to generate an interbank freeze in our model, we leave aside consideration of subsequent defensive actions and any feedbacks which may be associated with these.

To model the dynamics of funding contagion, we perturb the interbank network. Specifically,

let all banks be initially liquid according to condition (2) and suppose only a single bank suffers from a sufficiently adverse liquidity shock which forces it to hoard liquidity (so $\varepsilon_i = 0$ for all banks apart from one bank). This bank is assumed to raise the resources it needs by withdrawing funding from *all* of its counterparties. This assumption seems plausible – immediately accessible deposits are only likely to be available in relatively small amounts from each counterparty – and is more appropriate than assuming that funding withdrawal is concentrated on a small subset of its counterparties.

Recall that k_i denotes the number of outgoing links for bank i . Given the full withdrawal assumption, each linked bank loses a fraction $1/k_i$ of their interbank liabilities when a single counterparty hoards. From (3), it is clear that the only way hoarding can spread is if there is a neighbouring bank for which

$$\frac{A_i^L + L_i^N}{L_i^{IB}} < \frac{1}{k_i}. \quad (4)$$

We define banks that are exposed in this sense to hoarding by a single neighbour as *vulnerable* and other banks as *safe*. The vulnerability of a bank clearly depends on its out-degree, k . Specifically, a bank with out-degree k is vulnerable with probability

$$v_k = P \left[\frac{A_i^L + L_i^N}{L_i^{IB}} < \frac{1}{k} \right] \quad \forall k \geq 1. \quad (5)$$

Further, the probability of a bank having in-degree j , out-degree k and being vulnerable is $v_k \cdot p_{jk}$. Note that equations (4) and (5) specify a threshold rule that effectively introduces *local dependencies* and, hence, a network externality into our model. The effect that a single vulnerable counterparty will have on a given bank depends critically on whether the bank's other counterparties are hoarding or not.

Our assumption of full withdrawal acts as a key amplifying mechanism in the model. It means that banks further down the chain of contagion can be hit by fairly large shocks. In reality, hoarding banks may only withdraw some proportion of their deposits because their liquidity needs are limited or contractual obligations prevent them from withdrawing the full amount immediately. Arguably, however, if a borrowing bank has lost some portion of its deposits from a hoarding counterparty, it will only be a matter of time before the full amount is lost. Further withdrawals may occur when contractual obligations expire. And time consistency issues make it difficult for the lender to commit not to making further withdrawals in future. Moreover, even if current withdrawal is only partial, a forward looking bank may choose to act immediately *as if* it had lost its entire deposit, in order to limit the prospect of suffering liquidity problems in the future. As such, the full withdrawal assumption provides a simple mechanism for capturing

a rich set of dynamics operating through forward-looking expectations.

3.2. Funding Contagion

This structure presented above describes all of the key features of the model. In what follows, we first show that system-wide illiquidity cascades are possible when banks are not able to access new interbank funding sources to replace lost deposits (i.e. $L_i^N = 0$) by establishing that, with positive probability, a random case of liquidity hoarding at one bank can lead to the spread of hoarding across the entire interbank market. Section 4. then shows that when we relax this assumption, system-wide illiquidity cascades are still possible, but only occur very rarely.

If banks are unable to access new interbank funding (i.e. $L_i^N = 0 \forall i$), equation (5) simplifies to

$$v_k = P \left[\frac{A_i^L}{L_i^{IB}} < \frac{1}{k} \right] \quad \forall k \geq 1. \quad (6)$$

For any two banks, i and j , contagion can only potentially spread from i to j via an incoming link starting at bank j and ending at bank i , since this represents a situation where bank i lends to bank j . So, in sufficiently large networks, if contagion is to spread beyond the first neighbours of the initially hoarding bank, those neighbours must themselves have incoming links from other vulnerable banks (i.e. they must have lent to other banks).⁹

Define the probability generating function for the joint degree distribution of a vulnerable bank as

$$\mathcal{G}(x, y) = \sum_{j,k} v_k \cdot p_{jk} \cdot x^j \cdot y^k. \quad (7)$$

The generating function contains all the same information that is contained in the degree distribution, p_{jk} , and the vulnerability distribution, v_k , but in a form that allows us to work with sums of independent draws from different probability distributions. Importantly, it generates all the moments of the degree distribution of only those banks that are vulnerable. The appendix provides a detailed description of generating functions, highlighting the properties used in this paper.

We seek generating functions that characterise the distribution of the number of incoming links and the distribution of the sizes of connected components or ‘clusters’ in the random graph.

⁹If the number of nodes, n , is sufficiently large, banks are highly unlikely to have borrowed from more than one hoarding bank after the first round of contagion, meaning that safe banks will never start to hoard liquidity in the second round. This assumption clearly breaks down either when n is small or when contagion spreads more widely. However, the logic of this section still holds in both cases: in the former, the exact solutions derived for large n will only approximate reality; in the latter, the exact solutions will apply but the extent of contagion will be affected, as discussed further below.

From $\mathcal{G}(x, y)$, we define a single-argument generating function, $G_0(x)$, for the number of links ending up at a randomly chosen vulnerable bank. This is given by

$$\begin{aligned} G_0(x) &= \mathcal{G}(x, 1) \\ &= \sum_{j,k} v_k \cdot p_{jk} \cdot x^j. \end{aligned} \tag{8}$$

Note that

$$\begin{aligned} \mathcal{G}(1, 1) &= G_0(1) \\ &= \sum_{j,k} v_k \cdot p_{jk} \end{aligned} \tag{9}$$

so that $G_0(1)$ yields the fraction of banks that are vulnerable.

We define a second single-argument generating function, $G_1(x)$, for the number of links ending up at a bank which has been reached via a randomly chosen outgoing link from that bank. Because we are interested in the propagation of shocks from one bank to another, we require the degree distribution, $v_k \cdot r_{jk}$, of a vulnerable bank that is a random neighbour of our initially chosen bank. The larger the out-degree of this second bank, the more likely it is to be a neighbour and lie at the start of a randomly chosen inward link.¹⁰ So the probability of choosing it is proportional to kp_{jk} and the corresponding generating function is

$$G_1(x) = \sum_{j,k} v_k \cdot r_{jk} \cdot x^j = \frac{\sum_{j,k} v_k \cdot k \cdot p_{jk} \cdot x}{\sum_{j,k} k \cdot p_{jk}}. \tag{10}$$

We now describe the distribution of the *cluster* of vulnerable banks that can be reached by following a randomly chosen directed link, following an initial case of liquidity hoarding. The size and distribution of the vulnerable cluster characterises the spread of hoarding across the financial network.

Let $H_1(x)$ be the generating function for the probability of reaching an incoming cluster of given size by tracing back along a random incoming link after an initial case of liquidity hoarding. At this point, we assume that the links into the hoarding node are tree-like and contain no cycles or closed loops. This is solely to make an exact solution possible: the thrust of the argument goes through without this restriction and we do not apply it when conducting our numerical simulations.

As shown in Figure 6, the total probability of all possible forms can be represented self-consistently as the sum of probabilities of hitting a safe bank, hitting only a single vulnerable

¹⁰See Feld (1991) and Newman (2003) for a detailed analysis of this point.

bank, hitting a single vulnerable bank connected to one other component, two other components, and so on. Each component which may be arrived at is independent. Therefore, $H_1(x)$ satisfies the following self-consistency condition:

$$H_1(x) = \Pr[\text{reach safe bank}] + x \sum_{j,k} v_k \cdot r_{jk} \cdot [H_1(x)]^j, \quad (11)$$

where the leading factor of x accounts for the one vertex at the end of the initial edge and we have used the fact that if a generating function generates the probability distribution of some property of an object, then the sum of that property over m independent such objects is distributed according to the m^{th} power of the generating function (see the appendix). By using equation (10) and noting that $G_1(1)$ represents the probability that a random neighbour of a vulnerable bank is vulnerable, we may write equation (11) in recursive form as

$$H_1(x) = 1 - G_1(1) + xG_1(H_1(x)). \quad (12)$$

Finally, consider the distribution of incoming vulnerable cluster sizes to which a randomly chosen bank belongs. There are two possibilities that can arise. First, a randomly chosen bank may be safe. Second, it may have in-degree j and out-degree k , and be vulnerable, the probability of which is $v_k \cdot p_{jk}$. In this second case, each incoming link connects to a vulnerable cluster whose size is drawn from the distribution generated by $H_1(x)$. So the size of the vulnerable cluster to which a randomly chosen bank belongs is generated by

$$\begin{aligned} H_0(x) &= \Pr[\text{bank safe}] + x \sum_{j,k} v_k \cdot p_{jk} \cdot [H_1(x)]^j \\ &= 1 - G_0(1) + xG_0[H_1(x)]. \end{aligned} \quad (13)$$

And, in principle, we can calculate the complete distribution of cluster sizes by solving equation (12) for $H_1(x)$ and substituting the result into equation (13).

Although it is not usually possible to find a closed-form expression for the complete distribution of cluster sizes in a network, we can obtain closed form expressions for the average properties of clusters from equations (12) and (13). In particular, the average vulnerable cluster size, \mathcal{S} , is given by

$$\mathcal{S} = H_0'(1). \quad (14)$$

Noting that $H_1(x)$ is a standard generating function so that $H_1(1) = 1$ (see the appendix), it follows from equation (13) that

$$\begin{aligned} H_0'(1) &= G_0[H_1(1)] + G_0'[H_1(1)] H_1'(1) \\ &= G_0(1) + G_0'(1) H_1'(1). \end{aligned} \quad (15)$$

And we know from equation (12) that

$$H'_1(1) = \frac{G_1(1)}{1 - G'_1(1)}. \quad (16)$$

So, substituting equation (16) into (15) yields

$$\mathcal{S} = G_0(1) + \frac{G'_0(1)G_1(1)}{1 - G'_1(1)} \quad (17)$$

From equation (17), it is apparent that the points which mark the *phase transitions* at which the average vulnerable cluster size diverges are given by

$$G'_1(1) = 1, \quad (18)$$

or, equivalently, by

$$\sum_{j,k} j \cdot k \cdot v_k \cdot p_{jk} = z \quad (19)$$

where we have used equations (1) and (10).

The term $G'_1(1)$ is the average in-degree of a vulnerable first neighbour, counting only those links that join up to another vulnerable bank. If this quantity is less than one, all vulnerable clusters are small and contagion dies out quickly since the number of vulnerable banks reached declines. But if $G'_1(1)$ is greater than one, a “giant” vulnerable cluster – a vulnerable cluster whose size scales linearly with the size of the whole network – exists and occupies a finite fraction of the network. In this case, system-wide funding contagion is possible – with positive probability, a initial case of hoarding at one bank can lead to hoarding across the entire vulnerable portion of the interbank market.

3.3. The Probability and Spread of Contagion

From a system stability perspective, we are primarily interested in contagion within the giant vulnerable cluster. As z increases, the $\sum_{j,k} j \cdot k \cdot p_{jk}$ term in equation (19) increases monotonically but v_k falls. So equations (18) and (19) will either have two solutions or none at all. In the first case, there are two phase transitions and a continuous window of (intermediate) values of z for which contagion is possible. For values of z that lie outside the window and below the lower phase transition, the $\sum_{j,k} j \cdot k \cdot p_{jk}$ term is too small and the network is insufficiently connected for contagion to spread (consider what would happen in a network with no links); for values of z outside the window and above the upper phase transition, the v_k term is too small and contagion cannot spread because there are too many safe banks.

Contagion only emerges for intermediate values of z , and only when the initially hoarding bank is either in the giant vulnerable cluster or directly adjacent to it. The likelihood of contagion

is, therefore, directly linked to the size of the vulnerable cluster within the window.¹¹ Intuitively, near both the lower and upper phase transitions, the probability of contagion must be close to zero since the size of the vulnerable cluster is either curtailed by limited connectivity or by the presence of a high fraction of safe banks. The probability of contagion is thus non-monotonic in z : initially, greater connectivity increases the size of the vulnerable cluster and the probability of contagion; eventually, however, banks' funding becomes so diversified across different counterparties that the number vulnerable banks falls, and so does the probability of contagion.¹²

Near the lower phase transition, the *conditional* spread of contagion (i.e. conditional on contagion breaking out) corresponds to the size of the giant vulnerable cluster – average connectivity is so low at this point that no bank is safe and the spread of contagion is only contained by the fact that some nodes are completely unconnected to the network. But, for higher values of z , once contagion has spread through the entire vulnerable cluster, the assumption that banks are adjacent to no more than one hoarding bank breaks down. So 'safe' banks may become susceptible and contagion can spread well beyond the vulnerable cluster. Therefore, for high values of z , the fraction of banks affected by episodes of contagion can be greater than the probability of contagion breaking out, with the difference being magnified as z increases.

We illustrate the phase transitions and contagion window by simulating the model numerically. Since our theoretical model applies to random graphs with arbitrary degree distributions, we highlight results for both Poisson and geometric degree distributions. Our baseline results are, however, derived from a Poisson distribution to maximise clarity. With the Poisson distribution, each possible directed link in the graph is present with independent probability p . So the interbank network is constructed by looping over all possible directed links and choosing each one to be present with probability p – note that this algorithm does not preclude the possibility of cycles in the generated network. We also allow for the possibility that two banks can be linked to each other in both directions – no netting of exposures is assumed.

The interbank network in the simulation comprises 250 banks. Although the model readily applies to networks of fully heterogeneous banks, we treat the interbank liabilities, capital buffers and liquid asset holdings of each bank to be identical for the purpose of illustration.¹³ We use

¹¹Note that this is not given by (17) since this equation is derived on the assumption that there are no cycles connecting subclusters. This will not hold in the giant vulnerable cluster.

¹²In the special case of a uniform (Poisson) random graph in which each possible link is present with independent probability p , an analytical solution for the size of the giant vulnerable cluster can be obtained using techniques discussed in Watts (2002) and Newman (2003). Since this does not account for the possibility of contagion being triggered by nodes directly adjacent to the vulnerable cluster, it does not represent an analytical solution for the probability of contagion. However, it highlights that the size of the giant vulnerable cluster, and hence the probability of contagion, is non-monotonic in z .

¹³Incorporating heterogeneity in banks' assets would simply widen the contagion window. Intuitively, with heterogeneity, some banks will still be vulnerable despite having well diversified funding sources, because their liquid assets can be relatively low compared to their interbank liability position.

information in the published accounts of banks in developed countries to help calibrate the assumed balance sheet structure. The liability side of the balance sheet is comprised of 15% interbank liabilities and 4% capital, with the remainder being customer deposits. Since each bank’s interbank liabilities are evenly distributed over their outgoing links, interbank assets are determined endogenously within the network structure. Liquid assets are set to be 1% of total assets, with the asset side being ‘topped up’ by retail assets until the total asset position equals the total liability position.

In our simulation, we vary the average degree, z , drawing 1000 realizations of the network for each value. In each draw, we expose a bank to a liquidity shock at random, forcing it to hoard liquidity. As a consequence, its counterparties may also become hoarders if they are at the threshold defined by equation (4). The iterative process continues until no new banks are pushed into hoarding. We count as “systemic” those episodes in which at least 5% of banks are forced into hoarding liquidity, and our results identify the frequency of systemic liquidity hoarding and its extent in terms of the average fraction of the system affected in each systemic outbreak.

Figure 7 confirms our analytical results for the case of the Poisson distribution. Contagion occurs for values of z between 0 and 30.¹⁴ Within this range, the probability of contagion is non-monotonic in connectivity, and occurs with certainty when z lies between 5 and 20. The conditional spread of contagion in this range is approximately the same as the frequency of contagion. So when any bank in a relatively well connected interbank network experiences a liquidity shock and engages in precautionary hoarding, the probability that this is sufficient to induce system-wide hoarding is very high. Indeed, even when $z > 30$, although contagion never occurs more than 20 times in 1000 draws due to the extremely high diversity of funding sources, the decision by a single bank to hoard can still lead to a freeze across the entire interbank network.

Figure 8 illustrates the corresponding results using a geometric degree distribution. This distribution exhibits fat tails, with some banks having substantially more connections than the average degree, and arguably offers a more realistic description of the real-world banking network. In the version implemented, we draw in-degree and out-degree separately from the same distribution, implying that there is no correlation between the number of counterparties a bank lends to and borrows from.¹⁵

The contagion window in this case is substantially wider than in the Poisson case – but the probability of contagion is again clearly non-monotonic. Noticeably the probability of conta-

¹⁴By way of comparison, the average connectivity, z , for US interbank payment flows (the Fedwire system) is estimated at around 15. See Soramaki et.al (2006) for details.

¹⁵To construct the graph, we also need to ensure that the total number of outgoing links drawn equals the total number of incoming links. We follow the algorithm outlined by Newman et al (2001) to achieve this.

gion for very highly connected networks ($z \geq 30$) remains quite significant at around 0.4. The conditional spread of contagion near the lower phase transition is again curtailed by the limited connectivity of the network. But as z rises, precautionary hoarding quickly envelopes the entire network, with around 90% of banks in the network engaging in hoarding once z exceeds 10. As with the Poisson distribution, the scale of systemic hoarding remains high even as the frequency of hoarding begins to decline, with a complete collapse in interbank lending possible for values of z close to 40.

4. Interbank Market Seizures and Key Players

Our assumption that banks could not access funding markets when confronted with liquidity problems is clearly unrealistic. We, therefore, relax this restriction and allow banks to obtain new interbank deposits, L_i^N , with some probability when they lose funding from one of their counterparties.

4.1. The Uncommon Nature of Interbank Collapses

The possibility of some access to new money turns an illiquidity cascade induced by precautionary hoarding into a *rare event* which, when it does occur, triggers a complete freeze in the interbank market. The intuition for this result can be seen by considering two polar cases.

First, suppose that either all banks can borrow from the interbank market (with probability a), or no bank can borrow (with probability $1 - a$). This reflects a situation of *perfect correlation* across closure of interbank markets to banks. Suppose, additionally, that a is close to 1.

Obviously, if interbank markets are open, then there cannot be a cascade. But if interbank markets are shut, then the results of section 3.2. continue to apply. The scale of funding contagion is unchanged. Note, however, that the probability of contagion now equals the probability of the cascade when replacement of interbank deposits is never possible, multiplied by $1 - a$. Since a is close to 1, this is a low probability event (see also Figure 9).

A second case supposes *no correlation* across the closure of interbank markets – so that a bank which loses funding can borrow with independent probability, b , where that probability is close to, but strictly less than, 1, and either constant or, more realistically, decreasing in the number of banks in the system which are already hoarding. Again, an interbank freeze induced by precautionary behaviour can occur but is a rare event.

This can be demonstrated by supposing that the probability, b , is constant. Although the likelihood of obtaining external funding is independent across banks, it is possible, in theory,

for *no* bank to get a good draw which allows it to obtain external funding at any stage, albeit with an extremely low probability. But in these very rare cases, the results of section 3.2. continue to apply, so widespread liquidity hoarding remains a possibility. And supposing that the probability of obtaining external funding, b , declined with the number of hoarding banks would clearly strengthen this result by increasing the chances that subsequent banks in the chain cannot obtain external funding.

To illustrate this result numerically, we take the more realistic case of declining probabilities. Specifically, denoting the number of banks already hoarding when bank i receives a shock by h_i , we suppose that the probability, b_i , that bank i can obtain external funds is given by:

$$b_i = 0.95 - \frac{h_i + 1}{n - h_i} \quad (20)$$

for $\frac{h_i+1}{n-h_i} < 0.95$, and 0 otherwise. Figure 10 presents the results in this case. As is evident, the contagion window is considerably smaller ($5 < z < 18$), with the probability of contagion peaking when z is about 12. But it is also clear that systemic liquidity hoarding remains a possibility.

4.2. Super-spreaders

As Figure 4 makes clear, real-world networks often have skewed degree distributions. As a result, such networks are often resilient to random perturbations because many of the nodes typically have low degree and so the random removal of nodes with degree zero or one has little effect on the connectivity of the remaining nodes in the network. But, as emphasised by Anderson and May (1991) and Albert et al (2000), the targeted removal of nodes with the highest degree can frequently have devastating effects. These nodes lie on many of the paths between pairs of other nodes and their removal can destroy the connectivity of a network in short order.

Taking the case described above, in which banks can obtain external funding but with declining probability, we now show how highly connected banks are more instrumental in causing contagion in our model. Before presenting numerical results, it is helpful to clarify some intuition analytically by considering an extreme case of declining probabilities, whereby b takes a value close to 1 when only one bank is hoarding, but is 0 when more than one bank hoards. In this instance, cascades will still have the same scale as in section 3.2. because the logic of that section will apply once contagion has spread to a second bank. But cascades can only occur if hoarding spreads beyond the initial bank. This requires both a link from the initial hoarding bank to another vulnerable bank *and* that bank to be unable to borrow.

Setting the probability that a random bank is vulnerable, $E(v_k)$, to be π , it is clear that each

neighbour will *not* need to hoard with probability

$$(1 - \pi) + \pi b = 1 - \pi(1 - b) \quad (21)$$

Recalling that any random bank (including the initial hoarding bank) has z connections on average, the expected probability that no neighbour will need to hoard following a random shock to one bank is therefore

$$[1 - \pi(1 - b)]^z \quad (22)$$

and the probability that at least one neighbour will need to hoard is:

$$1 - [1 - \pi(1 - b)]^z \quad (23)$$

For large b , this is a small number, though it is increasing in z . Therefore, in this simple example, the connectivity of the initially hoarding bank influences the probability that widespread contagion will occur. In particular, a bank with a high in-degree will be more connected than the typical bank (i.e. its connectivity will be greater than z). Therefore, from equation (23), the likelihood that one of its neighbours will have to hoard is greater. So, in this simple example in which a second round of hoarding is sufficient to lead to systemic collapse, it is clear that more highly connected banks are more likely to cause widespread contagion, though they do not have a role in amplifying the extent of contagion.

To demonstrate these results numerically, we follow the assumptions of the previous section, with b_i set according to (20). But instead of perturbing the network by exposing a random bank to a liquidity shock, the shock is instead targeted at the bank with either (a) the highest in-degree, i.e the one with the largest number of interbank lending relationships; or (b) the bank with the most interbank assets.

Figure 11 shows that a liquidity shock on a highly connected bank (or ‘*super-spreader*’) almost doubles the probability of funding contagion, compared to the case depicted in Figure 10. The likely impact of contagion is, however, similar in both the targeted and baseline case, reflecting the fact that once contagion has broken out beyond a few banks, the ‘tipping point’ is crossed and it is almost impossible to contain, irrespective of the source of the initial shock.

Targeting the bank with the most interbank assets produces very similar results (Figure 12). The size of the contagion window and the scale of hoarding in this case again closely match the baseline simulation. But, as with the attack on the bank with highest in-degree, the frequency of systemic hoarding also doubles inside the contagion window.

It is interesting that we obtain these even though the network in these simulations is drawn from a Poisson distribution, implying that the most-connected banks are not significantly more

connected than the typical bank. Although not shown here, targeted attacks on networks with fat-tailed degree distributions would be much more debilitating relative to random shocks in such networks. It is also worth noting that if the banks with highest in-degree (or most interbank assets) are progressively removed, one would need only remove a small fraction to achieve a complete breakdown in interbank lending.

4.3. *Relaxing the Diversification Assumptions*

Our analysis has assumed that the total interbank liability position of each bank is independent of the number of outgoing links to that bank, and that these liabilities were evenly distributed over each link. In reality, we might expect a bank with a higher number of outgoing links to have a larger total interbank liability position. Intuitively, this would curtail the benefits of greater connectivity because the greater absolute exposure to hoarding associated with a higher number of links would (partially) offset the positive effects from greater diversification. But, as long as the total interbank liability position increases less than proportionately with the number of links, our main results continue to apply. In particular, v_k will still decrease in z , though at a slower rate. And equation (19) will continue to generate two solutions, though in an extended range of cases. The contagion window will be wider still.¹⁶

Assuming an uneven distribution of interbank liabilities over outgoing links also does not change any of our fundamental results. In particular, v_k would still decrease in z , maintaining the possibility of two solutions to equation (19). But an uneven distribution makes banks vulnerable to liquidity hoarding by particular counterparties for higher values of z than would otherwise be the case. As a result, the contagion window will again be wider.

5. Policy Implications

As the model makes clear, the failure of banks to internalise the consequences of their hoarding behaviour provides justification for the regulation and supervision of liquidity risk in support of the public policy goal of financial stability.

Our findings point to the need to hold banks to higher standards of liquidity risk management in their internal treasury operations than they would naturally adopt given their own individual incentives. In particular, requiring banks to increase their stock of liquid assets, such as cash and government bonds, can play an important role in mitigating funding contagion risk. From

¹⁶If the total interbank liability position increases more than proportionately with the number of links, v_k will increase in z and greater connectivity will unambiguously increase contagion risk. This latter case does not seem a particularly plausible description of reality.

equation (5), for a given k , v_k is lower when banks have more liquid assets relative to their interbank liabilities. As equation (19) suggests, this implies that the contagion window is narrower meaning that illiquidity cascades are less likely.

Goodhart (2008) and Jenkinson (2008, 2009) argue forcefully that the trend decline in banks' liquidity positions in the United States and the United Kingdom over recent years was a major factor underlying the financial turmoil (see Figures 13 and 14). During recent years, in particular, banks economised on high quality liquid assets based on the (erroneous) assumption that high-yielding risky assets such as mortgage-backed securities would remain liquid. So a requirement for banks to enhance their liquid asset holdings may be an important element for future financial stability policy. But the optimal degree of self-insurance needed for system stability will be difficult to gauge.

More diversified funding also lowers the potential for funding contagion. In our model, the granularity of interbank funding is captured by average degree in the network. From (5), $E(v_k)$ is lower when the average degree is higher and equation (19), implies that the contagion window will be narrower. There may even be a case for imposing large liability limits on banks, restricting the amount that they can borrow from any single institution.

Our results also point to a case for prudential regulation to provide a disincentive for bank to increase their liquidity risk in proportion to the marginal contribution of a bank to systemic risk. This view has recently been advocated by several policymakers, including BIS (2009), Jenkinson (2009) and Haldane (2009). Large interbank lenders and highly connected banks are more instrumental in causing contagion, as Section 4 shows. So standards of resilience (in the form of liquidity requirements) could be made more stringent for large banks that are very active in interbank markets than for small banks with few connections to the rest of the network.

6. Conclusion

Our paper offers a network-theoretic explanation of the recent breakdown in global interbank markets. In so doing, we characterise the probability and potential spread of funding contagion. Interbank collapses are uncommon events that can be explained by precautionary hoarding by banks, even in the absence of any significant increase in the credit risk of counterparties.

The failure of banks to internalise the potential consequences of their hoarding behaviour for the system as a whole points to an important role for prudential policy. Our results suggest that it may be sensible to target key players in the interbank network, subjecting them to tougher liquidity requirements. More broadly, our findings point to the need to impose more exacting

standards for banks' liquidity risk management than had been the case pre-crisis. But gauging the optimal balance between the benefits of taking on liquidity risk and establishing policies that contain (system-wide) liquidity risk requires further work.

Our work also represents a first step in applying novel network techniques from other disciplines to applications in economics. Although our use of generalised random graph methods with arbitrary degree distributions enables us to encompass most network structures, including those based on behaviour, it is clearly desirable to explicitly account for banks' decisions and strategies. These considerations require game-theoretic reasoning adapted to genuinely rich network settings. They have the potential to yield important new insights for the design of financial stability policy and we intend to explore this avenue in future research.

Acknowledgments

We thank Jason Dowson for outstanding research assistance. We are especially grateful to Lord Robert May for his valuable comments and considerable encouragement. The helpful suggestions of Emily Beau, Nigel Jenkinson, Andy Haldane, Ouarda Merrouche, Andrea Prat and seminar participants at the Bank of Canada, the Bank of England, and the European Summer Symposium in Economic Theory (Gersensee, July 2009) are also gratefully acknowledged.

References

- Acharya, V, Gale, D and T Yorulmazer (2009), Rollover risk and market freezes, *mimeo*, NYU.
- Acharya, V and O Merrouche (2009), Precautionary hoarding of liquidity and inter-bank markets: evidence from the sub-prime crisis, *mimeo*, Bank of England.
- Albert, R, Jeong, H and A-L Barabasi (2000), Error and attack tolerance of complex networks, *Nature*, 406, 378-382.
- Allen, F and A Babus (2008), Networks in finance, forthcoming in *Network-based Strategies and Competencies*, Kleindorfer, P and J Wind (eds.), Wharton School Publishing.
- Allen, F, Carletti, E and D Gale (2009), Interbank market liquidity and central bank intervention, *Journal of Monetary Economics*, 56, 639-652.
- Allen, F and D Gale (2000), Financial contagion, *Journal of Political Economy*, 108, 1-33.
- Anderson, R and R May (1991), *Infectious diseases of humans: dynamics and control*, Oxford University Press.

- Bank of England (2007), *Quarterly Bulletin*, 47(4).
- Bhattacharya, S and D Gale (1987), Preference shocks, liquidity and central bank policy, in W Barnett and K Singleton (eds), *New Approaches to Monetary Economics*, Cambridge University Press, 69-88.
- BIS (2009), *79th Annual Report*, Bank for International Settlements, Basel, June.
- Bolton, P, Santos, T and J Scheinkman (2008), Inside and outside liquidity, *mimeo*, Princeton University.
- Caballero, R and A Krishnamurthy (2008), Collective risk management in a flight to quality episode, *Journal of Finance*, 63, 2195-2230.
- Caballero, R and A Simsek (2009), Complexity and financial panics, *mimeo*, MIT.
- Callaway, D, Newman, M, Strogatz, S and D Watts (2000), Network robustness and fragility: percolation on random graphs, *Physical Review Letters*, 85, 5468-5471.
- Castiglionesi, F and N Navarro (2007), Optimal fragile financial networks, *Tilburg University Discussion Paper #2007-100*.
- Christensen, J, Lopez, J and G Rudebusch (2009), Do central bank liquidity facilities affect interbank lending rates? *Mimeo*, Federal Reserve Bank of San Francisco.
- Cifuentes, R, Ferrucci, G and H S Shin (2005), Liquidity risk and contagion, *Journal of the European Economic Association*, 3, 556-566.
- Diamond, D and R Rajan (2009), Fear of firesales and the credit freeze, *NBER Working Paper #14925*.
- Freixas, X, Parigi, B and J-C Rochet (2000), Systemic risk, interbank relations and liquidity provision by the central bank, *Journal of Money, Credit and Banking*, 32, 611-638.
- Gai, P and S Kapadia (2008), Contagion in financial networks. Paper presented at a special session of the Royal Economic Society Annual Conference, Warwick, 17-19 March.
- Goodhart, C A E (2008), The background to the 2007 financial crisis, *International Economics and Economic Policy*, 4, 331-346.
- Haldane, A G (2009), Rethinking the financial network, *Speech to the Financial Student Association*, Amsterdam, 28 April.

- Heider, F, Hoerova, M and C Holthausen (2009), Liquidity hoarding and interbank spreads: the role of counterparty risk, *mimeo*, European Central Bank.
- Jackson, M (2008), *Social and Economic Networks*, Princeton University Press, NJ, USA.
- Jackson, M and B Rogers (2007), Relating network structure to diffusion properties through stochastic dominance, *The B.E. Journal of Theoretical Economics*, 7, 1-13.
- Jenkinson, N (2008), Strengthening regimes for controlling liquidity risk: some lessons from the recent turmoil, *Speech at the Euromoney Conference on "Liquidity and Funding Risk Management"*, London, 24 April.
- Jenkinson, N (2009), Containing system-wide liquidity risks: some issues and challenges, *Speech to the University of Frankfurt Conference on "The Law and Economics of Money and Finance in Times of Financial Crisis"*, Frankfurt, 15 May.
- Leitner, Y (2005), Financial networks: contagion, commitment and private sector bailouts, *Journal of Finance*, 6, 2925-53.
- May, R, Levin, S and G Sugihara (2008), Ecology for bankers, *Nature*, 451, 893-895.
- Meyers, L A (2007), Contact network epidemiology: bond percolation applied to infectious disease prediction and control, *Bulletin of the American Mathematical Society*, 44, 63-86.
- Morris, S and H Shin (2008), Financial regulation in a system context, *Brookings Papers on Economic Activity*, 2, 229-61.
- Newman, M (2002), The spread of epidemic disease on networks, *Physical Review E*, 66, 016128.
- Newman, M (2003), Random graphs as models of networks, in *Handbook of Graphs and Networks*, Bornholdt, S and H G Schuster (eds.), Wiley-VCH, Berlin.
- Newman, M, Strogatz, S and D Watts (2001), Random graphs with arbitrary degree distributions and their applications, *Physical Review E*, 64, 026118.
- Repullo, R (2005), Liquidity, risk taking and the lender of last resort, *International Journal of Central Banking*, 1, 47-80.
- Soramaki, K, Bech M, Arnold, J, Glass, R and W Beyeler (2006), The topology of interbank payment flows, *Federal Reserve Bank of New York Staff Papers #243*.
- Strogatz, S (2001), Exploring complex networks, *Nature*, 410, 268-276.

Upper, C (2007), Using counterfactual simulations to assess the danger of contagion in interbank markets, *Bank for International Settlements Working Paper #234*, August.

Watts, D (2002), A simple model of global cascades on random networks, *Proceedings of the National Academy of Sciences*, 99, 5766-5771.

7. Appendix: Generating Functions

Let Y be a discrete random variable taking values in $\{0, 1, 2, \dots\}$ and let $p_r = P[Y = r]$ for $r = 0, 1, 2, \dots$

Then the (*probability*) *generating function* of the random variable Y of the distribution, p_r ($r = 0, 1, 2, \dots$), is

$$G(x) = E(x^Y) = \sum_{r=0}^{\infty} x^r P[Y = r] = \sum_{r=0}^{\infty} p_r x^r.$$

Note that

$$G(1) = \sum_{r=0}^{\infty} p_r = 1.$$

Theorem 1 The distribution of Y is uniquely determined by the generating function, $G(x)$.

Proof Since $G(x)$ is convergent for $|x| < 1$, we can differentiate it term by term in $|x| < 1$. Therefore

$$G'(x) = p_1 + 2p_2x + 3p_3x^2 + \dots$$

and so $G'(0) = p_1$. Repeated differentiation gives

$$G^{(i)}(x) = \sum_{r=i}^{\infty} \frac{r!}{(r-i)!} p_r x^{r-i}$$

and so $G^{(i)}(0) = i!p_i$. Therefore, we can recover p_0, p_1, p_2, \dots from the generating function. ■

Theorem 2

$$E[Y] = \lim_{x \rightarrow 1} G'(x)$$

and, provided that x is continuous at $x = 1$,

$$E[Y] = G'(1).$$

Proof

$$G'(x) = \sum_{r=1}^{\infty} r p_r x^{r-1}$$

Therefore, for $x \in (0, 1)$, $G'(x)$ is a non-decreasing function of x , bounded above by

$$E[Y] = \sum_{r=1}^{\infty} r p_r.$$

Choose $\varepsilon > 0$ and N large enough that $\sum_{r=1}^N rp_r \geq E[Y] - \varepsilon$. Then

$$\begin{aligned} \lim_{x \rightarrow 1} \sum_{r=1}^{\infty} rp_r x^{r-1} &\geq \lim_{x \rightarrow 1} \sum_{r=1}^N rp_r x^{r-1} \\ &= \sum_{r=1}^N rp_r \geq E[Y] - \varepsilon \end{aligned}$$

Since this is true for all $\varepsilon > 0$,

$$\lim_{x \rightarrow 1} G'(x) = E[Y].$$

Provided that x is continuous at $x = 1$, the second result follows immediately. ■

Theorem 3

$$E[Y(Y-1)] = \lim_{x \rightarrow 1} G''(x)$$

and, provided that x is continuous at $x = 1$,

$$E[Y(Y-1)] = G''(1).$$

Proof

$$G''(x) = \sum_{r=2}^{\infty} r(r-1)p_r x^{r-2}$$

and the remainder of the proof is the same as the proof of theorem 2. ■

Theorem 4 If Y_1, Y_2, \dots, Y_n are independent random variables with generating functions given by $G_1(x), G_2(x), \dots, G_n(x)$, then the generating function of $Y_1 + Y_2 + \dots + Y_n$ is $G_1(x) \cdot G_2(x) \cdot \dots \cdot G_n(x)$.


Proof

$$E[x^{Y_1+Y_2+\dots+Y_n}] = E[x^{Y_1} \cdot x^{Y_2} \cdot \dots \cdot x^{Y_n}] \tag{24}$$

Since Y_1, Y_2, \dots, Y_n are independent random variables, the standard result from probability theory that functions of independent random variables are also independent implies that $x^{Y_1}, x^{Y_2}, \dots, x^{Y_n}$ are independent. Therefore, using the properties of expectation, we can rewrite (24) as


$$\begin{aligned} E[x^{Y_1+Y_2+\dots+Y_n}] &= E[x^{Y_1}] \cdot E[x^{Y_2}] \cdot \dots \cdot E[x^{Y_n}] \\ &= G_1(x) \cdot G_2(x) \cdot \dots \cdot G_n(x). \blacksquare \end{aligned}$$

Figure 1: Twelve-month Interbank Rates Relative to Expected Policy Rates^(a)




Sources: British Bankers' Association, Bloomberg and Bank of England calculations.
(a) Spread of twelve-month Libor to twelve-month overnight index swap (OIS) rates.

Figure 2(a): Decomposition of the Sterling 12-month Interbank Spread^{(a)(b)}



Sources: British Bankers' Association, Bloomberg, Markit Group Limited and Bank of England calculations.
(a) Spread of twelve-month Libor to twelve-month overnight index swap (OIS) rates.
(b) Estimates of credit premia are derived from credit default swaps on banks in the Libor panel. Estimates of non-credit premia are derived by the residual. For further details on the methodology, see Bank of England (2007, pp. 498-499).

Figure 2(b): Decomposition of the Dollar 12-month Interbank Spread^{(a)(b)}




Sources: British Bankers' Association, Bloomberg, Markit Group Limited and Bank of England calculations.

(a) Spread of twelve-month Libor to twelve-month overnight index swap (OIS) rates.

(b) Estimates of credit premia are derived from credit default swaps on banks in the Libor panel. Estimates of non-credit premia are derived by the residual. For further details on the methodology, see Bank of England (2007, pp. 498-499).

Figure 2(c): Euro Decomposition of the Euro 12-month Interbank Spread^{(a)(b)}



Sources: British Bankers' Association, Bloomberg, Markit Group Limited and Bank of England calculations.

(a) Spread of twelve-month Libor to twelve-month overnight index swap (OIS) rates.

(b) Estimates of credit premia are derived from credit default swaps on banks in the Libor panel. Estimates of non-credit premia are derived by the residual. For further details on the methodology, see Bank of England (2007, pp. 498-499).

Figure 3: A Directed Network with Five Nodes





Figure 4: Network of large exposures^(a) between UK banks^{(b)(c)} (reproduced from Bank of England, 2009)



Source: FSA returns.

- (a) A large exposure is one that exceeds 10% of a lending bank's eligible capital during a period. Eligible capital is defined as Tier 1 plus Tier 2 capital, minus regulatory deductions.
- (b) Each node represents a bank in the United Kingdom. The size of each node is scaled in proportion to the sum of (1) the total value of exposures to a bank, and (2) the total value of exposures of the bank to others in the network. The thickness of a line is proportionate to the value of a single bilateral exposure.
- (c) Based on 2008 Q1 data.

Figure 5: Stylised Balance Sheet




Figure 6: Transmission of Contagion implied by Equation (11)




Figure 7: Systemic Liquidity Hoarding (Poisson Distribution; No Replacement of Interbank Deposits; Random Shock)




Figure 8: Systemic Liquidity Hoarding (Geometric Distribution; No Replacement of Interbank Deposits; Random Shock)




Figure 9: Systemic Liquidity Hoarding (Poisson Distribution; Perfect Correlation in Interbank Market Closure; Random Shock)




Figure 10: Systemic Liquidity Hoarding (Benchmark Case: Poisson Distribution; Interbank Market Closed to Each Institution with Probability Defined by Equation (20); Random Shock)




Figure 11: Super-Spreaders (1) (Benchmark Case compared to Targeted Shock to Bank with Highest In-Degree)




Figure 12: Super-Spreaders (2) (Benchmark Case compared to Targeted Shock to Bank with most Interbank Assets)






Figure 13: US Banks' Holdings of Treasury Bonds



Source: FDIC Statistics on Depository Institutions

Figure 14: Sterling liquid assets relative to total asset holdings of UK banking sector^(a) (reproduced from Bank of England, 2008)



Source: Bank of England calculations.

- (a) 2008 data are as of end-August 2008.
- (b) Cash + Bank of England balances + money at call + eligible bills + UK gilts.
- (c) Proxied by: Bank of England balances + money at call + eligible bills.
- (d) Cash + Bank of England balances + eligible bills.