



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Complexity, concentration and contagion

Prasanna Gai^a, Andrew Haldane^b, Sujit Kapadia^{b,*}^a Department of Economics, University of Auckland, 12 Grafton Road, Auckland 1142, New Zealand^b Bank of England, Threadneedle Street, London EC2R 8AH, United Kingdom

ARTICLE INFO

Available online 30 May 2011

ABSTRACT

This paper develops a network model of interbank lending in which unsecured claims, repo activity and shocks to the haircuts applied to collateral assume centre stage. We show how systemic liquidity crises of the kind associated with the interbank market collapse of 2007–2008 can arise within such a framework, with funding contagion spreading widely through the web of interlinkages. Our model illustrates how greater complexity and concentration in the financial network may amplify this fragility. The analysis suggests how a range of policy measures – including tougher liquidity regulation, macro-prudential policy, and surcharges for systemically important financial institutions – could make the financial system more resilient.

© 2011 Bank of England. Published by Elsevier B.V. All rights reserved.

1. Introduction

Herbert Simon spent over half a century teaching at Carnegie-Mellon University. In his classic study of 'The Architecture of Complexity', he laid the foundations for evaluating the resilience and evolution of complex systems (Simon, 1962). His analysis drew upon a wide spectrum of systems thinking—from physical, biological and social sciences. Out of this, Simon reached the powerful conclusion that even complex systems tended to exhibit a basic simplicity. Systems could be arranged in a natural hierarchy, comprising nested sub-structures. Non-hierarchical structures would tend to be deselected over time because of their inefficiency or lack of robustness relative to simpler, hierarchical structures. Whether the complex system was biological, physical, or social, it would be a case of survival of the simplest.

The modern financial system seems to have bucked this evolutionary trend. In recent years, it has become much more complex, concentrated and interconnected.¹ As much as two-thirds of the spectacular growth in banks' balance sheets in the years prior to the crisis reflected increasing claims within the financial system, rather than with non-financial agents. In Simon's words, the financial system has become less modular, less hierarchical and thus less decomposable. In consequence, it became markedly more susceptible to systemic collapse. This sowed the seeds of the global financial crisis of 2007/2008.

In this paper, we explore how the complexity and concentration of financial linkages can give rise to systemic liquidity crises that threaten financial system resilience. In keeping with the multi-disciplinary spirit of Simon (1962), our theoretical framework draws upon network techniques developed in epidemiology and statistical physics to identify 'tipping points' in complex systems, whereby a small change in the underlying parameters or shocks can make a very large difference

* Corresponding author. Tel.: +44 20 7601 5507.

E-mail addresses: p.gai@auckland.ac.nz (P. Gai), andy.haldane@bankofengland.co.uk (A. Haldane), sujit.kapadia@bankofengland.co.uk (S. Kapadia).

¹ See Haldane (2009). It is natural to ask why the evolutionary forces described by Simon may not have forced an earlier deselection of such a complex system. One possible explanation is policy. Successive crises have seen a progressive widening and broadening of the financial safety net—from last resort lending from the 19th century onwards, to deposit insurance from the 1930s, to the block-buster capital and liquidity support for both banks and non-banks during this century (Alessandri and Haldane, 2009). Until the end of the Bretton Woods era, increasing implicit support went hand-in-hand with tougher financial regulation and greater restrictions on global capital flows. The 25-year postwar period was remarkable for its absence of systemic banking crises. But an era of deregulation since the mid 1970s may have contributed to increasing complexity and concentration.



Fig. 1. Network of large exposures between UK banks, 2008.

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 darkness of a line is proportionate to the value of a single bilateral exposure. (c) Based on 2008 Q1 data.

to outcomes. In light of recent events, the paper focuses on the collapse of the interbank market. We demonstrate both analytically and via numerical simulations how repo market activity, haircut shocks, and liquidity hoarding in unsecured interbank markets may have contributed to the spread of contagion and systemic collapse.² The network lens also offers new perspectives on a broad range of recent and proposed policy measures aimed at tackling financial system risk.

Despite obvious parallels between financial systems and complex systems in other fields (May et al., 2008; Haldane, 2009; Haldane and May, 2011), the use of network techniques from epidemiology and statistical physics to the study of financial contagion is in its infancy. One important reason for the slow up-take amongst economists is that such network techniques are typically silent about behavioural considerations. But they have the advantage of eclecticism about network formation processes and can be calibrated to match real-world networks (such as those depicted in Fig. 1). They are particularly well suited for dealing with heterogeneity of agents, charting the dynamic propagation of shocks within the financial system and identifying the non-linearities that characterise financial instability in a parsimonious way.³

In what follows, we examine the interplay between complexity, concentration and stability using a simple network model of the banking system in which individual banks are randomly linked together by their interbank claims on each other. Although stylised, the set-up is sufficiently rich to permit the study of shocks that diminish the availability of interbank loans, so-called 'funding liquidity shocks'. During the crisis, this took the form of liquidity hoarding by banks. As one bank calls in or shortens the term of its interbank loans, affected banks in turn do the same. The connectivity and concentration of the players in the network play key roles in this propagation mechanism.

We start by exploring analytically how such funding contagion can generate systemic liquidity crises.⁴ In doing so, we articulate how tipping points may be embedded in the financial network and show how these depend on the level of liquid asset holdings, the amount of interbank activity, and the size of haircuts on banks' assets.

We then consider numerical simulations which illustrate the potential fragility of the banking system and shed light on the design and implementation of policies which might help to address this. The simulations consider two possible network configurations: one in which the interbank links in the network are distributed roughly evenly across different banks (Poisson configuration), and one in which some banks in the network are much more highly connected than the typical bank (geometric configuration). The simulations for the Poisson network are a useful benchmark, while the fat-tailed geometric configuration is more in keeping with real-world networks.

We consider six experiments and a further four policy exercises. First, to provide a baseline, we consider what happens when a random adverse haircut shock at a single bank forces it to start hoarding liquidity under the Poisson network

² A repo transaction entails borrowing money using securities as collateral. It is structured as the spot sale of a security for cash, coupled with an agreement to repurchase the same security at the initial price plus interest at a particular date in the future. When the cash lent on repo trades is lower than the current market value of the security used as collateral, the discount is referred to as the haircut.

³ See Jackson (2008) for a comprehensive study of network techniques and their applications to economics.

⁴ We build on Gai and Kapadia (2010, 2011) who model default contagion and liquidity hoarding by adapting techniques advanced by Newman et al. (2001) and Watts (2002). May and Arinaminpathy (2010) show how similar results on the likelihood and dynamics of default contagion can be obtained by assuming that all banks follow exactly 'average' behaviour, and we also draw on them in applying this type of (mean-field) approximation to our model.

configuration. We illustrate the frequency and contagious spread of systemic liquidity crises, identifying the tipping point of the system in the process. Our second experiment introduces an additional aggregate haircut shock which affects all banks. We use it to articulate how our model speaks to the collapse of interbank markets in the early part of the crisis during August and September 2007.

The third and fourth experiments focus on the role of concentration by assessing how the results change under a fat-tailed (geometric) network configuration, and also explore the differing consequences of a ‘targeted’ shock which affects the most interconnected interbank lender under both network configurations. These are old exercises in epidemiology and the analysis of network resilience (Anderson and May, 1991; Albert et al., 2000). But the rapid emergence of institutions that are too big, connected, or important to fail underlines their increasing importance in finance.

The fifth and sixth experiments explore different dimensions of complexity—what happens to financial vulnerability when unsecured interbank market activity increases and initial aggregate haircuts change. Since haircuts appear to exhibit cyclical behaviour, tending to be compressed in the upswing of a cycle as financial institutions become increasingly exuberant, the latter experiment speaks to how systemic risk in the system may change dynamically.

Our first policy exercise explores the consequences of imposing a uniform increase in liquid asset holdings on the resilience of the financial system. The second policy exercise then simulates the model under an alternative policy rule in which the *average* increase in liquid assets is identical to the first policy exercise but where the increase in liquid assets at each individual bank is positively related to its interbank assets. This allows us to assess the impact of targeting higher liquidity requirements on key players in the interbank network.

The third policy exercise analyses how imposing haircut-dependent liquidity requirements might affect the cyclicity of systemic risk. Finally, we consider what happens if the amplification mechanism underpinning the liquidity hoarding dynamic is attenuated somewhat. Since this amplification may be associated with uncertainty about the network among participants within it, this exercise speaks to the impact that greater network transparency could have on financial system resilience.

The parsimony of our framework allows us to incorporate insights from the recent literature on financial crises. It draws on more conventional theoretical and empirical approaches that view the financial system as a network.⁵ But it is able to assess resilience and policy interventions under a much more realistic network structure, while placing liquidity effects and funding contagion at centre stage and still maintaining the generality of an analytical model. Our model also embodies the amplification role of collateral in recognising that a key risk in repo transactions stems from shocks to the haircuts on the securities which serve as collateral, and complements the literature on the behavioural foundations of the interbank freeze by articulating how a strong focus on network contagion effects may help to improve understanding of the probability and impact of such events.⁶

The paper proceeds as follows. Section 2 presents some empirical motivation for a network approach to the study of systemic risk. Section 3 outlines the model and describes analytically how contagion can spread across a banking network. Section 4 presents the results of our six experiments and four policy exercises. Section 5 draws out the policy implications and discusses them in the context of the ongoing debate on regulatory reform. A final section concludes.

2. Stylised evidence on financial system complexity and concentration

Quantifying the complexity and concentration of financial systems is difficult. Nevertheless, data on interbank lending and intra-financial system activity provide some motivation for adopting a network approach, and for our focus on bank heterogeneity, liquidity hoarding dynamics, and non-linearities.

Fig. 1 depicts the network of bilateral large exposures between the major UK banks, with the nodes representing banks, their size representing each bank’s overall importance in the interbank network, and the darkness of the links reflecting the value of interbank exposures between institutions. The evident complexity highlights the need to move beyond stylised network representations when analysing systemic risk. Concentration in the network is also clearly visible and this is reflected in the fat-tailed nature of the underlying distribution of linkages and loan sizes implied by Fig. 1.⁷

Evidence on the recent evolution of concentration and complexity motivates our focus on these issues in relation to the interbank market collapse. Fig. 2 points to growing concentration within national financial boundaries, from already high

⁵ More conventional theoretical models include Allen and Gale (2000), Freixas et al. (2000) and Brusco and Castiglionesi (2007)—see Allen and Babus (2009) for a comprehensive survey. The empirical literature is surveyed by Upper (2011). This tends to use regulatory data on large exposures between banks to analyse default contagion in interbank markets—apart from a brief discussion by Furfine (2003), the role of liquidity effects in the contagion process is largely ignored in this line of the literature.

⁶ On the amplification role of haircut shocks, see Brunnermeier and Pedersen (2009), Adrian and Shin (2010a), Geanakoplos (2010) and Gorton and Metrick (2010). Key papers on the interbank freeze include Caballero and Krishnamurthy (2008), Allen et al. (2009), Caballero and Simsek (2010), Bolton et al. (2011), Diamond and Rajan (in press) and Acharya et al. (in press). Acharya and Skeie (this issue) is another contribution in this vein. Their model explores how banks’ uncertainty over their ability to roll over their own debt may cause them to restrict lending in interbank markets. Although they do not consider how such behaviour may propagate through the interbank network, their story of a precautionary motive for liquidity hoarding is consistent with the view of hoarding taken in this paper; indeed, it may be interpreted as providing a behavioural foundation for it.

⁷ See Boss et al. (2004) and Craig and von Peter (2010) for similar evidence on the Austrian and German interbank markets, respectively. Properties of cross-border banking networks are discussed by von Peter (2007), Kubelec and Sa (2010) and Garratt et al. (2011).

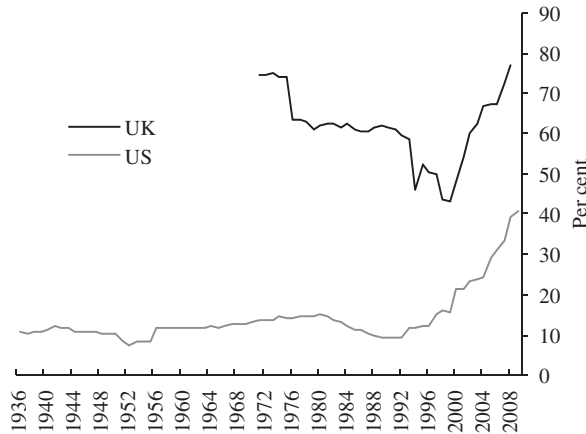


Fig. 2. Concentration of the UK and US banking systems.
 Sources: FDIC and Bank of England calculations. (a) Largest three banks by total assets as a percentage of total banking sector assets. (b) Data is to January 2009.

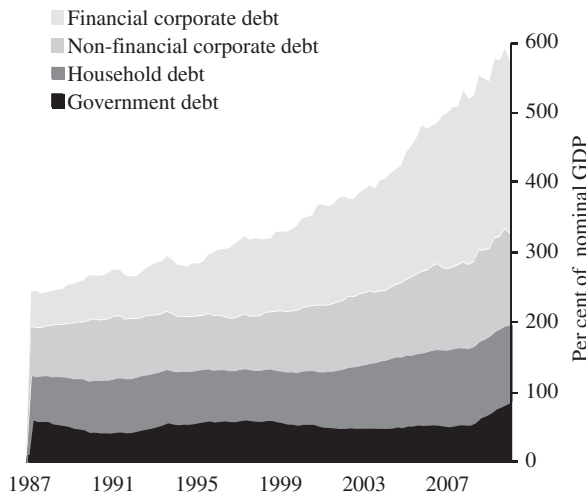


Fig. 3. Breakdown of UK Debt.
 Sources: ONS Blue Book and Bank of England calculations.

starting points. It shows the marked rise in the combined assets of the three largest banks by total assets as a percentage of total banking system assets in both the UK and US.⁸

Financial system complexity is likely to go hand-in-hand with intra-financial system activity which tends to increase the length of credit chains. Shin (2009) notes how the advent of securitisation markedly increased the complexity of the financial system by lengthening the intermediation chain in the lead up to the global financial crisis. In many cases, the same security was used repeatedly in repo markets, with the lender using the security received as collateral to borrow from others. As noted above, these transactions are subject to amplifying dynamics and cyclical fluctuations linked to the variability of collateral haircuts over the credit cycle. Fluctuations of this kind are illustrated by the dramatic rise and subsequent fall in the stock of repos and financial commercial paper as a percentage of broad money in the US.⁹ And, as Fig. 3 shows, the growth in intra-financial activity extended well beyond banks to other intermediaries, with financial corporate debt (including banks and non-banks) accounting for some two-thirds of the total growth in UK debt between 2003 and 2007.

Fig. 4 illustrates the sudden and sharp rise in the cost of unsecured interbank borrowing that followed the onset of difficulties at some institutions in August 2007 and, again in September 2008, following the collapse of Lehman Brothers. Precautionary hoarding by banks and growing counterparty risks prompted a freeze in interbank borrowing. Such a breakdown of the interbank market in the US, UK and Europe was an unprecedented event and the ensuing contagion

⁸ King (2010) also provides evidence on this.
⁹ See Fig. 10 in Adrian and Shin (2010b). For evidence of cyclicity in the underlying haircuts, see CGFS (2010), Geanakoplos (2010), Gorton and Metrick (2010), and Krishnamurthy (2010).

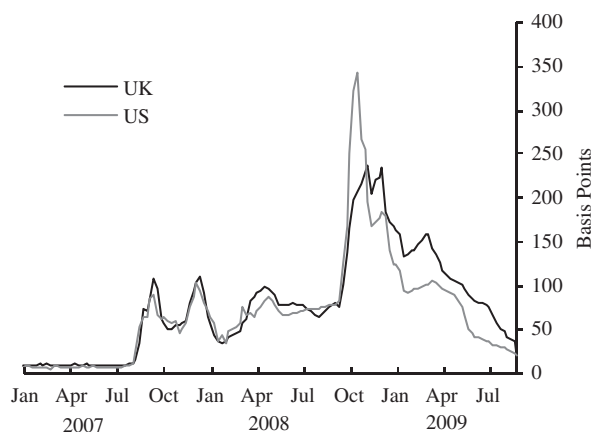


Fig. 4. Three-month interbank rates relative to expected policy rates.

Sources: Bloomberg and Bank of England calculations. (a) Spread of three-month Libor to three-month overnight index swap (OIS) rates. (b) Five-day moving average.

placed considerable funding pressure on banks in other jurisdictions. Prices alone do not tell the full story. The quantity of funding available, especially at maturities greater than overnight, declined dramatically. The counterpart to this was the sharp increase in banks' holdings of reserves with central banks. This meant that national financial systems effectively collapsed into a star network with central banks at their centre.

3. A network model of the banking system

In broad terms, our model explores the resilience of the financial system to liquidity shocks affecting a subset of banks under different network configurations, degrees of connectivity between financial institutions, haircut assumptions, and balance sheet characteristics. We start by describing how the network of interbank exposures and balance sheets are constructed, before discussing how shocks to haircuts which trigger liquidity hoarding at some institutions may potentially propagate across the system. We then provide some intuition for the tipping points that emerge in the simulations in Section 4 by applying a simplifying assumption to the model. This allows us to derive an explicit condition which identifies whether or not the system is vulnerable to a systemic liquidity crisis and how this is affected by the parameters of the model, including balance sheet characteristics and initial haircuts.

Throughout our analysis, the only sources of randomness in our model relate to the structure of the network, which is drawn from an exogenously set distribution, and the initial idiosyncratic shocks that trigger liquidity hoarding, which randomly affect any bank in the system. As the network fully determines the pattern of interbank linkages and because we also fix each bank's total interbank liability position exogenously, the value of each individual interbank linkage is determined endogenously, and so therefore is each bank's total interbank asset position. So randomness in the network structure maps into randomness in each institution's total interbank asset position.

Otherwise, the model is entirely non-stochastic: once the network and an initial shock have been drawn, the propagation of contagion is a purely deterministic process which depends on the parameters of the model, all of which are exogenous. Given these parameters, the set-up seeks to explain whether or not an initial shock generates more widespread liquidity hoarding and, if so, what fraction of the system is affected by outbreaks of contagion. But despite the deterministic nature of the contagion process, the randomness over the network and initiating shock imply uncertainty over whether contagion will occur from an *ex ante* perspective. So we also explore how the frequency of contagion is affected by changes in both the distributional assumptions that drive the network structure and in the nature of the initiating shocks.

3.1. The financial network

Table 1 provides a summary of all notation used in the paper. The financial network consists of n financial intermediaries, 'banks' for short, which are linked together randomly by their unsecured claims on one another. Each bank is represented by a node on the network, and the bilateral unsecured 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. The number of individual interbank exposures varies across banks. We suppose that each bank has j_i lending links representing its unsecured interbank assets (i.e. money lent to counterparties by the bank) and k_i borrowing links representing its unsecured interbank liabilities (i.e. money borrowed by the bank from counterparties). Since every interbank lending link for one bank is a borrowing link for another bank, the *average* number of lending links across all

Table 1
Description of parameters and calibration in baseline simulation.

Parameter	Description	Baseline calibration
n	Number of banks	250
j_i	Number of bilateral unsecured interbank lending links for bank i	Endogenous (depending on network)
k_i	Number of bilateral unsecured interbank borrowing links for bank i	Endogenous (depending on network)
z	Average degree or connectivity	Varies
L_i^{IB}	Unsecured interbank liabilities	15% of balance sheet
L_i^R	Repo liabilities (i.e. borrowing secured with collateral)	20% of balance sheet
L_i^D	Retail deposits	Endogenous (balancing item)
K_i	Capital	4% of balance sheet
L_i^N	New unsecured interbank borrowing raised after a shock	0
A_i^{IB}	Unsecured interbank assets	Endogenous (depending on network)
A_i^F	Fixed assets (e.g. individual corporate loans or mortgages)	Endogenous (depending on A_i^{IB})
A_i^C	'Collateral' assets (assets which may be used as collateral in repo transactions)	10% of balance sheet
A_i^{RR}	Reverse repo assets (i.e. collateralised lending)	11% of balance sheet
A_i^L	Unencumbered fully liquid assets	2% of balance sheet
h	Aggregate haircut applied to collateral used to obtain repo funding	0.1
h_i	Bank-specific haircut applied to collateral used to obtain repo funding	0
$-\varepsilon_i$	Idiosyncratic liquidity shock	0
μ_i	Fraction of banks linked to bank i which hoard (withdraw deposits from bank i)	na
λ	Proportion of deposits withdrawn by hoarding banks	1

banks in the network must equal the average number of borrowing links. We refer to this quantity as the *average degree or connectivity* of the system, and denote it by z .

The structure of bilateral claims and obligations in the financial network, as defined by the joint distribution of lending and borrowing links and its moments (including the average degree), plays a key role in determining how shocks spread through the network. Our general modelling framework applies for any arbitrary choice of joint distribution, though we adopt specific assumptions both when analysing the theoretical approximation to our model and in our numerical simulations.¹⁰ This means that our results are able to encompass all possible network structures, including any implied by a prior optimal network formation game and real-world networks like the UK interbank network depicted in Fig. 1.

3.2. Balance sheets and repo haircuts

Fig. 5 presents the composition of individual bank balance sheets in the model. The liabilities of each bank i are composed of unsecured interbank liabilities, L_i^{IB} ; repo liabilities (i.e. borrowing secured with collateral), L_i^R ; retail deposits, L_i^D ; and capital, K_i . We assume that the total unsecured interbank liability position of every bank is evenly distributed over each of its borrowing links and is independent of the number of links the bank has.¹¹

Since every interbank liability is another bank's asset, unsecured interbank assets at each bank, A_i^{IB} , are endogenously determined by the network links. So, although total unsecured interbank assets equal total unsecured interbank liabilities in aggregate across the network, each individual bank can have a surplus or deficit in their individual unsecured interbank position. Additionally, banks hold four further asset classes: fixed assets (e.g. individual corporate loans or mortgages), A_i^F ; assets which may be used as collateral in repo transactions (collateral assets), A_i^C ; reverse repo assets (i.e. collateralised lending), A_i^{RR} ; and fully liquid assets (e.g. cash, central bank reserves, high-quality government bonds), A_i^L .

In our model, fully liquid assets can always be used as collateral to obtain repo financing if required without any haircut (or alternatively sold without any price discount), i.e. borrowing can be obtained against the full value of the asset. On the other hand, we assume that fixed assets and unsecured interbank assets can never be used as collateral in repo transactions.

The aggregate haircut associated with using collateral assets to obtain repo funding is denoted by $h \in [0, 1]$. This haircut reflects the perceived underlying risk of the collateral. It is designed to protect lenders against losses that they may incur when trying to sell collateral in the event that they are left with it due to counterparty default. Such losses may be caused by market illiquidity, changing degrees of asymmetric information, or increases in the probability of default on the underlying security over the duration of the loan. So changes in the likelihood of any of these factors may affect aggregate haircuts.¹²

We further allow for the possibility of a bank-specific haircut, h_i , so that the maximum amount of repo funding that can be obtained from collateral assets is given by $(1-h-h_i)A_i^C$. This idiosyncratic haircut might reflect the greater default probability of a particular bank—if the lender perceives that there is a higher chance it will fail, then it might demand

¹⁰ The framework also assumes that there is no statistical tendency for highly connected banks to be either more or less likely to be linked to other highly connected banks or to poorly connected banks.

¹¹ If a bank has no borrowing links, $L_i^{IB} = 0$ for that bank.

¹² For a more detailed explanation of why repo collateral involves haircuts and why these haircuts might fluctuate, see Gorton and Metrick (2010).

Assets	Liabilities
Fixed Assets (A^F)	Retail Deposits (L^D)
‘Collateral’ Assets (A^C)	Repo (L^R)
Reverse Repo (A^{RR})	Unsecured Interbank Liabilities (L^{IB})
Unsecured Interbank Assets (A^{IB})	Capital (K)
Liquid Assets (A^L)	

Fig. 5. Stylised balance sheet.

a higher haircut as extra protection both because it is more likely to end up with the collateral in practice, and because there may be some legal risk in accessing the collateral in a timely fashion.

Reverse repo transactions are secured with collateral that commands the same aggregate haircut as on A_i^C . This implies that the amount of collateral that bank i receives on its reverse repo assets is given by $A_i^{RR}/(1-h)$. We allow for this collateral to be fully rehypothecated (i.e. reused in another unrelated transaction) to obtain repo funding with the same aggregate haircut, h . The maximum amount of repo funding that can be obtained from rehypothecating collateral obtained in reverse repo transactions is then given by $[(1-h-h_i)/(1-h)]A_i^{RR}$.

3.3. Liquidity shortages, replenishment and propagation

Given the balance sheet and haircut assumptions above, we now seek the condition under which bank i will remain liquid in each period. To simplify our analysis, we preclude the possibility of systematic retail deposit inflows or outflows and assume that banks cannot raise fresh equity. We also assume that the central bank never takes collateral at more generous terms than the market in its liquidity operations.

A bank is liquid if the total amount of collateral it has available to obtain repo funding (which includes its fully liquid assets) plus any new unsecured interbank borrowing, L_i^N , is sufficient to exceed the amount of existing repo funding it has and meet any idiosyncratic liquidity shock, $-\varepsilon_i$, or loss of interbank funding that it might experience. Liquidity shocks or shocks to aggregate or idiosyncratic haircuts thus have the potential to trigger a liquidity shortage at the bank.

If a bank faces such a liquidity shortage, it needs to take defensive action to avoid defaulting on required payments. In our model, we assume that the bank tries in the first instance to raise any resources needed by withdrawing (or, equivalently, refusing to roll over) unsecured interbank assets, A_i^{IB} , from counterparties it was lending to in the interbank network, i.e. it hoards liquidity.¹³ Clearly, there are other avenues open to the bank. For example, it may try to liquidate its fixed assets, A_i^F , in a fire sale. Alternatively, it might seek to raise the interest rate it is prepared to offer on new interbank borrowing.¹⁴ As the recent crisis demonstrates, however, liquidity hoarding was widely observed, whereas there was less evidence of widespread fire sales or of banks ‘paying up’ significantly in interbank funding markets. This is largely because banks typically viewed such actions as unattractive, last resort measures, given their large direct costs in terms of profitability and important adverse stigma effects due to their high visibility. By contrast, hoarding liquidity has fewer direct costs and since unsecured interbank transactions are over-the-counter, adverse stigma effects can be kept to a minimum. Indeed, in contrast to other options, the bank may not need to make any active decision if it chooses not to roll over interbank loans.¹⁵

¹³ Note that hoarding in this context is driven purely by the bank’s own liquidity needs; concerns over the solvency of its counterparty play no role in its decision. In the model presented here, we assume that such hoarding behaviour represents a genuine drain on the liquidity of the entire banking system, for example, ending up as increased reserves at central banks. As discussed by Gai and Kapadia (2011) in a similar model which abstracts from repo market activity, hoarding behaviour can also be interpreted as a switch from lending at long maturities (e.g. for three months or one year) to lending at much shorter maturities (e.g. overnight).

¹⁴ The bank could also try to cut lending to the real economy, A_i^F . While this was certainly observed during the crisis, we abstract from this possibility as it is only likely to generate liquidity slowly. But the potential for this type of credit crunch reaction is one of the key channels through which systemic liquidity crises can become so costly for society.

¹⁵ Diamond and Rajan (in press) observe that banks may be hesitant to enter into a fire sale because the alternative of holding on to the asset may be more beneficial. Adverse selection in the spirit of Stiglitz and Weiss (1981) points to the fact that banks may be unwilling to ‘pay up’ significantly in

Once we recognise liquidity hoarding, network effects take centre stage. In particular, liquidity shortages can propagate through the system via the network of interbank linkages. Suppose that a fraction, μ_i , of banks connected to bank i in the network ‘hoard’ liquidity from it, withdrawing a portion of their deposits held at bank i . Further, let us suppose that, on average, these hoarding banks withdraw a fraction λ of the deposits that they hold at bank i . Under these assumptions, bank i loses $\lambda\mu_i L_i^B$ of its liabilities due to liquidity hoarding by its counterparties in the network. Therefore, taking into account the potential for lost deposits due to liquidity hoarding and the haircut assumptions, its overall liquidity condition can be formally expressed as

$$A_i^L + (1-h-h_i)A_i^C + \frac{(1-h-h_i)}{(1-h)}A_i^{RR} + L_i^N - L_i^R - \lambda\mu_i L_i^B - \varepsilon_i > 0 \quad (1)$$

where the first four terms represent its available liquidity and the last three terms represent funding it needs to cover with collateral, and funding outflows.

The value of λ is a key determinant of the strength of amplification in the model—in particular, the higher the value of λ , the larger the shocks that can hit banks further down the chain of contagion. The extreme case of $\lambda = 1$ corresponds to full withdrawal, with lending banks withdrawing their entire deposit irrespective of their own liquidity ‘shortfall’ (their outstanding shortage of liquidity after all collateral and liquid assets have been used). At the other extreme, if banks’ liquidity shortfalls were the only determinant of the amount of hoarding, then λ would be fully endogenous within the model.

In practice, λ is likely to lie between these two extremes. In our baseline numerical simulations, however, we set $\lambda = 1$ to generate the sharpest possible results. Immediate full withdrawal may seem unlikely because contractual obligations may prevent banks from withdrawing their entire deposit straight away. On the other hand, if a bank has a liquidity shortfall because it has lost some portion of its deposits from a hoarding counterparty, it may consider it to be only a matter of time before the full amount is lost, for example when contractual obligations expire. So, 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 it suffering liquidity problems in the future. As such, assuming $\lambda = 1$ effectively captures a rich set of dynamics which may operate through forward-looking expectations.

The propagation of contagion also depends on the extent to which withdrawals are concentrated on particular counterparties or whether they are more evenly distributed. We assume the latter: banks raise any resources needed by withdrawing funding equally and proportionately from *all* of their counterparties. This assumption seems plausible as immediately accessible deposits are only likely to be available in relatively small amounts from each counterparty.

3.4. Contagion dynamics and ‘tipping points’

Eq. (1) makes clear that the decision by a single bank to hoard liquidity makes it harder for banks that were previously borrowing from it to meet their own liquidity condition without resorting to hoarding themselves. In particular, as each successive bank suffers a liquidity shortfall, its hoarding has the potential to trigger a liquidity shortage at other banks to which it is connected by interbank lending. This process will only die out if either no neighbours to newly distressed banks become distressed themselves or when every bank in the network is in distress. So hoarding can potentially spread across the system, with the structure and connectivity of the unsecured interbank network playing a critical role in determining the evolution of contagion.¹⁶

To illustrate the dynamics of contagion, recall that bank i has k_i borrowing links. So if a single counterparty to bank i hoards, $\mu_i = 1/k_i$ since interbank liabilities are evenly distributed across counterparties. Suppose we randomly perturb the network by assuming that a single bank suffers a haircut or idiosyncratic liquidity shock which is sufficiently large to cause it to start hoarding liquidity. Then, by substituting for μ_i in Eq. (1) and rearranging, we can see that for contagion to spread beyond the first bank, there must be at least one neighbouring bank for which

$$\frac{A_i^L + (1-h-h_i)A_i^C + [(1-h-h_i)/(1-h)]A_i^{RR} + L_i^N - L_i^R - \varepsilon_i}{\lambda L_i^B} < \frac{1}{k_i} \quad (2)$$

If this condition holds, then contagion starts to spread. In particular, a second bank is forced into hoarding liquidity which may, in turn, create liquidity shortfalls at other banks in the network, and so on. And the broader contagion dynamics can obviously also entail the possibility of banks being exposed to multiple hoarding counterparties, in which case similar equations derived from (1) determine the spread of liquidity hoarding across the network.

(footnote continued)

interbank markets as there is a risk that wholesale lenders may view it as a signal of underlying difficulties. Acharya and Skeie (this issue) provide behavioural foundations for why banks may choose to hoard liquidity for precautionary reasons.

¹⁶ As pointed out by Lo (this issue), a bank’s liquidity trigger is unlikely to be as deterministic as (1). His suggestion to introduce a probabilistic liquidity trigger, whereby the probability of needing to take defensive action increases as the liquidity position of the bank deteriorates, would certainly enrich the framework. In particular, it would introduce an explicit role for uncertainty and would also allow for heightened volatility to trigger liquidity events. But the main effect would be to introduce a greater dispersion in outcomes for given underlying parameters rather than fundamentally altering our main results.

Before simulating this process, we can gain further intuitive insight into the nature of contagion by making some further simplifying assumptions which allow us to obtain clean analytical results. Specifically, rather than taking the network to be randomly generated, let us suppose that each bank is connected to exactly z other banks as both a lender and borrower (which implies that $j_i = k_i = z$ for all banks).¹⁷ Let us also suppose that there are no idiosyncratic haircuts or shocks (so $h_i = \varepsilon_i = 0$) and that all banks have identical balance sheets, allowing us to drop all i subscripts. Given that every interbank asset is another bank's interbank liability, these assumptions also imply that $L^{IB} = A^{IB}$ for all banks. Then, assuming full withdrawal ($\lambda = 1$) and the absence of any new unsecured interbank funding ($L_i^N = 0$), we can rewrite (2) as

$$z < \frac{A^{IB}}{A^L + (1-h)A^C + A^{RR} - L^R} \quad (3)$$

This expression is identical for every bank in the network under these stark assumptions. It yields a 'tipping point' condition, determining when contagion may break out across the financial network. In particular, if (3) is satisfied, then provided that z is greater than or equal to 1 so that there is sufficient network connectivity, any liquidity hoarding by a single bank will cause all neighbouring banks to become distressed and start hoarding. Because neighbours of neighbours face the same liquidity condition, hoarding behaviour will cascade through the entire network. By contrast, if (3) is violated, an initial case of liquidity hoarding by one bank will have no systemic consequences at all. As such, the 'tipping point' condition clearly illustrates how very small changes in the underlying parameters of the model can lead to fundamentally different outcomes.¹⁸

Eq. (3) also clarifies the conditions under which systemic liquidity crises may occur. In particular, it highlights the way in which low liquid asset holdings, large adverse aggregate haircut shocks, a high amount of repo borrowing, and a high level of unsecured interbank lending are all likely to increase the susceptibility of the system to a widespread liquidity crisis. Our simulations demonstrate that these results are borne out under more general assumptions.

4. Model simulations

We now draw upon numerical simulations to offer further insight into the role of concentration and complexity in contributing to systemic liquidity crises, and to assess a range of possible financial stability policy interventions.

Table 1 summarises the baseline simulation parameters. The system comprises 250 banks. In the initial step in each realisation, the network of unsecured interbank linkages between banks is drawn randomly from an underlying distribution which is assumed to characterise the structure of the network. In what follows, we focus on two characteristic network structures: one in which the links connecting banks are distributed roughly uniformly (Poisson); and one in which the underlying distribution describing linkages results in some banks in the network being much more highly connected than the typical bank (geometric). And when drawing the network, we allow for the possibility that two banks can be linked to each other via both lending and borrowing links—no netting of exposures is assumed.

Although the model applies to fully heterogeneous banks, for the purpose of illustration we take the liability side of the balance sheet of all banks to be identically composed of a capital buffer of 4% of the balance sheet, unsecured interbank liabilities (15%), repo liabilities (determined as described below) and retail deposits. Since each bank's interbank liabilities are evenly distributed over each of their borrowing links, interbank assets are determined endogenously within the (random) network structure and thus vary across banks. In addition to interbank assets, the bank has liquid assets (2% of the balance sheet), reverse repos and collateral assets. The asset side is further 'topped up' by fixed assets until the total asset position equals the total liability position.

For the repo aspects of the balance sheet, we assume that, in the initial state, all collateral assets and assets received as part of reverse repo transactions are used as collateral to obtain repo funding so that

$$L_i^R = (1-h-h_i)A_i^C + \frac{(1-h-h_i)}{(1-h)}A_i^{RR} \quad (4)$$

We suppose that reverse repo assets are 11% of the balance sheet, collateral assets are 10%, the initial aggregate haircut, h , is 0.1, and there are no idiosyncratic haircuts. This implies that repo liabilities comprise 20% of the balance sheet initially.

To keep the number of experiments manageable, we assume throughout that banks can never raise any new deposits in the unsecured interbank market (i.e. $L_i^N = 0$ for all banks).¹⁹ Apart from in one of the policy experiments, full withdrawal by hoarding banks ($\lambda = 1$) is assumed throughout.

In most of what follows, we vary the average connectivity between banks, z , drawing 1000 realisations of the network for each value, and then shock the network in different ways according to the experiment in question. The dynamics of contagion

¹⁷ In formal terms, this represents a mean-field approximation to the model.

¹⁸ Gai and Kapadia (2011) show that although a precise threshold in the form of Eq. (3) cannot normally be identified, it is possible to prove the existence of such 'tipping points' (or, more formally, phase transitions) in this class of model without taking a mean-field approximation and while still maintaining heterogeneous balance sheets (though the $\lambda = 1$ assumption is still necessary in their formal analysis).

¹⁹ As Gai and Kapadia (2011) show, although allowing for replacement of lost interbank deposits can significantly reduce the likelihood of systemic liquidity crises, it generally does not otherwise change the fundamental properties of the model and its funding contagion dynamic—in particular, the tipping point is still present and contagion can still spread to the entire system on some occasions.

follow the process described in Section 3.3, with Eq. (1) at the centre of the propagation dynamics. For each realisation, we follow these dynamics iteratively until no new banks are forced into hoarding liquidity or until every bank is hoarding. We count as ‘systemic’ those episodes in which at least 10% of banks are forced into hoarding liquidity. Our results identify the frequency of systemic liquidity crises and their impact in terms of the average fraction of the system affected in each systemic outbreak (i.e. how widely contagion spreads, conditional on it spreading to at least 10% of banks in the system).

4.1. Experiment 1: a stylised systemic liquidity crisis

In our first experiment, we assume that the links in the network are spread roughly uniformly, with each possible directed link in the network being present with independent probability p (a Poisson network). We hold aggregate haircuts constant but shock the system by assuming that a single bank receives a very large adverse idiosyncratic haircut shock which causes it to start hoarding liquidity.

Fig. 6 (baseline) presents the results. Contagion occurs for values of z between 0 and 20 and its probability is non-monotonic in connectivity, at first increasing before falling. But when contagion breaks out, it invariably spreads to the entire network.

These results accord well with the analytical approximation at the end of Section 3.4, which is not too surprising because this approximation is most reasonable for a Poisson network. Specifically, given the parameters chosen in the baseline simulation, Eq. (3) suggests that contagion will occur for $z < 7.5$ but not for $z > 7.5$. From Fig. 6 (baseline), it is evident that $z=7.5$ is the point around which the probability of contagion starts to fall from close to one. The reason it remains positive for higher values of z is due to the randomness of the network structure which means that contagion can still break out under certain configurations. And the reason contagion is not always certain for smaller values of z is that the initial shock may hit a bank which either has no interbank assets, and is therefore unable to trigger any contagion by hoarding liquidity, or is in an isolated subset of the network.

4.2. Experiment 2: adding aggregate haircut shocks

Our second experiment repeats the first but also incorporates an aggregate haircut shock that increases the haircut from 0.1 to 0.2. This acts like a bank run in draining liquid assets from every repo borrower in the system as they are all required to post more collateral to obtain the same amount of repo funding. It is clear from (3) that this shifts the tipping point to around $z=15$. This is borne out by the results presented in Fig. 6 (with aggregate haircut shock).

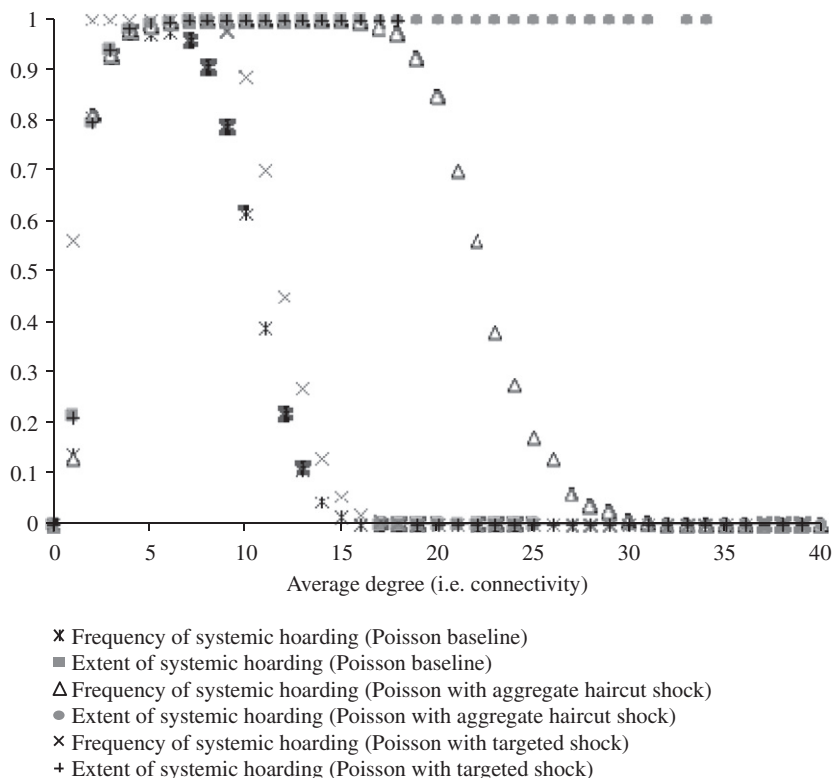


Fig. 6. Systemic liquidity hoarding (Poisson network: single random idiosyncratic haircut shock; aggregate haircut shock and single random idiosyncratic haircut shock; and single haircut shock targeted to bank with most interbank lending links).

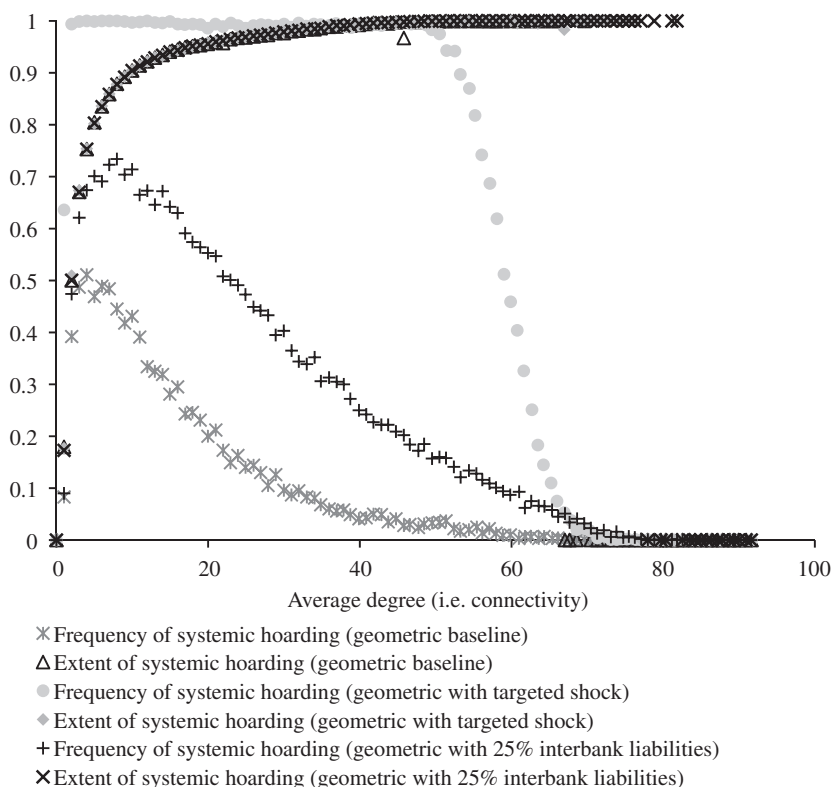


Fig. 7. Systemic liquidity hoarding (geometric network: single random idiosyncratic haircut shock; single haircut shock targeted to bank with most interbank lending links; single random idiosyncratic haircut shock with 25% unsecured interbank liabilities).

In some respects, this type of experiment mimics the behaviour of interbank markets in the early part of the crisis during August and September 2007. In light of both bad news and greater uncertainty on subprime mortgages and other types of collateral which were being used to back repo and other secured forms of funding, aggregate haircuts increased. As a result, some banks found themselves short of liquidity. This was especially true of banks which were forced to take back assets from off-balance sheet vehicles for which liquidity had dried up. In response to their own funding liquidity stress, some banks started to hoard in the unsecured interbank market.

As highlighted by Lo (this issue), uncertainty over the value of assets may also have exacerbated hoarding behaviour. The result was paralysis in the interbank market as reflected by the sharp increase in spreads depicted in Fig. 4. This will have increased counterparty risk which may have intensified problems for some institutions. Although this framework abstracts from the counterparty risk dimension, it makes clear how a seemingly small shock to a limited set of assets which are being used as collateral can lead to a collapse in both secured and unsecured interbank markets.²⁰

4.3. Experiment 3: systemic liquidity crises in a concentrated network

Real-world financial networks do not appear to be particularly uniform; instead they appear to exhibit fat tails, with a small number of key players who are very highly connected both in terms of the number of interbank relationships they have and in terms of the overall value of those relationships. This reflects the underlying concentration in the banking sector.

To explore the impact of concentration on our results, we repeat our first simulation exercise but draw the network from a (geometric) distribution which embeds fat tails by allowing some banks to have substantially more connections than the typical bank.²¹ Fig. 7 (baseline) presents the results. Contagion is less likely and less severe for low values of z than under the Poisson distribution. This reflects the well-known result that fat-tailed networks tend to be more robust to random shocks (Anderson and May, 1991; Albert et al., 2000). On the other hand, it is also clear that contagion can occur

²⁰ In the model, hoarding entails withdrawal of funds while, in reality, much of the hoarding behaviour early in the crisis involved banks dramatically reducing the maturity of their lending. But, as noted above, the broad framework can be interpreted as speaking to this behaviour as well.

²¹ In the version implemented, we draw the number of borrowing and lending links separately from the same distribution, implying that there is no correlation between the number of counterparties a bank lends to and borrows from. To construct the network, we also need to ensure that the total number of borrowing links drawn equals the total number of lending links. We follow the algorithm outlined by Newman et al. (2001) to achieve this.

for much higher values of z , albeit rarely. So, for a broad range of connectivity, higher concentration in the network makes the system more susceptible to a systemic liquidity crisis.

4.4. Experiment 4: the impact of targeted shocks in concentrated and less concentrated networks

Experiment 3 masks what is perhaps the key difference between the fragility of concentrated and less concentrated networks. Thus far, we have assumed that when the initial idiosyncratic haircut shock occurs, it hits any bank in the network at random. Now suppose instead that the initial shock hits the bank with the largest number of unsecured interbank lending relationships. Figs. 6 and 7 (with targeted shock) show the results for each type of network.

When the shock is targeted at the most connected interbank lender, contagion occurs more frequently in both cases. But for the less concentrated (Poisson) network, a targeted shock only makes a relatively small difference to the results. By contrast, a targeted shock under the concentrated (geometric) network has catastrophic consequences, making contagion a near certainty for a very wide range of z . In the Poisson network, the most connected bank is not that much more connected than the typical bank. But under the fat-tailed (geometric) network, the most connected bank is likely to be connected to a very large portion of the other banks in the network, so if it becomes distressed, it has the potential to spread contagion very widely. This is again consistent with the results of Anderson and May (1991) and Albert et al. (2000), who both demonstrate how fat-tailed networks are particularly susceptible to shocks targeted at key participants.

It is clear that banks who are heavily involved in repo activity are more likely to face liquidity shortages from aggregate haircut shocks. So the most dangerous banks for the stability of the network are those which are both heavily involved in repo activity and big lenders in the unsecured interbank market—the former because it makes them highly susceptible to the initial shock, the latter because they propagate the shock widely. Because the banks that are heavily involved in repo activity are typically the same large, complex financial institutions that are big players in the unsecured interbank market, we can immediately see how the structure of the modern financial system may be particularly prone to systemic collapse. It also demonstrates why the seemingly small shocks of early/mid 2007 could have had such catastrophic systemic consequences as they affected banks that were both highly susceptible to liquidity disturbances and central to the structure of the network.

4.5. Experiment 5: the impact of greater complexity

Financial system complexity has no single definition and is difficult to measure. But it seems likely that, on any definition, it will be increased by intra-financial system activity, both in unsecured and secured markets. Therefore, in this experiment, we consider a random shock in a concentrated network (as in Experiment 3) but now suppose that unsecured interbank liabilities comprise 25% rather than 15% of the balance sheet. As can be seen in Fig. 7 (with 25% interbank liabilities), contagion occurs more frequently than in the baseline. Intuitively, this is because the overall rise in interbank liabilities increases the likelihood of larger funding withdrawals which cannot be absorbed by liquid assets.

4.6. Experiment 6: cyclical in haircuts and the likelihood of systemic liquidity crises

It is perhaps more interesting to consider complexity in a dynamic setting. If unsecured interbank lending were to increase over the economic cycle, then the system would become increasingly vulnerable to shocks. And compression in aggregate haircuts, which often occurs during the upswing of a cycle as the financial system becomes increasingly exuberant, may influence the amount of repo market activity by allowing more secured funding to be generated from a fixed amount of collateral.²² Rehypothecation of collateral may also increase in upswings, serving to increase the money multiplier and expand balance sheets. All of this may alter the vulnerability of the financial system over time.

We provide an indicative simulation of such time-varying risks in our model by varying the initial aggregate haircut from 0.25 to 0 (which, in turn, affects the initial amount of repo funding via Eq. (4)). We then assess the effect of this on the frequency of systemic crises in response to a combined shock in which the aggregate haircut jumps to a stressed level, taken to be 0.25, and a single bank suffers a sufficiently large idiosyncratic shock that it is forced to start hoarding. Fig. 8 (diamonds) presents the results for a concentrated network with average connectivity (z) fixed at 50. It is clear that risk in the system increases as the initial aggregate haircut falls, pointing towards the potential for systemic risk to increase as the credit cycle evolves. The non-linearity is due to the tipping point embedded in the model which implies that around critical thresholds, the system suddenly becomes much more vulnerable to collapse.

4.7. Policy Exercise 1: tougher liquidity requirements

We now turn to a series of policy experiments. Eq. (3) implies that an increase in liquid asset holdings should make the system directly less susceptible to systemic liquidity crises. So, in our first policy exercise, we repeat Experiment 3

²² For more on the links between haircuts and repo activity over the economic cycle, see Adrian and Shin (2010a) or CGFS (2010).

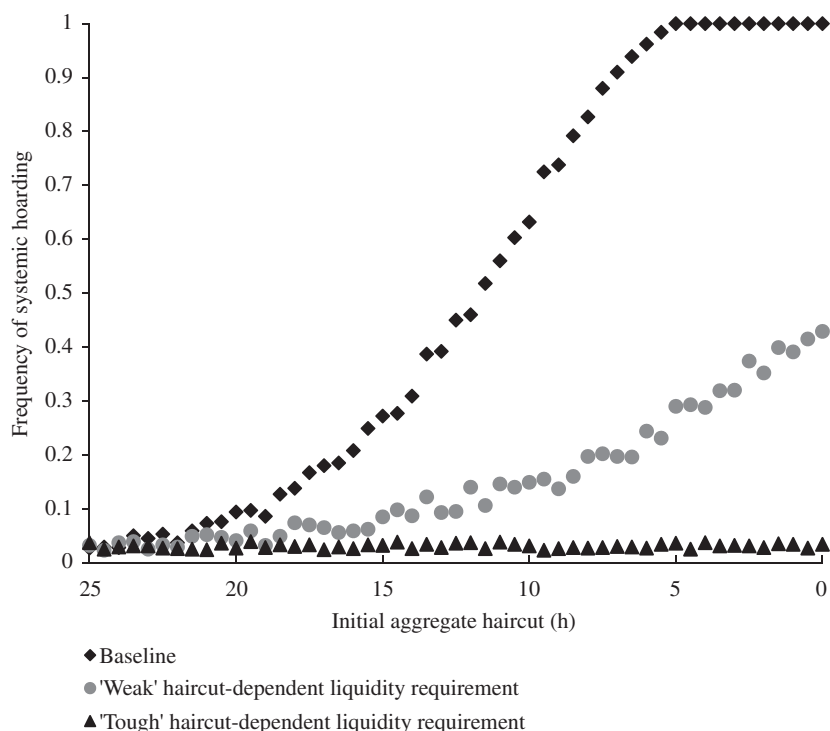


Fig. 8. Aggregate haircuts and the frequency of systemic liquidity hoarding (geometric network with average degree of 50: baseline; 'weak' haircut-dependent liquidity requirement; 'tough' haircut-dependent liquidity requirement).

(random idiosyncratic shock in a concentrated network) but suppose that 3.5% of assets are liquid rather than 2%. It is clear from Fig. 9 (with additional liquid assets) that this makes the system less prone to collapse.

4.8. Policy Exercise 2: systemic liquidity requirements

Experiment 4 revealed that concentrated networks were highly vulnerable to shocks to key players. Therefore, it is interesting to consider the effects of targeting higher liquidity requirements at such banks. So we simulate the model under a liquidity rule in which banks are required to hold a minimum of 2% liquid assets, plus an amount equal to 10% of their total interbank assets. Since interbank assets comprise 15% of total assets *on average*, this implies that the *average* liquid asset holding is 3.5%, thus making the experiment directly comparable to Policy Exercise 1. But banks with higher-than-average interbank assets will hold more liquid assets, while banks with lower-than-average interbank assets will hold less.

Fig. 9 (with systemic liquidity surcharge) presents the results. Comparing with the other experiment in Fig. 9, it is clear that the augmented liquidity rule is more effective in reducing the probability and spread of contagion than an equivalent across-the-board increase in liquid asset requirements. The intuition is simple—targeting liquidity requirements at the banks which are most instrumental in spreading contagion is more potent than requiring peripheral players to hold extra liquid assets.

4.9. Policy Exercise 3: haircut-dependent (time-varying) liquidity requirements

Experiment 6 demonstrated that the financial system might become increasingly vulnerable to collapse as aggregate haircuts fall. One possible policy response would be to compensate for this by linking liquidity requirements to aggregate haircuts—given the cyclical tendencies of haircuts, this policy can also be interpreted as a time-varying liquidity requirement. The solid circles and triangles in Fig. 8 illustrate the effects of amending Experiment 6 by introducing a 'tough' policy rule, in which liquid assets are required to rise from 2% of total assets at an aggregate haircut of 0.25 to 4.5% when haircuts are zero, and a 'weak' rule in which the liquid asset requirement only rises to 3.25%. In both cases, policy intervention works to offset the systemic risks created from lower aggregate haircuts by increasing resilience. Under the 'tough' policy rule, there is no increase in risk. Although not captured in the simulations, such policy rules may also reduce the incentives of banks to cut haircuts in the first place via effects operating through expectations.

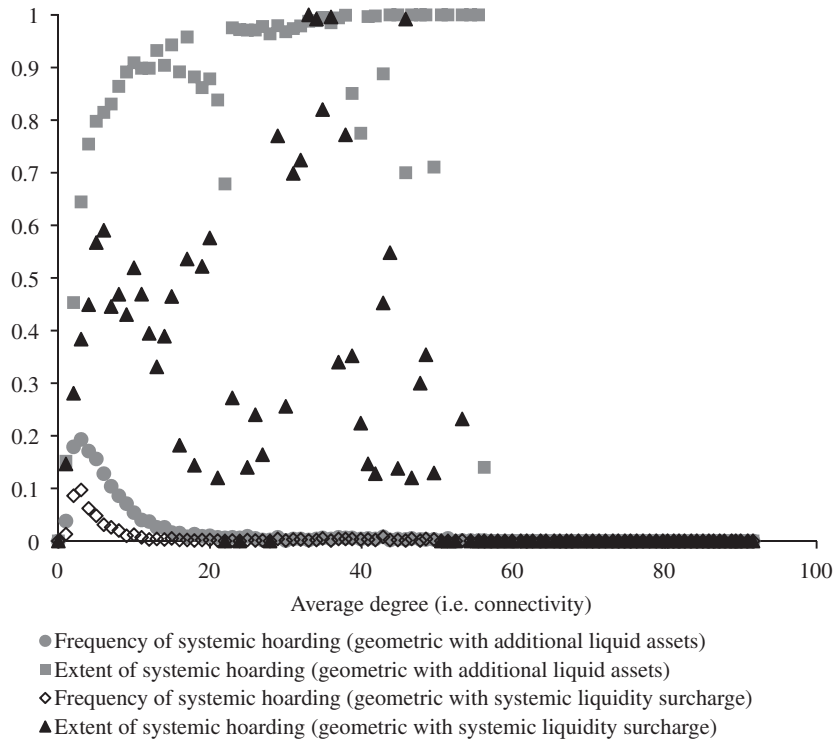


Fig. 9. Systemic liquidity hoarding (geometric network; single random idiosyncratic haircut shock with 3.5% liquid asset holdings; single random idiosyncratic haircut shock with 3.5% average liquid asset holdings but set via systemic liquidity surcharge rule).

4.10. Policy Exercise 4: greater network transparency

The assumption that hoarding banks fully withdraw their interbank deposits held with counterparties acts as a key amplifier in the model. Policy measures which help to promote greater network transparency could help to mitigate such incentives to hoard liquidity. In the context of the model, this would reduce the value of λ . To illustrate this, Fig. 10 presents results from a simulation in which, rather than withdrawing all their interbank lending when under liquidity stress, banks only withdraw an amount equal to their liquidity shortfall plus half of any remaining interbank assets. It is clear that this greatly diminishes the likelihood of collapse.

5. Policy implications

The literature on market failures in banking systems justifies the need for static (micro-prudential) regulation and dynamic (macro-prudential) policy intervention.²³ At the heart of the systemic collapses modelled in this paper is an underlying network externality: banks fail to internalise the consequences of their hoarding behaviour on others in the network.²⁴ The scale of this externality depends on the network and balance sheet structure of the financial system and behavioural assumptions such as the amount of liquidity hoarding when a bank is in distress (the value of λ). As the simulations involving targeted shocks make clear, these externalities vary across banks. Given these externalities, there is a case for public policy intervention, along a range of potential dimensions. These include:

(a) *Tougher micro-prudential liquidity regulation.* One immediate implication of our results is that banks need to have a larger stock of genuinely high-quality liquid assets than they would naturally choose given their own individual incentives. As noted above, it is clear from Eq. (3) that an increase in liquid asset holdings makes the system directly less susceptible to systemic liquidity crises. This is supported by the simulation results from the first policy exercise. But to the extent that greater liquid asset buffers also reduce the amount of repo activity and overall balance sheet size, they may also indirectly improve the resilience of the system by making banks less susceptible to haircut shocks. In Eq. (3), this

²³ See, for example, Allen and Gale (2004). For an overview of market failures which may justify a macro-prudential policy response, see Bank of England (2009).

²⁴ Note that the externality in our model differs from 'conventional' network externalities in economics which typically refer to the spillovers on users of particular services or products from the adoption of the service or product by others, and are often discussed in the context of telecommunications networks. But it is a network externality in the sense that the behaviour of individual participants can affect other participants via the links joining them together.

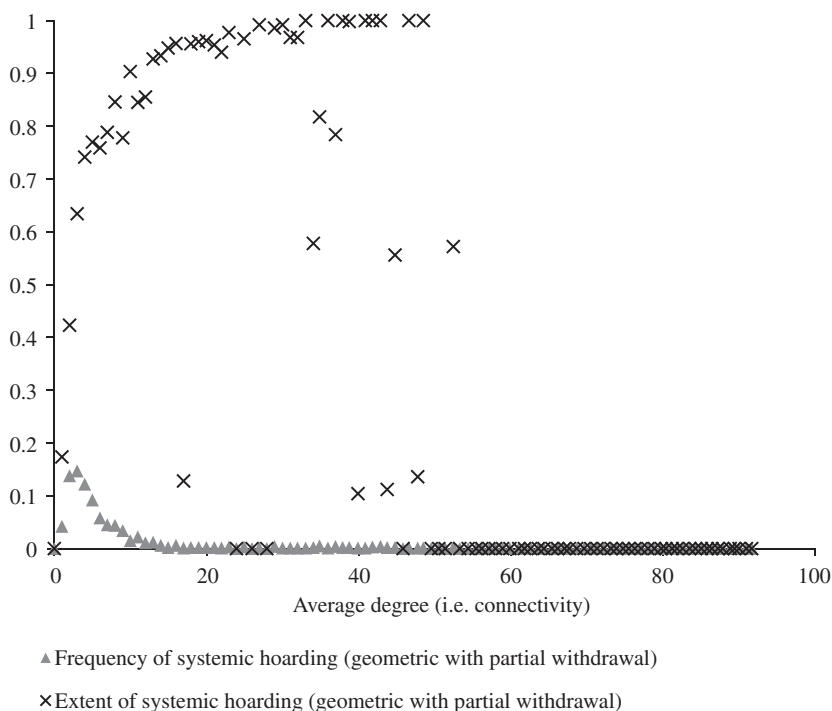


Fig. 10. Systemic liquidity hoarding (geometric network; single random idiosyncratic haircut shock; partial withdrawal).

would show up via a reduction in the amount of repo liabilities. So tougher liquidity requirements may have an additional *ex ante* benefit for system resilience by dampening the money multiplier.

The model takes liquid assets to be those which can either be sold without any price discount or used as collateral to obtain repo financing without any haircut. This clearly distinguishes them from ‘collateral assets’ which can only be used to obtain repo funding at a positive haircut. The model suggests that any asset which is likely to have a large haircut in times of stress is much less useful as a buffer against systemic liquidity crises. This points towards the importance of micro-prudential liquidity regulation maintaining a relatively tight definition of what constitutes a liquid asset. In particular, it suggests that holdings of bank debt (e.g. certificates of deposit, covered bonds, etc.) should not constitute the majority of banks’ prudential liquid asset requirements because such assets are likely to have large haircuts and thus relatively little value in a situation of systemic stress. They are ‘inside’ liquidity. By contrast, genuine ‘outside’ liquidity is likely to prove much more useful. From an *ex ante* perspective, the model also suggests that maturing interbank assets should not be allowed to contribute to liquid asset requirements because if they are, then there is again less overall ‘outside’ liquidity available in stress. Indeed, one of the central points of the model is to demonstrate how withdrawals of interbank assets may be a key amplifier in precipitating a systemic liquidity crisis.²⁵

Regulators internationally have just agreed, for the first time, international standards for liquidity regulation. This includes a requirement for banks to hold sufficient genuinely liquid securities to cover prospective outflows of funds. The definition of liquid assets is narrowly drawn and seeks explicitly to exclude ‘inside’ liquidity. Some of the key parameters of this liquidity regime can be justified analytically using the model developed here.

(b) *Macro-prudential policy and systemic surcharges.* The results on concentration and complexity highlight two different dimensions of risk linked to intra-financial system activity. For a given level of concentration, lower aggregate haircuts or a more complex network are likely to make the system more prone to collapse. Similarly, for a given level of complexity, a more concentrated network is more vulnerable to shocks to key banks. Past experience suggests that fluctuations in haircuts and the complexity of the network are likely to be at least partially cyclical, while changes in network concentration are likely to be more structural.

By the Tinbergen principle, two distinct policy instruments (one time-varying and one structural) are thus likely to be needed to help address intra-financial system risks. These form the bedrock of so-called macro-prudential policy. With the advent of the Financial System Oversight Committee in the US, the European Systemic Risk Board in the EU and the Financial Policy Committee in the UK, macro-prudential policy is about to become a reality.

²⁵ This second issue is less clear-cut *ex post* once a liquidity crisis has begun. In this case, again for reasons made clear by the model, it is better from a systemic perspective for banks to use their ‘outside’ liquidity to meet liquidity needs rather than withdrawing interbank funding. Therefore, care is needed to ensure that banks do not have strong incentives to withdraw interbank funding rather than run down liquid assets *ex post*, for example because such action helps to maintain prudential liquid asset buffers at a high level.

In terms of *structural* macro-prudential instruments, one regulatory response to the too concentrated to fail problem is to levy a tax on institutions in line with their contribution to systemic risk. Because systemic risk is an externality, such an approach is equivalent to the imposition of a Pigouvian tax. Within the international regulatory community, the G20 and Financial Stability Board (FSB) are drawing up a blueprint for dealing with so-called systemically important financial institutions (SIFIs). Among the options on the table are additional, graduated capital charges and regulatory limitations on the extent of exposures between these SIFIs (a 'large exposures' regime).

To target concentration and the structural aspects of interconnectivity effectively, additional capital or liquidity could be set aside according to banks' market shares in interbank activity (e.g. the share of a bank's interbank assets as a fraction of total interbank assets across all banks) and potentially other systemically important markets. These additional buffers would then help both to enhance the resilience of SIFIs to shocks and create incentives for the network to become less concentrated in the first place.

Our second policy exercise highlights the potential improvements in system resilience which may be achieved by targeting higher liquidity requirements at the most connected banks in the network. Importantly, it shows that even a mean-preserving redistribution of liquid assets towards the more connected banks is stability-enhancing for the system as a whole. And this exercise may understate the potential benefits of systemic liquidity policies because it does not account for any reduction in concentration that might be incentivised by them.²⁶

The *cyclical* aim of macro-prudential policy is to lean against the wind of the credit cycle, much like monetary policy does in respect of the business cycle (Bank of England, 2009; Hanson et al., 2011). There are a variety of instruments that might be used, including varying headline regulatory capital and liquidity ratios, changing the risk weights which attach to some of the components of lending, and adjusting haircuts on secured financing to financial institutions (CGFS, 2010).

Our third policy exercise demonstrates how haircut-dependent liquidity requirements, which can be interpreted as time-varying liquidity requirements given the cyclical tendencies of haircuts, can help to deliver a less fragile financial network. Increasing the amount of liquidity held by banks as complexity increases may reduce the effects of liquidity hoarding in the event of stress. Such a policy may also help to mitigate the cyclical build-up in complexity by dampening the money multiplier. Similarly, adjusting the haircuts on secured financing in a pro-cyclical fashion could also help defuse the cyclical dynamics highlighted in our sixth experiment. More generally, time-varying macro-prudential policy could be targeted at changes in intra-financial system activity given its inherent cyclicity (and perhaps greater cyclicity than real economy lending—see Fig. 3). And raising the risk weights on intra-financial system exposures would serve as another way of disincentivising network complexity.

(c) *Network transparency*. At present, very little is known about the dynamic properties of the financial network. This is in part a result of data deficiencies as financial data have tended to be collected and analysed on an institution by institution basis (Haldane, 2009). Such data do not allow an effective mapping of the entire financial web or a simulation of its properties, though small parts of its sub-structures such as the domestic payments and interbank networks have been mapped in some countries. This may be all about to change. There are efforts internationally to begin collecting systematically much greater amounts of data on evolving financial network structure, potentially in close to real time. For example, the introduction of the Office of Financial Research (OFR) under the Dodd-Frank Act will nudge the United States in this direction.

This data revolution potentially brings at least two benefits. First, it ought to provide the authorities with data to calibrate and parameterise the sort of network framework developed here. An empirical mapping of the true network structure should allow for better identification of potential financial tipping points and cliff edges across the financial system. It could thus provide a sounder, quantitative basis for judging remedial policy actions to avoid these cliff edges.

Second, more publicly available data on network structures may affect the behaviour of financial institutions in the network. Armed with greater information on counterparty risk, banks may feel less need to hoard liquidity following a disturbance. A reduction in network uncertainty and counterparty risk could make banks' liquidity decisions less hair-trigger. This would have a collective benefit, lessening the potential for destabilising liquidity spirals.

In the context of the model, a lower incentive to hoard liquidity would reduce the value of λ . Our final policy exercise shows how this diminishes the frequency of contagion across the network.

(d) *Netting and central clearing*. The financial system is a dense cats-cradle of exposures. These gross exposures can easily exceed a bank's capital. And it is gross exposures which matter when gauging the virulence of contagion through a network. One means of reducing this contagion is by netting-off gross exposures between participants within the financial system as this would lower the value of interbank connections relative to balance sheet size. In graphical terms, this would have a similar effect to moving from the case with 25% interbank liabilities to the baseline in Fig. 7.

There have been recent attempts by regulators internationally to achieve such netting benefits across the financial system, in particular in respect of derivative contracts. Algorithms have been developed which either eliminate (in the case of redundant transactions) or net-down (in the case of perfectly offsetting transactions) outstanding derivative contracts.

²⁶ There is an interesting analogy to be drawn here with the use of targeted vaccination programmes to control the outbreak of disease. It is a well-known result in the epidemiology literature that if vaccine stocks are limited, programmes should be targeted at the most highly connected individuals (Anderson and May, 1991). The same logic applies to banks here. But there is also an important difference: while vaccinating particular individuals is unlikely to reduce the size of their social networks, applying surcharges to highly connected banks may create real costs for them which may affect the structure of the network in a desirable way. Thus targeted approaches may be even more useful in relation to financial networks than they are in epidemiology.

A more ambitious, and far-reaching, regulatory initiative is the drive from the G20 and FSB to centrally clear a much larger proportion of, in particular, over-the-counter products through central counterparties (CCPs). Our model speaks to the rationale for such action given that CCPs may affect both complexity and concentration.

On complexity, a CCP simplifies the network of bilateral exposures. Higher-order, unobservable counterparty credit risk is replaced by first-order, observable counterparty risk with respect to the CCP. This, too, ought to reduce the sensitivity of banks' liquidity hoarding to disturbances, thereby lowering the fragility of the system as a whole, in line with the fourth policy exercise.

On concentration, a CCP alters the structural configuration of the network. This now resembles a star formation rather than a cats-cradle. From a resilience perspective, this is a double-edged sword. On the one hand, concentration among the key financial institutions is effectively eliminated, reducing the contagion risk from SIFIs. In the context of the model, this is the limiting case of moving towards a more fat-tailed distribution, comparing the baseline in Figs. 6 and 7. On the other, concentration risk is, in an important sense, relocated rather than eliminated, from SIFIs to the CCP. This underscores the importance of ensuring that CCPs are bullet-proof moving forward as they clear larger numbers of transactions.

At the height of the crisis, the financial network degenerated to a hub-and-spoke configuration. The central bank became the de facto central counterparty. To avoid that outcome ex post, a set of bullet-proof CCPs could enforce a simplified hub-and-spoke configuration ex ante. From an incentives and counterparty uncertainty perspective, this ex ante structure seems preferable.

6. Conclusion

This paper has developed a network model of interbank interactions which seeks to embody many of the key structural features of the financial system. The framework generates knife-edge liquidity dynamics which have strong echoes with the systemic liquidity crisis of 2007–2008. We showed how concentration and complexity may be key amplifiers of this fragility by potentially leaving the system highly vulnerable to distress at key institutions and increasing risk over the economic cycle. The paper has also identified some public policy measures which could mitigate these fragilities, many of which are currently being designed or implemented internationally. These include tougher micro-prudential liquidity regulation, countercyclical liquidity requirements, and surcharges for systemically important financial institutions.

The model could be improved in several respects. It could be made more practically relevant by incorporating a stronger role for uncertainty. Calibrating it to the data would allow it to match more closely the unfolding of the crisis, though the lack of available information on the interbank network and haircuts would present significant challenges. It would also be interesting to enrich the framework by incorporating a stronger role for behavioural considerations, for example in relation to the formation of links. This would allow for endogeneity in the network structure.

Another interesting extension of the network approach would be to consider the effects of policy measures which act on financial system structure more directly—a quantity rather than price-based approach to regulation. Following Weitzman (1974), this directive approach may be optimal if there is uncertainty about the elasticities associated with price-based instruments. In the spirit of Simon (1962), the aim of such structural interventions would be to create a financial system that was decomposable and modular. A number of regulatory initiatives and ideas are currently in train which speak to this agenda.

One is 'living wills' or, more formally, recovery and resolution plans (RRPs). These require banks to draw up plans for how they would be wound-down in an orderly fashion in the event of failure. The newly passed Dodd-Frank Act in the US requires that US firms draw up credible RRP. Pilot RRP are currently being drawn up for the world's largest financial institutions. A second idea, which is being debated within academic and policymaking circles, is the potential for ring-fencing of financial activities, either within firms or across them. The Volcker rule, enacted in US Legislation recently, is one example, in this case in respect of proprietary trading activity; the Glass-Steagall Act, which separated investment and commercial banking in the US between 1933 and 1999, is another.

These structural measures have the advantage that they act directly on the topology of the network. As such, they may be less subject to implementation error. Firebreaks and firewalls are familiar fail-safe devices against systemic problems in other networks, from infectious diseases spreading across people to infectious viruses spreading across the web, from transport networks to meteorology maps, from forest-fire control to military control (Haldane, 2009). The model developed here could be adapted to assess the systemic implications of these structural policy measures. In some respects, this would represent a truer test of Herbert Simon's complexity hypothesis in a financial system context.

Disclaimer

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

Acknowledgements

The paper was prepared for the November 2010 Carnegie-Rochester Conference on Public Policy. We thank Clare Rogowski and, particularly, Jason Dowson for outstanding research assistance. We are also grateful to John Elliott, Marvin

Goodfriend, Andrew Lo, Lavan Mahadeva and seminar participants at the Bank of England, Imperial College London, the University of Auckland, the Carnegie-Rochester Conference on Public Policy (Pittsburgh, 12–13 November 2010) and the Reserve Bank of New Zealand Workshop on Macro-prudential Policy (Wellington, 21 March 2011) for helpful comments and suggestions.

References

- Acharya, V., Gale, D., Yorulmazer, T. Rollover risk and market freezes. *Journal of Finance*, in press.
- Acharya, V., Skeie, D. A model of liquidity hoarding and term premia in inter-bank markets. *Journal of Monetary Economics*, this issue. doi:10.1016/j.jmoneco.2011.05.006.
- Adrian, T., Shin, H.S., 2010a. Liquidity and leverage. *Journal of Financial Intermediation* 19 (3), 418–437.
- Adrian, T., Shin, H.S., 2010b. The changing nature of financial intermediation and the financial crisis of 2007–09. *Annual Review of Economics* 2, 603–618.
- Albert, R., Jeong, H., Barabasi, A.-L., 2000. Error and attack tolerance of complex networks. *Nature* 406, 378–382.
- Alessandri, P., Haldane, A.G., 2009. Banking on the state. Paper from Material Presented at the Federal Reserve Bank of Chicago/World Bank Event—International Financial Crisis: Have the Rules of Finance Changed? 24 September 2009.
- Allen, F., Babus, A., 2009. Networks in finance. In: Kleindorfer, P., Wind, Y., Gunther, R. (Eds.), *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World*. Wharton School Publishing.
- Allen, F., Carletti, E., Gale, D., 2009. Interbank market liquidity and central bank intervention. *Journal of Monetary Economics* 56 (5), 639–652.
- Allen, F., Gale, D., 2000. Financial contagion. *Journal of Political Economy* 108 (1), 1–33.
- Allen, F., Gale, D., 2004. Financial intermediaries and markets. *Econometrica* 72 (4), 1023–1061.
- Anderson, R., May, R., 1991. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press.
- Bank of England, 2009. *The Role of Macroprudential Policy*. November.
- Bolton, P., Santos, T., Scheinkman, J., 2011. Outside and inside liquidity. *Quarterly Journal of Economics* 126 (1), 259–321.
- Boss, M., Elsinger, H., Thurner, S., Summer, M., 2004. Network topology of the interbank market. *Quantitative Finance* 4 (6), 677–684.
- Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. *Review of Financial Studies* 22 (6), 2201–2238.
- Brusco, S., Castiglionesi, F., 2007. Liquidity coinsurance, moral hazard, and financial contagion. *Journal of Finance* 62 (5), 2275–2302.
- Caballero, R., Krishnamurthy, A., 2008. Collective risk management in a flight to quality episode. *Journal of Finance* 63 (5), 2195–2230.
- Caballero, R., Simsek, A., 2010. Fire sales in a model of complexity. Mimeo. MIT.
- CGFS, 2010. The role of margin requirements and haircuts in procyclicality. CGFS Papers, Number 36.
- Craig, B., von Peter, G., 2010. Interbank tiering and money center banks. Bank for International Settlements Working Paper, Number 322.
- Diamond, D., Rajan, R. Fear of fire sales, illiquidity seeking and credit freezes. *Quarterly Journal of Economics*, in press.
- Freixas, X., Parigi, B., Rochet, J.-C., 2000. Systemic risk, interbank relations and liquidity provision by the central bank. *Journal of Money, Credit and Banking* 32 (3), 611–638.
- Furfine, C.H., 2003. Interbank exposures: quantifying the risk of contagion. *Journal of Money, Credit and Banking* 35 (1), 111–128.
- Gai, P., Kapadia, S., 2010. Contagion in financial networks. *Proceedings of the Royal Society A* 466 (2120), 2401–2423.
- Gai, P., Kapadia, S., 2011. Liquidity hoarding, network externalities, and interbank market collapse. Mimeo. Bank of England.
- Garratt, R.J., Mahadeva, L., Svirydzenka K., 2011. Mapping systemic risk in the international banking network. Bank of England Working Paper, Number 413.
- Geanakoplos, J., 2010. Solving the present crisis and managing the leverage cycle. Federal Reserve Bank of New York Economic Policy Review, August 16 (1).
- Gorton, G., Metrick, A., 2010. Haircuts. Federal Reserve Bank of St. Louis Review, November/December 92 (6), 507–520.
- Haldane, A.G., 2009. Rethinking the financial network. In: Speech at the Financial Student Association, Amsterdam, 28 April.
- Haldane, A.G., May, R., 2011. Systemic risk in banking ecosystems. *Nature* 469, 351–355.
- Hanson, S.G., Kashyap, A.K., Stein, J.C., 2011. A macroprudential approach to financial regulation. *Journal of Economic Perspectives* 25 (1), 3–28.
- Jackson, M., 2008. *Social and Economic Networks*. Princeton University Press, NJ, USA.
- King, M., 2010. Banking: From Bagehot to Basel, and back again. In: Speech at the Second Bagehot Lecture Buttonwood Gathering, New York City, 25 October.
- Krishnamurthy, A., 2010. How debt markets have malfunctioned in the crisis. *Journal of Economic Perspectives* 24 (1), 3–28.
- Kubelec, C., Sa, F., 2010. The geographical composition of national external balance sheets: 1980–2005. Bank of England Working Paper, Number 384.
- Lo, A. Complexity, concentration and contagion: a comment. *Journal of Monetary Economics*, this issue. doi:10.1016/j.jmoneco.2011.06.001.
- May, R., Levin, S., Sugihara, G., 2008. Ecology for bankers. *Nature* 451, 893–895.
- May, R., Arinaminpathy, N., 2010. Systemic risk: the dynamics of model banking systems. *Journal of the Royal Society Interface* 46, 823–838.
- Newman, M., Strogatz, S., Watts, D., 2001. Random graphs with arbitrary degree distributions and their applications. *Physical Review E* 64, 026118.
- Simon, H., 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* 106 (6), 467–482.
- Shin, H., 2009. Financial Intermediation and the post-crisis financial system. Paper Presented at the 8th BIS Annual Conference, June.
- Stiglitz, J., Weiss, A., 1981. Credit rationing in markets with imperfect information. *American Economic Review* 71 (3), 393–410.
- Upper, C., 2011. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability* 7 (3), 111–125.
- von Peter, G., 2007. International banking centres: a network perspective. *BIS Quarterly Review*, December. <http://www.bis.org/publ/qtrpdf/r_qt0712.htm>.
- Watts, D., 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences* 99 (9), 5766–5771.
- Weitzman, M.L., 1974. Prices vs. quantities. *Review of Economic Studies* 41 (4), 477–491.