BANK OF ENGLAND

Working Paper No. 465

Size and complexity in model financial systems

Nimalan Arinaminpathy, Sujit Kapadia and Robert May

October 2012

BANK OF ENGLAND

Working Paper No. 465
Size and complexity in model financial systems

Nimalan Arinaminpathy,[(1)] Sujit Kapadia[(2)] and Robert May[(3)]

Abstract

The global financial crisis has precipitated an increasing appreciation of the need for a systemic perspective towards financial stability. For example: What role do large banks play in systemic risk? How should capital adequacy standards recognize this role? How is stability shaped by concentration and diversification in the financial system? We explore these questions using a deliberately simplified, dynamical model of a banking system which combines three different channels for direct spillovers from one bank to another: liquidity hoarding, asset price contagion, and the propagation of defaults via counterparty credit risk. Importantly, we also introduce a mechanism for capturing how swings in 'confidence' in the system may contribute to instability. Our results highlight that the importance of relatively large, well-connected banks in system stability scales more than proportionately with their size: the impact of their collapse arises not only from their connectivity, but also from their effect on confidence in the system. Imposing tougher capital requirements on larger banks than smaller ones can thus enhance the resilience of the system. Moreover, these effects are more pronounced in more concentrated systems, and continue to apply even when allowing for potential diversification benefits which may be realised by larger banks. We discuss some tentative implications for policy, as well as conceptual analogies in ecosystem stability, and in the control of infectious diseases.

Key words: Systemic risk, financial crises, contagion, network models, liquidity risk, confidence.

JEL classification: D85, G01, G21, G28.

(1) Princeton University. Email: nimpathy@princeton.edu
(2) Bank of England. Email: sujit.kapadia@bankofengland.co.uk (corresponding author)
(3) Oxford University. Email: robert.may@zoo.ox.ac.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or Financial Policy Committee members. The paper is forthcoming in *Proceedings of the National Academy of Sciences*. The authors thank David Barr, Andy Haldane, Simon Hall, four reviewers, and seminar participants at the Bank of England, Princeton University and the Geneva Finance Research Institute Conference on 'Financial Networks' (Geneva, 10 June 2011) for helpful comments. Nimalan Arinaminpathy was supported by NIH grant R01 GM083983-01, and the Bill and Melinda Gates Foundation. This paper was finalised on 26 September 2012.

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

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

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

ISSN 1749-9135 (on-line)

Contents

Summary

What role do large banks play in systemic risk and financial instability? How should capital adequacy standards recognize this role? How is stability shaped by concentration and diversification in the financial system? This paper explores these questions using a deliberately simplified, dynamical model of a banking system.

Developing methods used in epidemiology and ecology, we adopt network techniques which are well suited for such questions, particularly in modelling 'contagion' that is transmitted through linkages in the financial system. Specifically, we bring together three important transmission channels into a unified framework: (i) liquidity hoarding, where banks may cut their lending to each other as a defensive measure; (ii) asset price contagion linked to the falls in market prices which may be generated by asset sales by banks in distress; and (iii) the propagation of losses which may occur if banks default on their obligations to other banks in the interbank market (the network of lending exposures amongst banks). Importantly, we also integrate a mechanism for capturing how broader swings in 'confidence' in the system may contribute to instability, with the overall state of the system potentially influencing an individual bank's actions, and *vice versa*.

The interaction of such network and confidence effects arguably played a major role in the collapse of the interbank market and global liquidity 'freeze' that occurred during the financial crisis. Interbank loans have a range of maturities, from overnight to a matter of years, and may often be renewed, or 'rolled over', at the point of maturity. A pronounced feature of the 2007-08 crisis was that, as the system deteriorated, banks stopped lending to each other at all but the shortest maturities. The bankruptcy of Lehman Brothers in September 2008 transmitted distress further across the financial network. The effects extended well beyond those institutions directly exposed to Lehman Brothers, with banks throughout the system withdrawing interbank lending outright and propagating distress to the real economy by sharply contracting household and corporate lending.

Several specific motivating factors have been proposed to explain 'liquidity hoarding' (the maturity-shortening and ultimate withdrawal of interbank lending): precautionary measures by lending banks in anticipation of future liquidity shortfalls; counterparty concerns over specific borrowing banks; or collapses in overall system confidence. Our framework parsimoniously

incorporates all of these mechanisms, while also capturing the idea that a bank's distress may affect not just those directly exposed or linked to it, but also confidence in the market at large.

We use our model to explore the effects of shocks to the system, such as the failure of banks or big losses on certain types of lending. We focus particularly on the adverse feedback dynamics arising from each of the contagion channels included, the effects of size disparity amongst banks and system concentration, and the effects of diversification. Our results highlight the disproportionate importance of large, well-connected banks for system stability: the impact of their collapse arises not only from their connectivity, but also from their effect on confidence in the system. Moreover, we show that while diversification may serve to limit the risk of failure of an individual bank, it does not mitigate the importance of that bank to systemic risk, and may indeed exacerbate it. Overall, these results illustrate the different approaches needed for regulation focused at the level of individual banks, and that focused on a systemic level. While sound microprudential regulation remains important for the former, the latter, macroprudential perspective, supports the notion of regulatory requirements concomitant with bank size, interconnectedness or (more generally) systemic importance. In particular, imposing tougher capital requirements on larger banks than smaller ones can enhance the resilience of the system. Furthermore, such requirements may also have the beneficial side-effect of providing disincentives for financial institutions to become 'too big to fail'. Our findings have conceptual analogies in ecosystem stability, and in the control of infectious diseases, which we also discuss briefly.

As with any theoretical approach, there are important caveats to our model. In particular, a key empirical challenge for future work is to quantify the confidence processes which we model. Incorporating uncertainty, for example over the underlying health of individual institutions or the system as whole, would also be a useful extension. Another key question is how the vulnerabilities in financial systems modelled in this paper emerge, and potentially grow, over time. Finally, while this paper focuses on one aspect of the regulatory response relating to capital requirements, other policy responses, such as the use of liquidity requirements or the implementation of effective resolution regimes, are also likely to be important in enhancing the resilience of the financial system.

1 Introduction

While global financial systems have seen considerable growth in size, concentration and complexity over the past few decades (Gai *et al* (2011)), our understanding of the dynamical behaviour of such systems has not necessarily kept pace. Indeed, the current financial crisis has presented a stark demonstration of the potential for modern financial systems to amplify and disseminate financial distress on a global scale. From a regulatory perspective, these events have prompted fresh interest in understanding financial stability from a system level. In particular, while pre-crisis regulation (as typified by the Basel II accords) sought to minimize the risk of failure of individual banks irrespective of systemic importance, new regulation will seek to target the systemic *consequences* of bank collapse as well. To quote Haldane and May (2011b), "What matters is not a bank's closeness to the edge of the cliff; it is the extent of the fall".

In this context, a clear feature of interest is the presence of large, highly connected banks. These have conceptual parallels in biology: simple models have been influential in underlining the importance of 'superspreaders' in the spread and control of infectious diseases (Anderson and May (1991); Lloyd-Smith *et al* (2005)), while 'keystone' species are thought to serve a valuable role in ecosystem stability (Paine (1966); Kareiva and Levin (2003)). Here we develop dynamical models to apply and extend these lessons to financial systems. Our approach is theoretical, and our models necessarily oversimplified. Nonetheless, by considering transmission mechanisms specific to modern financial systems, our approach recognizes some important differences between these and other complex systems. We show how, even with such distinctions, the basic insights deriving from our model allow us to draw certain parallels with other situations where size and complexity are important.

If financial crises may be compared with forest fires, causes for the 'initiating sparks' pose important questions in their own right: for example the role of excessive leverage and credit growth (Bank of England (2011)), or the pricing for complex financial instruments (Caccioli *et al* (2009); Haldane and May (2011a)). Here, however, our focus is on the role of large banks in the 'flammability' of the system, or its capacity for amplification and propagation of an initiating shock. We ask the following questions: how does the impact of a bank's collapse scale with its size? How might capital adequacy standards seek to mitigate this impact? More broadly, what is the effect of concentration and diversification on system stability?

Network approaches (Strogatz (2001); Jackson (2008); Allen and Babus (2009); Battiston *et al* (2009); Kirman (2010)) are well-suited for such questions, particularly in modelling 'contagion' that is transmitted through linkages in the financial system. Here we adopt such an approach to bring together three important transmission channels into a unified framework: (i) liquidity hoarding, where banks cut lending to each other as a defensive measure (Brunnermeier (2009); Gai *et al* (2011)), (ii) asset price contagion linked to market illiquidity (Cifuentes *et al* (2005); Coval and Stafford (2007); Adrian and Shin (2010)), and (iii) the propagation of defaults via counterparty credit risk (Nier *et al* (2007); Gai and Kapadia (2010); Upper (2011)).

While the 'network effects' listed above act on defined webs of connectivity, 'confidence effects' can operate more broadly, with the overall state of the system potentially influencing an individual bank's actions, and *vice versa*. This motivates a special feature of our model, which explicitly integrates network dynamics with confidence effects.

The interaction of such network and confidence effects arguably played a major role in the collapse of the interbank market (a network of lending exposures amongst banks) and global liquidity 'freeze' that occurred during the crisis (Gorton and Metrick (2012)). Interbank loans have a range of maturities, from overnight to a matter of years, and may often be renewed, or 'rolled over', at the point of maturity. A pronounced feature of the 2007-08 crisis was that, as the system deteriorated, banks stopped lending to each other at all but the shortest maturities (Bank of England (2011); Kapadia *et al* (2012)). The bankruptcy of Lehman Brothers in September 2008 transmitted distress further across the financial network, while signalling that there was no guarantee of government support for institutions in distress. The effects extended well beyond those institutions directly exposed to Lehman Brothers, with banks throughout the system withdrawing interbank lending outright and propagating distress to the real economy by sharply contracting household and corporate lending (Ivashina and Scharfstein (2010)). At the time of writing, ongoing events illustrate the potential for similar dynamics in the context of sovereign and banking sector distress in some eurozone countries.

Several specific motivating factors have been proposed to explain 'liquidity hoarding' (the maturity-shortening and ultimate withdrawal of interbank lending): precautionary measures by lending banks in anticipation of future liquidity shortfalls; counterparty concerns over specific borrowing banks; or collapses in overall system confidence (Caballero and Krishnamurthy (2008); Acharya and Skeie (2011)). Our framework parsimoniously incorporates all of these mechanisms, while also capturing the idea that a bank's distress may affect not just those directly exposed or linked to it, but also confidence in the market at large.

In what follows we summarize essential features of the model structure, with details provided in Appendix 1, and a summary of model parameters and their default values, given in Table 1. We use this model to explore the impact of an initiating shock, with particular reference to the non-linearities arising from each of the contagion channels modelled, the effects of size disparity amongst banks and system concentration, and the effects of diversification. We then outline tentative implications for regulatory capital requirements before discussing important caveats to our work. Throughout the paper, we abstract from extraordinary policy intervention in crisis, so that liquidity cannot be obtained more easily from the central bank than from the market, and failing institutions are not bailed out.

2 Model overview

Figure 1 shows a minimal balance sheet representation of an individual 'bank', or node in the system. On the asset side of Figure 1, this bank lends to other banks in the system and holds a small proportion l of assets as 'liquid assets' (eg cash, central bank reserves and high-quality government bonds). The remainder of the asset side consists of holdings in (and thus exposures to) a range of distinct 'external' asset classes held against the real economy, such as mortgages, corporate lending and commercial real estate lending. The liability side is even simpler, consisting of retail deposits (taken to be external to the system) and interbank borrowing. The 'capital buffer' is the excess of assets over (debt) liabilities and if this falls below zero, the bank is insolvent ('capital default'). In our treatment, we think of capital in simple terms consisting only of common equity, thus excluding any form of unsecured term debt (eg subordinated debt) or contingent capital. Given that we also abstract from risk-weighting of assets, the capital-to-assets ratio, γ, can be thought of as a simple leverage ratio, with our baseline choice of 4% conservatively reflecting typical leverage ratios seen prior to the crisis (Bank of England (2011)).

We take a system of 200 banks, interconnected in two distinct ways: (i) the interbank market, represented by a directed network with 'nodes' being individual banks and each 'edge' being a loan from one bank to another; and (ii) a system of shared exposures to a set of external assets, such that two different banks may hold some external asset classes in common, but not necessarily all: see Figure 2 for a schematic illustration of this overall scheme. For the interbank network we label half of loans, at random, as initially having 'long-term' maturity, and the remainder as being 'short-term': in our discrete time simulations we assume that long-term loans can be made short-term from one time step to the next, while only short-term loans can be withdrawn in the same interval.

We assume for simplicity that there are two sizes of banks, where 'big' banks are λ times larger on average than 'small' ones, but are λ times fewer. Thus big banks issue and receive λ times as many loans, and – by holding the same number of external asset classes as small banks in our baseline setup – they hold λ times as much of each single asset class, including interbank assets. We subsequently allow large banks to be more diversified (holding more asset classes) than smaller ones. In both setups, bank size and interconnectivity are correlated, in agreement with the data (Drehmann and Tarashev (2011)). This formulation also yields 'heavy-tailed' bank size and connectivity distributions, consistent with available evidence (Boss *et al* (2004); Soramaki *et al* (2007)). Note also that λ can be interpreted as a 'concentration' parameter, as the proportion of assets in the system held by big banks, divided by the proportion of banks that are big, is easily shown to be ½ (1+ λ). While this framework lends itself to a straightforward comparison of 'big' and 'small' banks, we explore an alternative bank size distribution in Appendix 2.

So far we have described the structure of the system, and its interconnectedness. To describe its dynamics, we now explicitly link confidence effects to fundamental balance sheet characteristics such as capital and liquidity strength. We define the 'system confidence', C, and the 'individual health', h_i, of bank i as follows:

$$C = AE; \qquad h_i = c_i\, m_i; \qquad \text{with} \quad 0 \le h_i \le 1 \text{ and } 0 \le E, A \le 1, \qquad \textbf{(1)}$$

where, at a given time, A is the total value of all remaining assets in the system (at the current market price) as a proportion of its initial level (reflecting confidence in terms of solvency); E is the fraction of interbank loans not withdrawn (reflecting confidence in terms of liquidity); c_i is the capital of bank i as a proportion of its initial value, and m_i reflects the bank's liquidity position, the fraction of its short-term liabilities that it can settle immediately, through its liquid and short-term assets (see Appendix 1 for details of how c_i and m_i are calculated). Since the model only analyses the (crisis) dynamics of the system after adverse shocks, asset prices (and thus A) cannot increase above their initial values, as discussed further in Appendix 1. Similarly, E, c_i and m_i cannot exceed 1.

We interpret C as reflecting the condition of the system, equally felt by all participants, given that A and E are taken to be common knowledge. Equation **(1)** is of course one of many possible functional forms. In particular, it neglects the effect of maturity shortening on system confidence C. Results of an alternative formulation, aiming to incorporate this effect, are presented in Appendix 2.[1]

We discussed above how liquidity hoarding could arise from a combination of factors, including a bank's own health; that of its counterparties; and more broadly, confidence in the system as a whole. Terms in equation **(1)** can be used to capture all three factors. Specifically, we suppose simply that a long-term link from bank i to bank j (these being lender and borrower, respectively) is made short term whenever

$$h_i h_j < (1 - C). \qquad \textbf{(2)}$$

As described in Appendix 1, such behaviour raises m_i, thus allowing an individual bank to improve its own health. We also suppose that a short-term loan is withdrawn altogether (ie the corresponding 'link' is removed from the network) whenever

$$h_i h_j < (1 - C)^2. \qquad \textbf{(3)}$$

Withdrawals can also propagate through the interbank network if borrowers need to recall their own interbank lending to meet their obligations, with banks experiencing 'liquidity default' if they have insufficient liquidity to do so.

The intuitive basis for the rules embodied in **(2)** and **(3)** is as follows: when $C = 1$, there is no hoarding. With a perturbation to C, however, hoarding may potentially be triggered anywhere in the system either because of precautionary motives (driven by the bank's own health, h_i) or counterparty concerns (driven by h_j), both causing and exacerbated by further deterioration in C – in particular, falls in A, which might be associated with market illiquidity (see below), drive

[1] In principle, this framework could also be extended to incorporate exogenous confidence shocks driven by a deterioration in market conditions or heightened volatility, perhaps as might be embodied in the VIX; or uncertainty over A and E in equation **(1)**.

down *C* and thus contribute to funding illiquidity (a reduction in *E*). In equation **(3)**, the exponent on the right-hand side ensures that withdrawals – being a last resort – happen only under more extreme stress than shortening of loan maturities. While clearly stylized, these rules are broadly consistent both with observed behaviour during the crisis and with notions of optimizing balance sheet management by *individual* banks (Acharya and Skeie (2011); Kapadia *et al* (2012)), though as will become evident below, such behaviour has clear adverse spillovers for other banks in the system. As illustrated in the simulations, equations **(2)** and **(3)** can also drive large fluctuations in balance sheets operating through changes in debt rather than equity, thus allowing the model to generate the procylicality in leverage that many financial institutions are seen to exhibit in the data (Kim *et al* (2012)).

While liquidity hoarding acts on the interbank network, we also incorporate shocks transmitted through the system of external assets. As a bank fails, the sale of its assets to outside investors has the potential to drive down the price of those assets in the market, adversely affecting the capital position of other banks also holding these assets. Previous work (Cifuentes *et al* (2005); Nier *et al* (2007); Gai and Kapadia (2010); May and Arinaminpathy (2010)) has illustrated the potential for this process to push these additional banks to default, thus reinforcing the downward pressure on asset prices. The 'intensity' of this process is related to 'market liquidity', or the ability of assets to be sold without significant price movements. In our framework, we directly link market liquidity to system confidence *C*, as detailed in Appendix 1. This captures the idea that confidence effects – for example, as driven by social contagion (Marsili *et al* (2010)) – are likely to cause market liquidity to deteriorate in times of crisis; it also implies that funding illiquidity, as captured by a reduction in *E*, can exacerbate market illiquidity. Overall, then, asset price contagion is determined both by the amount of assets liquidated and by overall market conditions at that time. In Appendix 2, we also consider the behaviour of the model under two alternative formulations of asset price contagion – one in which banks can sell their external assets prior to failure in an attempt to avoid liquidity default, and one in which system confidence does not affect market illiquidity.

Finally, we include the potential for cascades of capital default through the interbank network, as explored by other authors (Nier *et al* (2007); Gai and Kapadia (2010); May and Arinaminpathy (2010)): should bank *i* undergo capital default, it cannot settle its debt in its entirety, and thus its creditors, in turn, lose value on their interbank lending, eroding their own assets and thus their own capital. Further details are given in Appendix 1. In summary, our model combines channels for ‘direct’ transmission between banks (that is, of asset contagion and counterparty default) with the ‘indirect’, system-wide effects represented by equations **(1)-(3)**. It also captures how the interplay of market and funding illiquidity can generate a downward spiral during crises (Brunnermeier and Pedersen (2009)).

3 Results

For our baseline simulations, we take $\lambda = 24$, giving a system with 8 ‘large’ banks and 192 ‘small’ ones. We assume that all banks have the same capital and liquidity ratios (broadly in line with observed median pre-crisis ratios (Bank of England (2011)), with parameters given in Table 1. Clearly the results presented here are dependent on the type of shock and initial parameters chosen, so should be interpreted as an illustration of systemic effects only.

To demonstrate the significance of liquidity hoarding and confidence effects, we apply a shock to the capital buffer of a randomly chosen, large ‘index bank’ by setting the value of one of its external assets, also selected at random, to zero (other banks holding this asset are unaffected). As a measure of impact, we then record the loss in total assets across the system as a fraction of its initial value, thus counting large banks in proportion to their size. Figure 3 shows mean results, both with and without liquidity hoarding. Counterparty credit risk and asset price contagion generate losses in both cases, but hoarding has a clear negative externality on the system: although representing defensive behaviour on the part of individual banks to improve their own liquidity position, it leads to a decrease in E and thus C in equation **(1)**. Its effects are also amplified by an adverse feedback with asset price contagion.

Considering the evolution of bank balance sheets is a useful approach for unpicking the relative contributions of different contagion channels illustrated in Figure 3, and comparing this to data. Figure 4 (panel A) illustrates recent empirical findings (Kim *et al* (2012)), that contractions in balance sheets tend to be associated primarily with debt (linked to the withdrawal of interbank lending), rather than with diminishing equity (eg from asset contagion or counterparty credit risk). Figure 4 (panel B) shows corresponding model results, illustrating behaviour qualitatively consistent with panel A, owing to the inclusion of confidence effects and liquidity hoarding. Omitting hoarding, and withdrawal of lending in particular, yields the converse outcome in which the blue line has zero gradient, and the red accounts for all changes in bank assets.

Next, we explore connections between idiosyncratic and systemic risk, in particular how the systemic importance of a given bank depends on its size. Here, we measure the 'importance' of a given bank as the overall impact to the system arising from its idiosyncratic failure: we choose an 'index' bank of given size at random, and force it to default by setting its capital buffer to zero.[2] Without affecting any other part of the balance sheet, this initial condition ensures that, beyond the system's exposures to the failing bank (through shared assets and counterparty links), there is no exogenous difference between the collapse of a small bank and a big one: in particular, while the initial capital loss is greater when a big index bank fails, this does not affect simulation outcomes since (for example) system confidence, C, depends on asset, and not capital, positions. Figure 5 (panels A and B) shows frequency distributions for the resulting number of failed banks, comparing the cases of small and big index banks. They demonstrate that, while the failure of a small bank affects a relatively small group of other banks, the impact of a large bank collapse scales *more* than proportionately with size, entailing a non-zero probability of whole-system collapse.

To explore this systematically, as a measure of 'impact' we write f_S for the mean fraction of total assets lost following the collapse of a small index bank, and correspondingly f_B in the case of a large index bank. The ratio $R = f_B / f_S$ thus gives some measure of the disparity between the impact of big and small bank collapse, and Figure 5 (panel C) plots this ratio for a range of λ. If the impact of bank collapse scales in proportion with bank size, we would expect $R = \lambda$, as represented by the bottom, grey line. In agreement with panels A and B, however, the upper (blue) curve illustrates that R consistently exceeds proportional scaling. Intermediate curves show corresponding results in reduced models where either liquidity hoarding or asset price contagion are excluded: for the parameter ranges considered here and when the initiating shock is only a *single* bank failure, these channels in isolation are relatively limited in their capacity to precipitate system-wide collapse following the failure of a large index bank. (The strength of asset contagion, for example, grows too slowly here with declining C to spread significantly beyond the index bank, whether big or small.) However, these results illustrate how, in combination, these channels mutually reinforce each other in a non-linear way: for example, liquidity hoarding has a negative externality realised in a reduction in C, which in turn can

[2] Although constructed with a focus on exploring the systemic importance of individual institutions, this shock may be seen in practical terms as arising, for example, from the crystallization of operational risk (eg fraud) or from an aggregate shock that has particularly adverse consequences for one institution. In later simulations (Figure 6), we consider results following aggregate shocks in which several small banks fail alongside a single large bank.

exacerbate market illiquidity to intensify asset shocks, further pushing remaining banks nearer to collapse.

Next, how can loss-absorbing capital serve to *mitigate* the systemic effects of an asset shock? Taking $\lambda = 24$ once again, we now consider a scenario in which banks suffer losses because the value of a specific asset class, taken at random subject to being held by 9 small banks and one large one, falls to zero. Figure 6 (panel A) illustrates results, exploring different capital ratios for small and big banks. Higher levels of capital promote system stability in general, but there is an asymmetry with respect to bank size. In particular, the diagonal in the *xy* plane running from the foreground (at (0.15, 0.15)) to the hidden origin (at (0,0)) represents the case where all banks have the same capital ratio irrespective of size. Along this line, contagion may be contained as long as capital is sufficiently high. If not, however, there is a sharp transition in which much of the system collapses. A system in which larger banks have higher capital ratios (ie, where the absolute size of the capital buffer scales *more* than proportionately with size) lies in the region to the left of this line, and conversely for the region to the right. A comparison of these regions illustrates the essential result that contagion is better mitigated by well-capitalized big banks, than by well-capitalized small ones: arguably the converse of the pattern of capital ratios prior to the current financial crisis.

The differences between these regions, while broadly illustrative, will of course depend on concentration, λ, in the system. Recent data (FDIC (2011)) indicates that in the first quarter of 2011, 79% of US banking system assets were held by the largest 1.4% of banks. Translated into our simple framework, this roughly corresponds to a highly concentrated scenario in which the big banks are 250 times the size of small ones, while numbering 3 in a system of 200. Figure 6 (panel B) shows results of simulations adopting these parameters; although highly stylized, these illustrate nonetheless how increasing concentration serves to widen the disparity between the two regions described above. Figure 9 (discussed in Appendix 2) presents results from an alternative model that seeks more faithfully to reproduce the actual distribution of US bank sizes.

So far these results take big banks as simply proportionately scaled versions of small ones and thus neglect the potential for larger banks to mitigate their own risk by having more diversified asset portfolios. Could such behaviour also serve to limit the systemic importance of big banks? Returning to the case $\lambda = 24$, Figure 6 (panel C) allows big banks to hold twice as many asset classes as small banks, under the same asset-specific initiating shock as in panel A. It illustrates the risk-mitigating effects of diversification on the part of individual banks: large banks remain solvent with lower levels of capital than in panel A because they are less vulnerable to losses on any individual asset class. Importantly, however, the *steepness* of the surface is not mitigated by increasing diversification, and is in fact increased. This is because greater diversification increases the overall number of exposures through shared assets, exacerbating the role of asset contagion in system collapse. This effect is underscored by Figure 6 (panel D), which repeats Figure 5 (panel C) while incorporating diversification (see Appendix 1 for details). For sufficiently concentrated systems, this illustrates the potential for asset contagion by itself (red curve) to cause system collapse following the failure of a large index bank, an outcome not apparent in Figure 5 (panel C). Indeed, in the limit, perfectly diversified banks holding an equal fraction of all available assets will also be perfectly correlated (see also Wagner (2011)). In addition, it is clear from these results that the systemic consequences of large bank collapse continue to scale more than proportionately with size.

4 Discussion

The systemic importance of large, well-connected banks has gained recognition in policy discussion, most notably in proposals from the Financial Stability Board and Bank for International Settlements (2011), the UK Independent Commission on Banking (2011), and in the so-called 'Swiss finish' (State Secretariat for International Financial Matters (2010)). However, policy recommendations for requirements concomitant with systemic importance (as reflected in metrics such as bank size and interconnectedness) have drawn some contention from the financial industry.

Our work aims to contribute to this discussion from a dynamical perspective, drawing together different channels of contagion into a unified framework, while incorporating system-level confidence effects. We demonstrate a key consequence of the resulting non-linearities: the disproportionate importance of large, well-connected banks for system stability. Moreover, we show that while asset portfolio diversification may serve to limit the risk of failure of an individual bank, it does not mitigate the importance of that bank to *systemic* risk, and may indeed exacerbate it. Overall, these results illustrate the different approaches needed for regulation focused on idiosyncratic risk, and that focused at a systemic level. While sound microprudential regulation remains important for the former, the latter, relating to the so-called ‘structural’ dimension of macroprudential policy (Bank of England (2009)), supports the notion of regulatory requirements concomitant with bank size, interconnectedness or (more generally) systemic importance. Furthermore, such requirements may also have the beneficial side-effect of providing much-needed disincentives for financial institutions to become ‘too big to fail’.

Our findings echo familiar concepts in other complex systems. Like keystone species in ecosystems (Paine (1966); Kareiva and Levin (2003)), large banks can perform a stabilizing function as long as they remain healthy (Figure 6). Conversely their failure can adversely affect the entire system (Figure 5). In the context of infectious diseases, the largest banks – by their connectivity – are comparable with ‘superspreaders’ of infection (Anderson and May (1991); Lloyd-Smith *et al* (2005)). There too, targeted intervention pays dividends: concentrating vaccination in the most well-connected or the most infectious individuals achieves disease eradication with lower coverage than is required in the case of random vaccination (Anderson and May (1991); Lloyd-Smith *et al* (2005)). As we have explored here, however, balance sheet linkages and pervasive confidence effects can play a distinctive role in financial systems, intensifying these dynamics (Figure 5, panel C).

As with any theoretical approach, there are important caveats to our model. First, our representation of confidence dynamics is a necessarily phenomenological approach for an important, yet poorly understood mechanism. A key empirical challenge for future work is to quantify these confidence processes, for example the relative weights of the different factors in equations **(1)-(3)**. Second, we assume individual and system fundamentals are fully transparent to all in the system. We thus potentially neglect the effects of *uncertainty*, for example over the actual health of counterparties or the extent of system-wide liquidity hoarding. Nonetheless, given that such uncertainty would intensify with deteriorating fundamentals and confidence, we would expect these effects to accentuate the dynamics we have explored. Third, it is also important to consider how the vulnerabilities in financial systems modelled in this paper emerge, and potentially grow, over time. Fourth, while this work has concentrated on capital ratios, liquidity requirements are potentially also an important policy response. Future work may seek to treat these more systematically, for example by considering the impact of ‘haircut shocks’, which can exogenously generate liquidity shortfalls (Gai *et al* (2011); Gorton and Metrick (2012)). Similarly, effective resolution mechanisms could also enhance the resilience of the system to the risks posed by systemically important institutions. Finally, our work contributes towards the identification of indicators for the systemic importance of institutions by exploring the role of size and interconnectedness, as correlated attributes. Future refinements of this approach may seek to examine more closely the separate effects of these two important factors.

To conclude, market confidence and unprecedented interconnectivity make for far-reaching, complex dynamics in modern financial systems. Any attempt at regulating on a systemic scale can, therefore, only benefit from a deeper understanding of these dynamics. Simple dynamical models, for all their limitations, can offer a valuable starting point for guiding such essential insights.

Appendix 1: Details of the model

Network details

We represent the lending network as a directed graph, taking each individual loan to be the same size; we allow multiple loans in both directions between any given pair of banks, which are then aggregated to give the total bilateral exposure between pairs of banks. Small banks have an in- and out-degree (equivalently, numbers of interbank borrowing and lending links) drawn from a Poisson distribution with mean z_s (we typically take $z_S = 5$), while big banks have mean λz_S (to begin we take $\lambda = 24$). Thus big banks have systematically higher levels of interbank exposure than small banks. In the baseline model, we assume a purely random structure for the interbank lending network; in Appendix 2, however, we show results allowing for 'preferential lending' between banks of different sizes.

For external assets, we define a sharing scheme by assuming that there exists a fixed number Γ of distinct asset classes. Of these, big banks hold n_B distinct asset classes each, in equal value, and small banks hold n_S asset classes each. In the baseline setup we assume simply that $n_B = n_S = 10$. Moreover, every asset class is held in common by g banks (typically 10). This implies that $\Gamma = (N_B n_B + N_S n_S)/g$, where N_B, N_S are the number of big and small banks respectively.

Developing the baseline setup to allow for diversification (Figure 6, panels C and D), we write: $n_B = n_S\, \lambda^x$, such that $x = 0$ recovers the baseline scenario (all banks being equally diversified), and $x > 0$ yields $n_B > n_S$ so that larger banks are diversified over a wider range of asset classes than small banks (this also allows for the possibility that some asset classes are only held by large banks.). Note that, with other parameters fixed, specifying x fixes n_B and n_S. In particular, n_B, n_S are respectively increasing and decreasing functions of x. Where any of these calculations imply non-integer values for n_B, n_S, etc, we adopt the nearest non-integer value.

Health expressions

Hoarding is driven by health and confidence effects, as outlined in the main text, with an alternative formulation presented in Appendix 2. In particular, to calculate an individual bank's liquidity position, m_i, in equation **(1)**, we assume that both short-term interbank loans and liquid assets l_i are available on demand to meet funding outflows. An expression for m_i is thus:

$$m_i = \min\left[1, \frac{A_i^{ST} + l_i}{L_i^{ST}}\right]$$

where L_i^{ST} is the total value of the bank's short-term interbank liabilities; A_i^{ST} is the total value of the bank's short-term interbank assets; and l_i is the amount of liquid assets held by the bank.

Initiating shocks

With the system thus constructed, we initiate contagion by either: (i) an idiosyncratic capital shock applied to a single, randomly selected bank (Figures 3 and 5) or (ii) an aggregate shock, simultaneously affecting the capital positions of all *g* banks holding a randomly selected ('distressed') asset (Figure 6). We then calculate the system-level confidence and individual 'healths', *C* and h_i, to propagate shocks in discrete time as follows.

Shock propagation and bank failure

A bank fails for solvency reasons if shocks to its assets erode its capital buffer to zero. A bank may also fail for liquidity reasons: that is, having insufficient cash to meet immediate obligations. We describe here how both of these may occur with reference to the propagation of: (i) lending withdrawals ('liquidity hoarding'), (ii) asset price contagion, and (iii) counterparty defaults.

(i) Lending withdrawals: Banks may cut lending in the interbank market as a result of effects acting through *C* and h_i. When a bank withdraws lending from a debtor bank *j*, this amounts to the loss of short-term borrowing on the liability side of bank *j*. In such an event, bank *j* must settle these claims immediately and – if liquid assets are insufficient – will raise cash by withdrawing a necessary and sufficient amount of its own, short-term lending, A_j^{ST} (asset side, Figure 1). If, however, even these measures are insufficient, bank *j* is assumed to undergo liquidity default, given the difficulty in realizing the full value of external assets quickly. Thus, for example, if $m_j < 1$ in equation **(1)**, bank *j* is vulnerable to liquidity default if all of its interbank creditors withdraw funding simultaneously (a wholesale 'bank run', which would be prompted as $h_j \rightarrow 0$ in equation **(3)**). In Appendix 2, we extend this framework to allow for banks selling their external assets to try to avoid liquidity default.

(ii) Asset price contagion: We assume that, by liquidating its external assets, a failing bank may depress the price of those assets in the market. This causes the capital position of other banks holding these same assets to be eroded. Previous work (Cifuentes *et al* (2005); Nier *et al* (2007); Gai and Kapadia (2010); May and Arinaminpathy (2010)) has modelled this process by assuming that the price of asset *i* is diminished to a fraction $\exp(-\alpha x_i)$ of its original value, where x_i is the proportion of that asset being sold by the failing bank, and $\alpha > 0$ is a constant denoting systemic 'market illiquidity'. In our baseline model, market illiquidity is also linked directly to confidence effects, writing $\alpha = 1 - C$, though we relax this assumption in Appendix 2. Thus confidence effects on asset prices are mild or negligible when $C = 1$, but become more severe as C declines. Note that for simplicity, we suppose that assets are sold to outside investors rather than to other banks in the network, and also make the conservative assumption that the prices of different assets are uncorrelated, aside from their common dependence on C – allowing for the increase in asset correlations which typically occurs during crises would amplify our results.

(iii) Counterparty defaults: These propagate through the interbank network in the opposite direction to liquidity hoarding, that is from borrower to lender. Suppose that a bank *i* suffers a shock of size S exceeding its capital γ_i, causing insolvency. Its z creditors will then suffer a loss on their lending to this bank: in particular, neglecting bankruptcy costs, we assume that they each receive an asset-side shock of size $(S - \gamma_i)/z$, bounded above by the total size of their exposure to the failing bank. This erodes their own capital position and can cause cascades of capital defaults (May and Arinaminpathy (2010)).

Appendix 2: Extensions to the model

In what follows we present extensions of the model formulation described in the main text and Appendix 1.

1. Alternative form for confidence expression, C

As noted in Appendix 1, we have adopted a simple form for system confidence C that captures liquidity health at a given time by the proportion of interbank loans in the system that have not been withdrawn. However, systemic 'shortening' of lending maturities – even in the absence of any outright withdrawals – may also play a significant role in system health by making banks more vulnerable to subsequent liquidity outflows. To incorporate this, we use an alternative form:

$$C = ½ (E + V)A,$$

where, at a given time, V is the proportion of initially long-term interbank lending that remains long-term. As in the main text, A is the sum of all remaining assets in the system (at the current market price) as a proportion of its initial value, while E is the fraction of interbank loans not withdrawn.

Thus, if all loans are shortened (ie $V = 0$) and none withdrawn (ie $E = 1$), then $C = 0.5A$, while if all loans have also been withdrawn (ie additionally $E = 0$), then $C = 0$. Figure 7 (panel A) repeats the simulations presented in Figure 5, illustrating qualitatively similar results. Figure 7 (panel B) additionally compares this framework with the 'baseline' model from the main text, plotting (as in Figure 3) the extent of contagion when the value of a single asset class of a randomly selected big bank is set to zero. As loan shortening precedes outright withdrawal, system-level confidence deteriorates more quickly when it also depends on the extent of shortening, and the amount of contagion is concomitantly increased.

2. Interbank lending network: 'preferential mixing' between banks

Here we relax the assumption that the interbank lending network is a random web. A study of the Fedwire payments system (National Research Council (2007); May *et al* (2008)) suggests that the financial lending network may in fact be disassortative, that is with low-degree nodes (equivalently, small banks) tending to connect more to high-degree nodes (ie large banks) than other low-degree nodes. We consider both this and its converse, the ‘assortative’ case.

In particular, write N_B, N_S for the number of large and small banks respectively, and z_B, z_S for the mean number of loans made by large and small banks respectively. We now assume that a proportion P of all loans are made between banks of different size classes: that is, the loans made by big banks to small ones, or vice versa. It is straightforward to show, under the size distribution described in the main text and parameterized by λ, that $P = ½$ corresponds to a random web; $P < ½$ yields an assortative network; and $P > ½$ yields a disassortative one. Figure 8 (panels A and B) show results for the latter two cases, illustrating that essential results shown in Figure 5 are qualitatively unchanged.

3. Non-Poisson bank size distributions

As described in Appendix 1, the number of loans made by a given bank is drawn from a Poisson distribution, whose mean depends on whether the bank is large or small. Here we adopt a different scheme, in particular relaxing the assumption that there are two distinct sizes of banks (as parameterized by λ in the main text). Figure 9 (panel A) shows data relating to the US banking sector, drawn from the Federal Deposit Insurance Corporation (FDIC). The closest fit to a Pareto distribution is shown. Originally applied to the distribution of wealth amongst individuals in a society, this distribution has the cumulative function:

$$P(X < x) = 1 - \left(x_0/x\right)^{\alpha}$$

defined for all $x \geq x_0$, for a given $x_0 > 0$. With the caveat that only three points are available here, the data suggest a Pareto shape parameter α of 0.83. Consistent with the main text, we take a system of 200 banks with the smallest banks making 5 interbank loans (recalling our assumption that this constitutes 20% of their balance sheet – see Table 1). Thus we take $x_0 = 5$ for the interbank ‘degree distribution’. Moreover we assume a cut-off on bank sizes such that the biggest bank is no more than 10^4 times larger than the smallest.

This approach necessitates some minor adaptations of the plotting scheme used in Figure 5. First, for given parameters, the present approach has some variability in the realized composition of bank sizes from one simulation to the next. To accumulate the results of successive simulations on a single plot, therefore, we measure the size of the index bank, not in absolute terms, but as a proportion of the system's initial, total assets. Second, in the present framework, the index bank can be so large as to account for a significant proportion of the total system's assets by itself. Accordingly, we measure 'impact' here as the proportion lost of total initial assets *excluding* the index bank. Similar to Figure 5 in the main text, we select a random index bank (in this case, of any size) and force it to fail by setting its capital to zero, without affecting the remainder of its balance sheet.

Figure 9 (panels B-D) plot results from this approach. Here different combinations of channels are shown on different plots, to provide a better illustration of respective sets of individual simulation outcomes ('dots' in grey), that give rise to the mean impact plotted. The full model (panel B) illustrates strong non-linearities: while index banks smaller than about 15% of the whole system do not initiate contagion, those above this threshold are liable to bring down the whole system. Indeed, the 'intermediate' points in the figure represent outcomes where only a single bank remains, always being the largest bank originally present. Panels C and D explore less extreme cases, where either asset shocks or liquidity hoarding are deactivated. Here it is possible for system failure to be only partial (eg panel D). Nonetheless, the essential result is in accordance with that shown in Figure 5: the effect of index bank collapse scales more than linearly with index bank size.

4. Alternative formulations of asset price contagion

In the main text, banks only liquidate their external assets upon failure, and the price effects of such action increase as system-wide confidence deteriorates. Here we relax these assumptions in turn.

First, we allow banks to sell their external assets prior to failure to try to avoid liquidity default. In particular we assume that, if withdrawals of interbank lending are insufficient to meet funding shortfalls, a bank may additionally sell a necessary and sufficient amount of its external assets, distributed evenly across asset classes in its portfolio. However, if this too is insufficient, the bank undergoes liquidity default. Figure 10 (panel A) shows that allowing for asset sales prior to default yields qualitatively similar results to Figure 5 (panel C) in the main text. And Figure 10 (panel C, blue curve) shows that for the parameter values employed here, this change has a minimal effect on overall outcomes when the value of one of a random large index bank's external asset classes is set to zero. This arises because, while liquidity hoarding in the 'baseline model' plays a significant role in the model dynamics (see eg Figures 3 and 5 (panel C)), only a minority of banks ultimately fail for liquidity reasons; for those for which asset sales could help prevent this, it is likely that they will subsequently fail either for solvency reasons or due to ever-greater funding withdrawals, in instances of systemic collapse. In other words, for the parameters adopted here, avoiding illiquidity tends merely to delay failure. So allowing banks to sell their assets for liquidity reasons prior to failure has only a limited effect on overall outcomes for the system.

Second, we uncouple market liquidity from system confidence C, fixing $\alpha = 0.1$, a comparatively low value for market illiquidity (recalling that $\alpha = 0$ corresponds to a perfectly liquid market). Although the blue curve in Figure 10 (panel B) does not reach the same plateau as in Figure 5 (panel C) (as there is a lower probability of big index banks causing system collapse in this framework), the greater-than-proportional scaling with respect to bank size is still evident. But as Figure 10 (panel C, green curve) confirms, with α fixed at 0.1, contagion is indeed weaker under this assumption than in the baseline model, as would be expected given that one of the amplifying dynamics of the model is switched off.

Tables and Figures

Table 1: Parameters and their default values

Category	**Symbol**	**Meaning**	**Default value**
Global	N	Total number of banks	200
	λ	Bank size disparity	24
	π_S	Initial proportion loans being short term	0.5
Balance sheet	z	Number of outgoing loans, small banks	5
	θ	Proportion of total assets initially in interbank lending	0.2
	l_i	Proportion of bank *i*'s total assets initially liquid	0.01
	γ_i	Bank *i*'s initial capital (to assets) ratio	0.04
Asset class	Γ	Total number of distinct asset classes in the system	200
	g	Number of banks sharing an asset class	10
	n_B	Number of asset classes held by each big bank	10
	n_S	Number of asset classes held by each small bank	10

Figure 1: Schematic, balance sheet representation of a single bank. See Figure 2 for a schematic of how this bank interacts with others in the network.

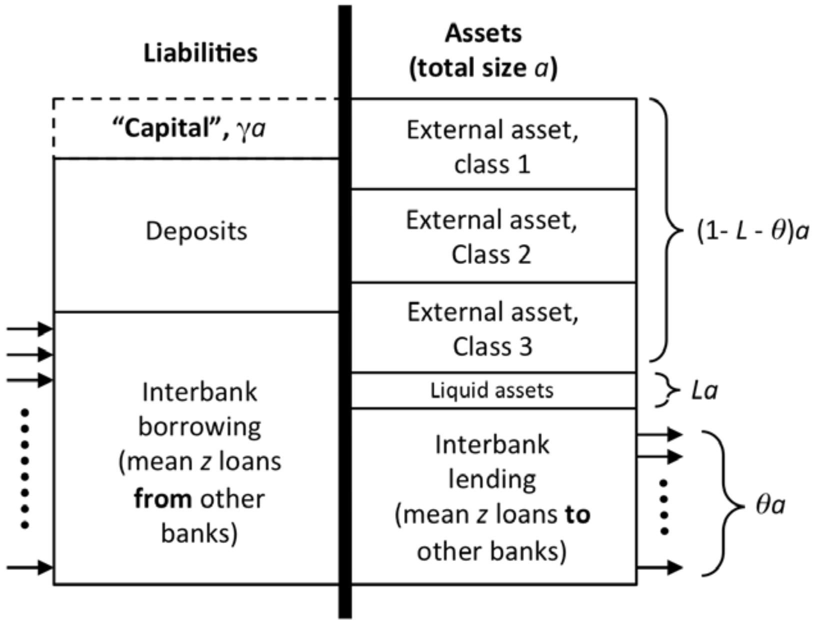

Figure 2: Schematic representation of different modes of connectivity. A schematic illustrating two different modes of connectivity in the model, showing for simplicity the case where all banks have the same size. Arrows represent loans in the interbank network, pointing from lender to borrower. Independently of this network, banks may also hold external assets in common (shown 'hatched' in different patterns). For example, bank 1 holds one asset in common with bank 2, and two assets in common with bank 3.

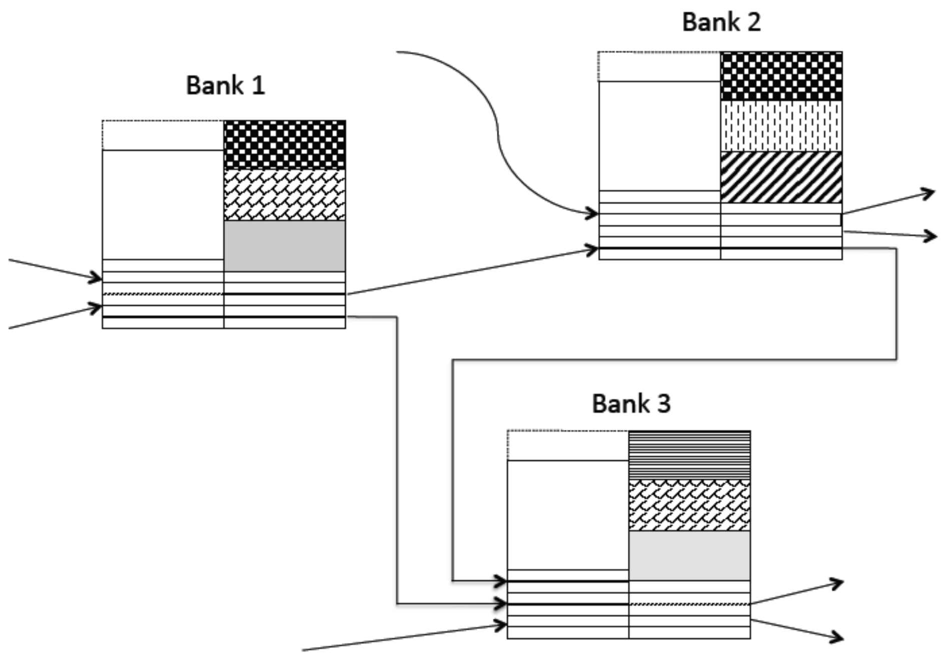

Figure 3: Baseline results and the effect of liquidity hoarding. Although intended as a defensive action (increasing h_i), hoarding imposes negative externalities on other banks in the system.

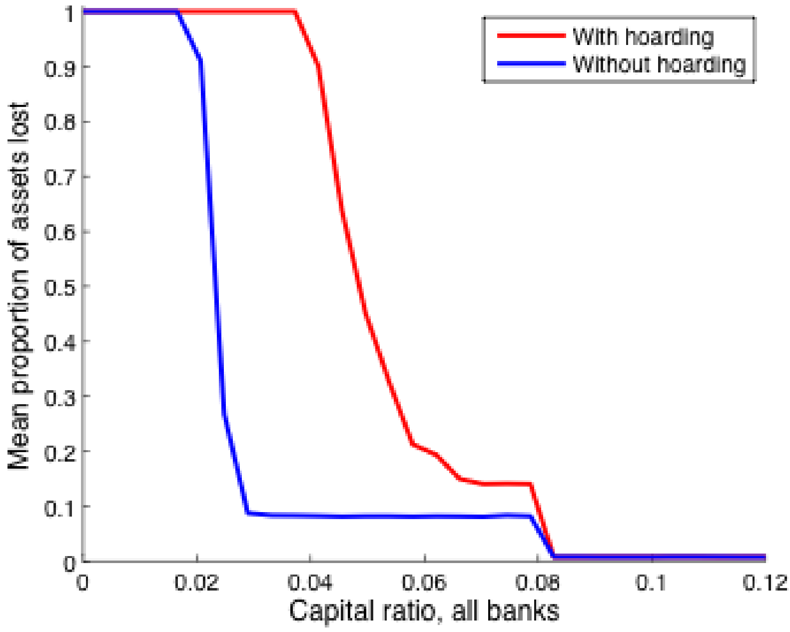

Figure 4: Relating changes in assets to changes in debt and equity.

Panel A: Empirical findings (Kim *et al* (2012)), illustrating that the dynamics of a bank's total assets (x-axis) arise predominantly from changes in debt (blue line) rather than equity (red line). Although shown for the example of Morgan Stanley, qualitatively similar behaviour applies for other banks. Reproduced with permission from Hyun Shin.

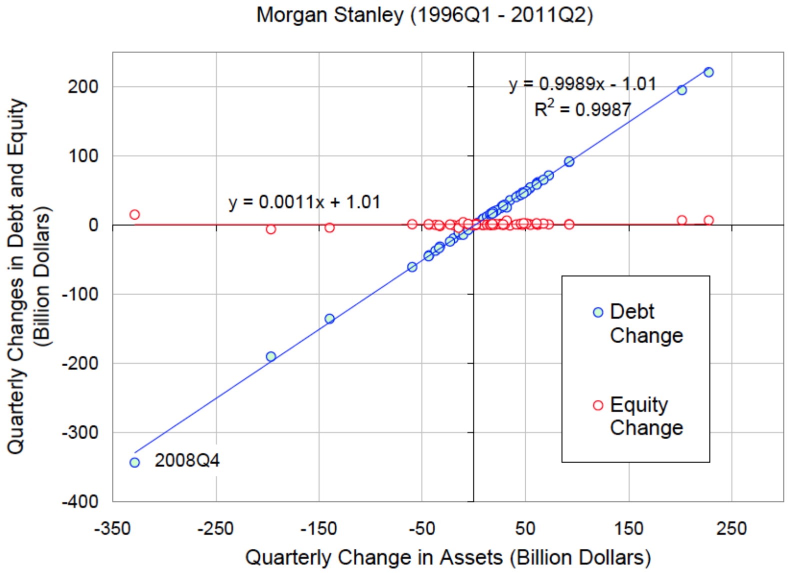

Panel B: Simulation results corresponding to the lower left quadrant of Panel A. Plotted points are accumulated by following the evolution of a random bank's balance sheet through a simulation, repeating, and superimposing all results over 50 simulations. Note that both axes are negative since we simulate here only the post-shock, contractionary phase.

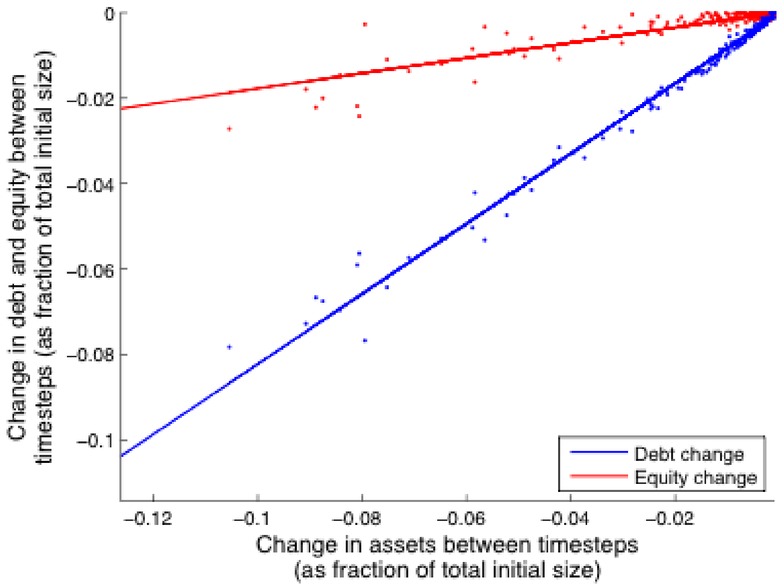

Figure 5: Small and big bank failure in the baseline setup. Frequency distributions for numbers of banks failing in the baseline setup ($\lambda = 24$), following the failure of a single index bank of given size. Comparing distributions from: (A) small index bank collapse, and (B) big index bank collapse. Panel C: comparing the effects of large bank collapse, relative to those of small bank collapse, as a function of λ. The upper (blue) curve reaches a plateau at the maximum possible value for a finite system with 200 banks. See text for details.

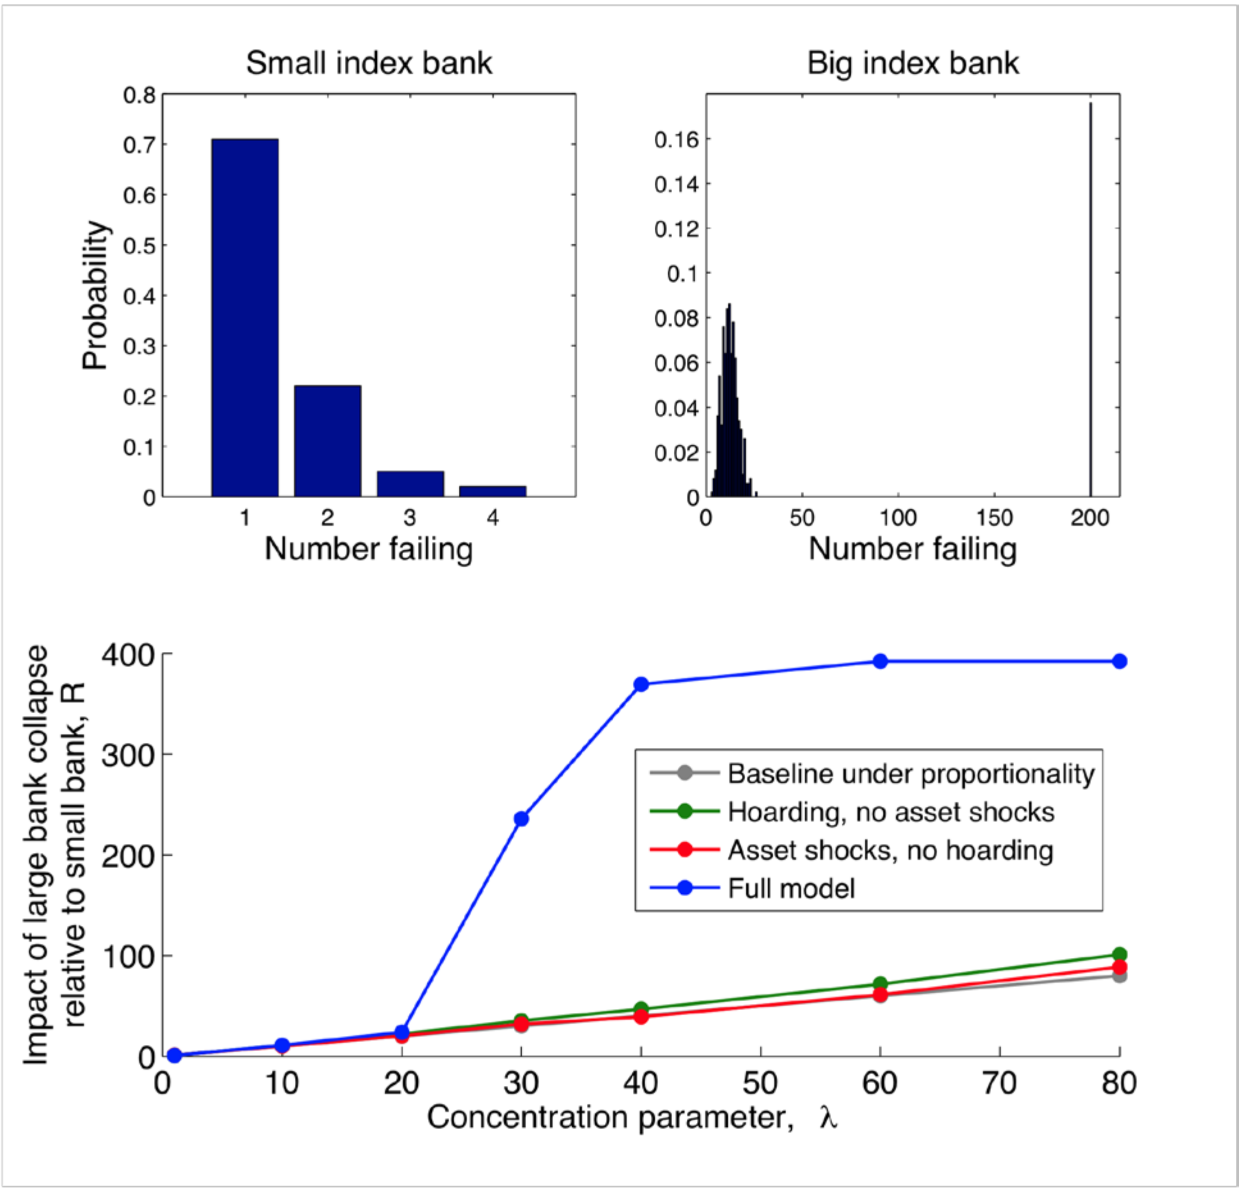

Figure 6: Exploring the relative importance of big and small bank capital ratios, under different scenarios. Here, γ_B, γ_S represent capital ratios of big and small banks respectively. As described in the text, n_B, n_S are the numbers of asset classes held by each big and small bank respectively, and λ is the 'concentration parameter', denoting both the relative size and number of big and small banks. (A) λ = 24, giving 8 big banks and 192 small ones (B) Big banks are 250 times the size of small ones, but 70 times less numerous, giving 3 big banks and 197 small ones. (C) λ = 24, with big banks having more diversified assets. Here n_B, n_S denote the number of asset classes held by big and small banks, respectively. (D) Repeating Figure 5 (panel C), in the case where relative diversification (n_B/n_S) scales with λ as indicated. (See Appendix 1 for details.)

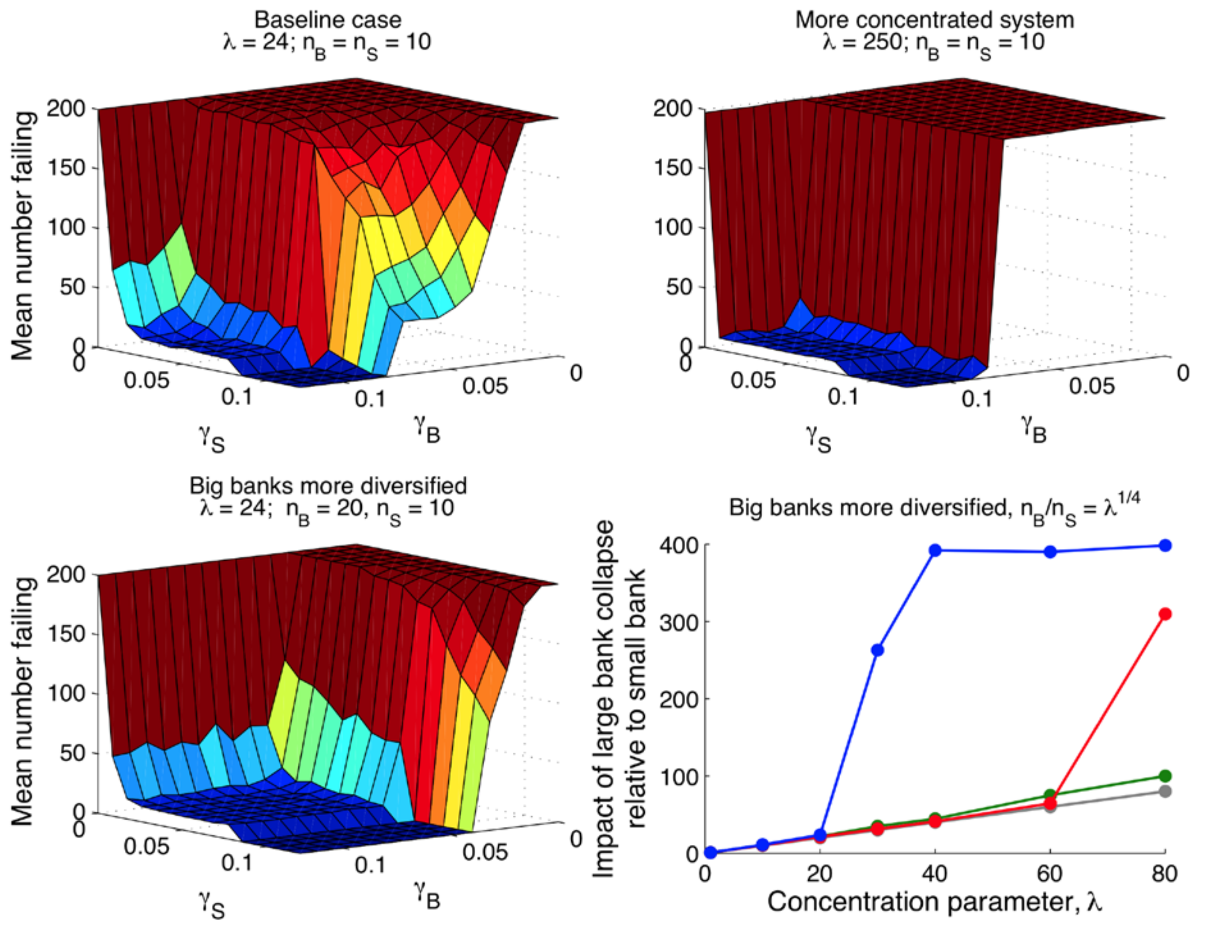

Figure 7: Model results under alternative form for confidence expression. Panel A repeats Figure 5 (panel C) under an alternative expression for system confidence (see Appendix 2, Section 1). Panel B compares the extent of contagion under this specification with the 'baseline' model in the main text – in analogy to Figure 3, the initiating shock is applied to a randomly selected large bank, setting the value of one of its external asset classes to zero without affecting other banks holding the same asset class.

Panel A

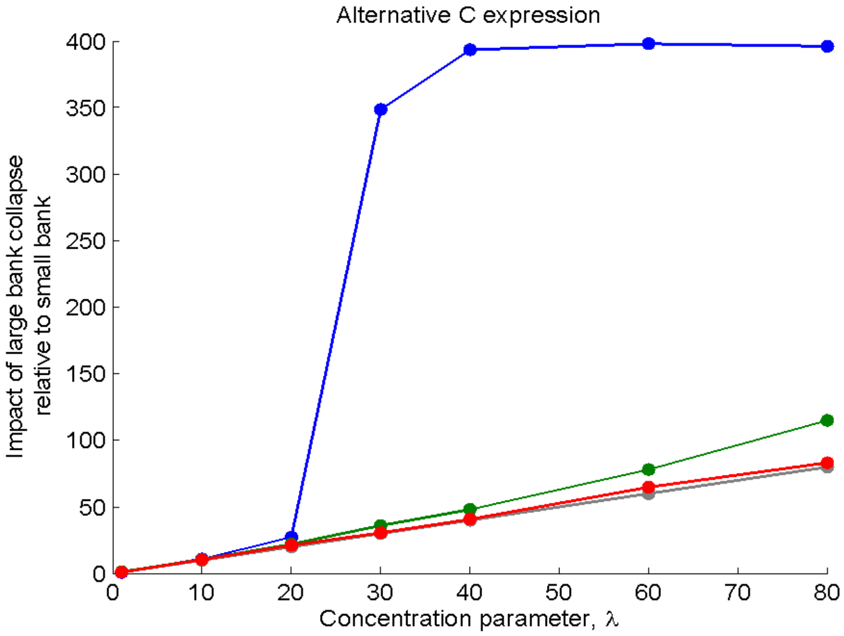

Panel B

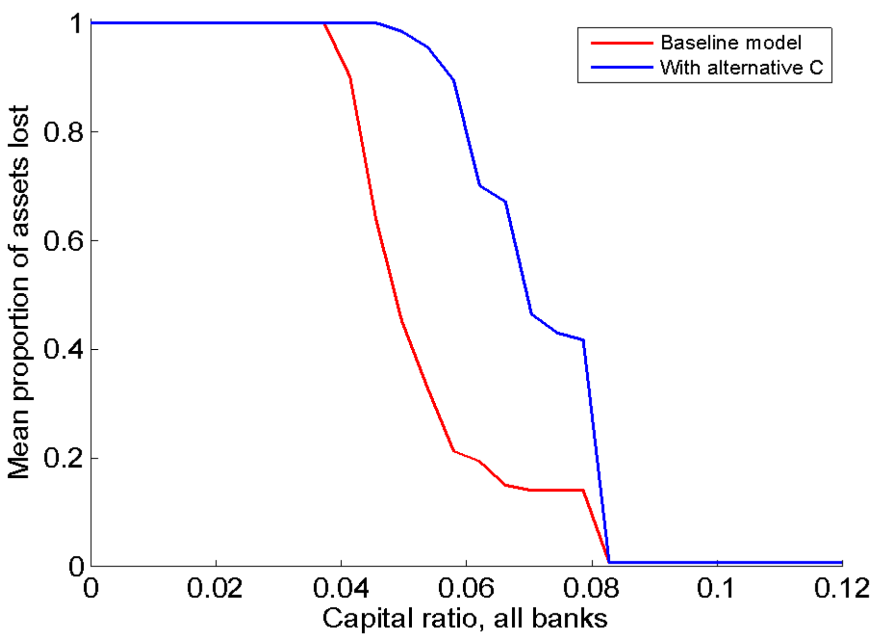

Figure 8: Model results when allowing a non-random lending network between big and small banks (see Appendix 2, Section 2)

Panel A: assortative network

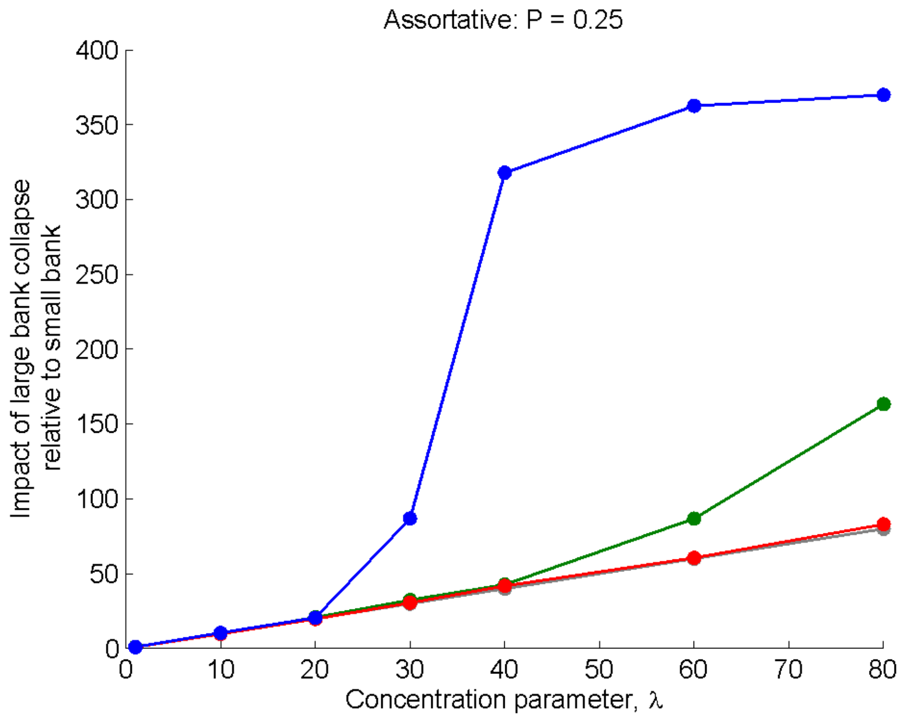

Panel B: dissassortative network

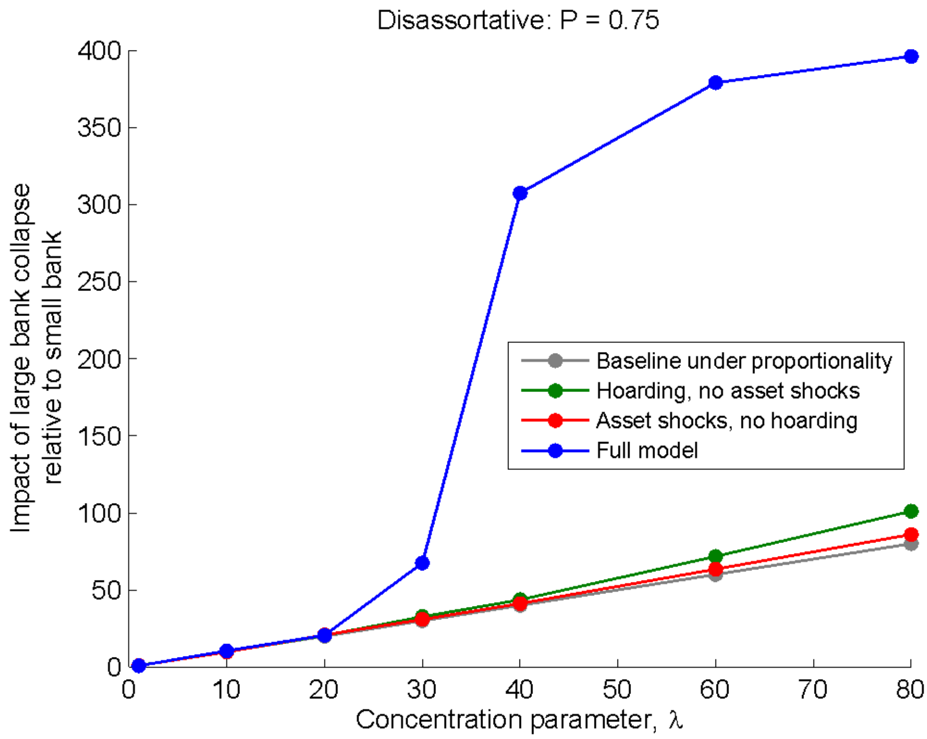

Figure 9: Model results in the case of a Pareto distribution for bank sizes (see Appendix 2, Section 3). (A) Cumulative distribution of bank sizes in the US banking sector (FDIC), and its closest fit with a Pareto distribution, implying a 'shape parameter' of 0.83. (B) Impact of index bank collapse according to the full model, as a function of the index bank size. Grey points indicate individual simulation outcomes, while the blue line shows the mean outcome. Note that 'system impact' here refers to the proportion of the initial system's assets, excluding the index bank, that is ultimately lost as a result of index bank collapse. 'Intermediate' points all correspond to the survival of a single large bank, as described in Section 3. (C, D) As for panel B, but individual channels acting alone (as indicated).

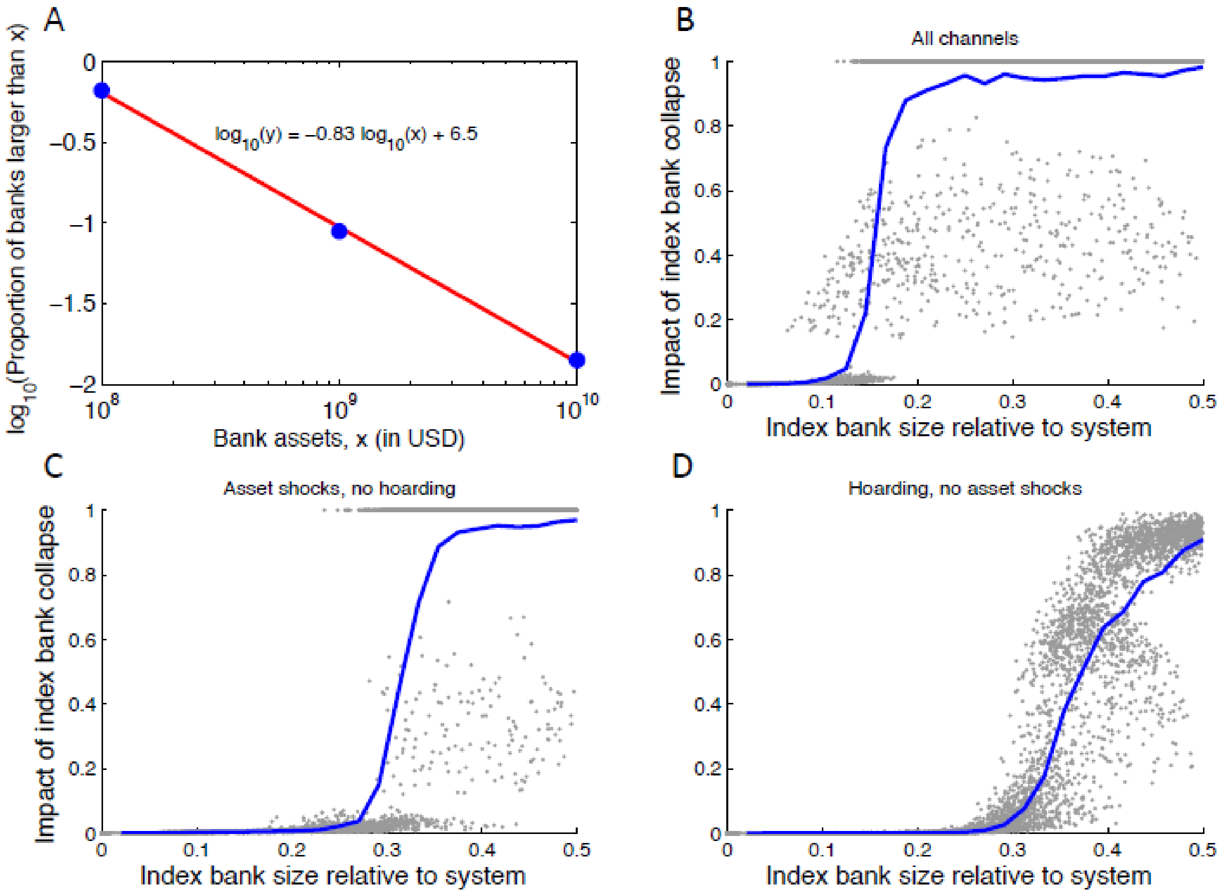

Figure 10: Model results under alternative formulations of asset price contagion (see Appendix 2, Section 4). Repeating Figure 5 (panel C) for the 'full model' (ie including hoarding), in the case of: (A) banks being able to sell external assets to avoid liquidity default, and (B) a 'milder' contagion scenario of constant market liquidity not linked to system confidence, here taking α = 0.1. (C) As for Figure 3, comparing the extent of contagion under different models presented in Section 4 of Appendix 2.

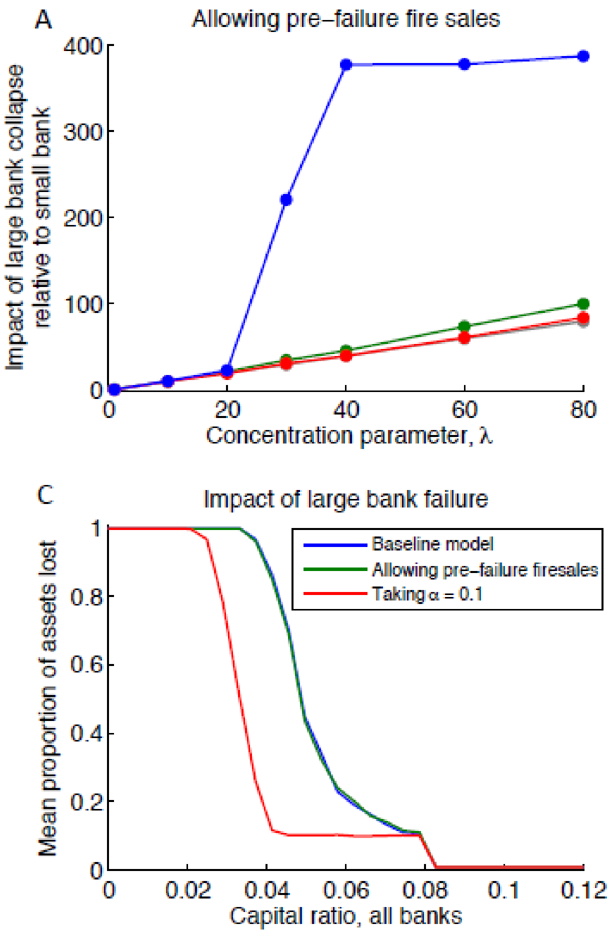

References

Acharya, V and Skeie, D (2011), 'A model of liquidity hoarding and term premia in inter-bank markets', *Journal of Monetary Economics*, Vol. 58(5), pages 436-47.

Adrian, T and Shin, H S (2010), 'Liquidity and leverage', *Journal of Financial Intermediation*, Vol. 19(3), pages 418-37.

Allen, F and Babus, A (2009), 'Networks in finance', in Kleindorfer, P, Wind, Y and Gunther, R (eds), *The network challenge: strategy, profit, and risk in an interlinked world*, Wharton School Publishing.

Anderson, R and May, R (1991), *Infectious diseases of humans: dynamics and control*, Oxford University Press.

Bank for International Settlements (2011), 'Global systemically important banks: assessment methodology and the additional loss absorbency requirement', available from: http://www.bis.org/publ/bcbs207.pdf

Bank of England (2009), *The role of macroprudential policy*. Discussion paper, November.

Bank of England (2011), *Instruments of macroprudential policy*. Discussion paper, December.

Battiston, S, Gatti, D D, Gallegati, M, Greenwald, B C and Stiglitz, J E (2009), 'Liaisons dangereuses: increasing connectivity, risk sharing, and systemic risk', *NBER Working Paper No. 15611*.

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

Brunnermeier, M K (2009), 'Deciphering the 2007-08 liquidity and credit crunch', *Journal of Economic Perspectives*, Vol. 23, pages 77–100.

Brunnermeier, M K and Pedersen, L H (2009), 'Market liquidity and funding liquidity', *Review of Financial Studies*, Vol. 22(6), pages 2,201-38.

Caballero, R J and Krishnamurthy, A (2008), 'Collective risk management in a flight to quality episode', *Journal of Finance*, Vol. 63, pages 2,195–230.

Caccioli, F, Marsili, M and Vivo, P (2009), 'Eroding market stability by proliferation of financial instruments', *The European Physical Journal B-Condensed Matter and Complex Systems*, Vol. 71, pages 467–79.

Cifuentes, R, Ferrucci, G and Shin, H S (2005), 'Liquidity risk and contagion', *Bank of England Working Paper No. 264*.

Coval, J D and Stafford, E (2007), 'Asset fire sales (and purchases) in equity markets', *Journal of Financial Economics*, Vol. 86, pages 479–512.

Drehmann, M and Tarashev, N (2011), 'Systemic importance: some simple indicators', *BIS Quarterly Review*, available from: http://www.bis.org/publ/qtrpdf/r_qt1103e.htm.

FDIC (2011) ‘Quarterly Banking Profile, First Quarter 2011’, available from: http://www2.fdic.gov/qbp/2011mar/qbp.pdf.

Gai, P, Haldane, A and Kapadia, S (2011), ‘Complexity, concentration and contagion’, *Journal of Monetary Economics*, Vol. 58(5), pages 453-70.

Gai, P and Kapadia, S (2010), ‘Contagion in financial networks’, *Proceedings of the Royal Society A*, Vol. 466, No. 2120, pages 2,401-23.

Gorton, G B and Metrick, A (2012), ‘Securitized banking and the run on repo’, *Journal of Financial Economics*, Vol. 104(3), pages 425-51.

Haldane, A G and May, R M (2011a), ‘Systemic risk in banking ecosystems’, *Nature*, Vol. 469, pages 351–55.

Haldane, A G and May, R M (2011b), ‘The birds and the bees, and the big banks’, *Financial Times*, available from: http://www.ft.com/intl/cms/s/0/5c7fa72e-3d20-11e0-bbff-00144feabdc0.html#axzz28Fr58fK4 (accessed 3 Aug 2012).

Independent Commission on Banking (2011), *Final report*, September.

Ivashina, V and Scharfstein, D S (2010), ‘Bank lending during the financial crisis of 2008’, *Journal of Financial Economics*, Vol. 97(3) pages 319-38.

Jackson, M O (2008), *Social and economic networks*, Princeton University Press.

Kapadia, S, Drehmann, M, Elliott, J and Sterne, G. (2012), ‘Liquidity risk, cash flow constraints, and systemic feedbacks’. Forthcoming in Haubrich, J G and Lo, A W (eds), *Quantifying Systemic Risk*, University of Chicago Press.

Kareiva, P M and Levin, S A (2003), *The importance of species: perspectives on expendability and triage*, Princeton University Press.

Kim, H G, Shin, H S and Yun, J (2012), ‘Monetary aggregates and the Central Bank's financial stability mandate’, paper presented at Federal Reserve conference 23-24 March 2012, available from: http://www.princeton.edu/~hsshin/www/monetary_aggregates.pdf.

Kirman, A (2010), *Complex economics: individual and collective rationality (The Graz Schumpeter Lectures)*, Routledge.

Lloyd-Smith, J O, Schreiber, S J, Kopp, P E and Getz, W M (2005), ‘Superspreading and the effect of individual variation on disease emergence’, *Nature*, Vol. 438, pages 355–59.

Marsili, M, Kirman, A and Anand, K (2010), ‘Epidemics of rules, information aggregation failure and market crashes’, mimeo, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1624803.

May, R M and Arinaminpathy, N (2010), ‘Systemic risk: the dynamics of model banking systems’, *Journal of the Royal Society Interface*, Vol. 7, pages 823-38.

May, R M, Levin, S A and Sugihara, G (2008), ‘Ecology for bankers’, *Nature*, Vol. 451, pages 893–95.

National Research Council (2007), 'New directions for understanding systemic risk: a report on a conference cosponsored by the Federal Reserve Bank of New York and the National Academy of Sciences'.

Nier, E, Yang, J, Yorulmazer, T and Alentorn, A (2007), 'Network models and financial stability', *Journal of Economic Dynamics and Control*, Vol. 31, pages 2,033–60.

Paine, R T (1966), 'Food web complexity and species diversity', *American Naturalist*, Vol. 100, pages 65–75.

Soramaki, K, Bech, M L, Arnold, J, Glass, R J and Beyeler, W E (2007), 'The topology of interbank payment flows', *Physica A: Statistical Mechanics and its Applications*, Vol. 379, pages 317–33.

Strogatz, S H (2001), 'Exploring complex networks', *Nature*, Vol. 410, pages 268–76.

Switzerland State Secretariat for International Financial Matters (2010), 'Final report of the Commission of Experts for limiting the economic risks posed by large companies', available at: http://www.sif.admin.ch/dokumentation/00514/00519/00592/index.html?lang=en.

Upper, C (2011), 'Simulation methods to assess the danger of contagion in interbank markets', *Journal of Financial Stability*, Vol. 7(3), pages 111-25.

Wagner, W (2011), 'Systemic liquidation risk and the diversity–diversification trade-off', *Journal of Finance*, Vol. 66, pages 1,141–75.