BANK OF ENGLAND

Staff Working Paper No. 909

Solvency distress contagion risk: network structure, bank heterogeneity and systemic resilience

Kumushoy Abduraimova and Paul Nahai-Williamson

February 2021

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.

BANK OF ENGLAND

Staff Working Paper No. 909

Solvency distress contagion risk: network structure, bank heterogeneity and systemic resilience

Kumushoy Abduraimova[1] and Paul Nahai-Williamson[2]

Abstract

We systematically analyse how network structure and bank characteristics affect solvency distress contagion risk in interbank networks. As interbank networks become more connected and more regular in structure, the contagion risk of system-wide shocks and individual bank defaults initially increases and then decreases, all else being equal. The low density heterogeneous network structures that typify real interbank networks are particularly prone to solvency distress contagion risk, when banks are similar in balance sheet size and capitalisation. However, when networks are calibrated to UK data, the higher capitalisation of large, highly-connected banks relative to their interbank exposures significantly increases the resilience of the system and reduces the importance of network structure. These findings reinforce the importance and effectiveness of imposing higher capital buffers on systemically important banks and of policies that limit interbank exposures. We also demonstrate that for real-world interbank networks, simple network metrics other than individual bank connectedness do not provide robust indicators for monitoring solvency contagion risk, suggesting that policymakers should continue efforts to model these risks explicitly rather than rely on simple aggregate indicators.

Key words: Financial networks, systemic risk, financial contagion, banking policy.

JEL classification: C54, D85, G01, G21, G28.

(1) Imperial College Business School and Bank of England. Email: kumushoy.abduraimova11@imperial.ac.uk
(2) Bank of England. Email: paul.nahai-williamson@bankofengland.co.uk

Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee. The authors are grateful to Marco Bardoscia, Paolo Barucca and Anne-Caroline Hüser for their comments, and to seminar and session participants at the Department of Finance at Imperial College Business School, the Bank of England, CIRANO Montreal 2017 conference and the IFABS Oxford 2017 conference.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email enquiries@bankofengland.co.uk

© **Bank of England 2021**
ISSN 1749-9135 (on-line)

1 Introduction

One of the lasting lessons of the global financial crisis – reinforced by the "dash for cash" in March 2020 – is the importance of addressing risks to the financial system as a whole, rather than solely focusing on individual institutions' risks. The crisis revealed that we had limited knowledge of the structure of the financial network, and of the ways in which this structure interacts with contagion mechanisms and systemic risk.

There are a number of amplification mechanisms that are relevant to systemic risk, which can operate along a number of different financial networks. For example, these include amplification risks arising from indirect exposures to the same tradable securities (the bipartite asset-institution networks); and liquidity risks in short-term lending markets (funding networks). The risks associated with each mechanism interact both with the network structure relevant to that mechanism and to the other mechanisms and financial networks. To understand the way in which network structure contributes to systemic risk, it is important both to explore this relationship for individual mechanisms, and how different networks interact.

In this paper, we examine in detail the role of network topology and institutional characteristics in making the financial network more or less resilient to the risk of solvency distress contagion in particular. Solvency distress contagion occurs when, in the presence of uncertainty around banks' creditworthiness, credit risk is incorporated into the valuation of interbank claims. As a bank's solvency deteriorates, the value of claims on that bank falls, spreading distress to banks holding those claims. In this way, solvency distress amplifies and spreads throughout the interbank network.

We model solvency distress contagion risk using the model of Bardoscia *et al.* (2019), a realistic model of the contagion mechanism that has been used in regulatory stress testing (see e.g. Bank of England (2016)). We simulate a large number of networks (real-world, random and regular) and analyse how particular network patterns and characteristics relate to solvency distress contagion risk under various shock scenarios, including system-wide shocks and individual bank defaults. We further consider how the characteristics of the banks in these networks impact the level of amplification risk - focusing on the impact of heterogeneous capital and exposure levels. We then use these results to investigate econometrically whether there are summary network metrics that could be used by policymakers to monitor solvency contagion risk in interbank networks.

This study makes several contributions to the literature. First, we find that while the core-periphery structure of the interbank exposure network is particularly prone to solvency distress contagion risk, this risk is mitigated in the real network of heterogeneous banks due to more capital being held on average relative to interbank exposures by core banks. This supports the policy of setting extra capital buffers for systemically important banks as a way to mitigate systemic risk via the solvency contagion channel. This finding is consistent with and expands upon previous studies that have identified the importance of capital levels for the system as a whole in reducing *default* contagion risk (in which solvency shocks propagate upon counterparty default, see e.g. Nier *et al.* (2008)).

Second, solvency distress contagion risk is sensitive to the level of exposures in the system - larger total exposures relative to capital both increase the magnitude of solvency conta-

gion risk, and make it less sensitive to network structure. This supports the application of large exposure limits, and the importance of monitoring interbank exposures (consistent with the findings of Nier *et al.* (2008); Halaj & Kok (2013); Georg (2013) for default contagion risk).

Further, the results of our econometric analysis can inform regulators on the extent to which network measures can reliably be used for identifying and monitoring systemic risk. For networks consisting of banks of homogeneous size and with homogeneous capital levels, several network metrics are informative for monitoring systemic risk – with system-wide measures of connectedness (degree) and concentration of links among a subset of banks (assortativity, k-nearest neighbour) being statistically significant in explaining contagion outcomes for networks of different densities (consistent with the findings of Hurd *et al.* (2017) on the importance of assortativity for default cascade risk). For more realistic networks however – featuring real-world heterogeneity in bank balance sheets – these simple network metrics become unreliable (statistically insignificant) indicators of solvency contagion risk. At an individual bank level, a simple measure of a bank's connectedness to the network (its degree) does however correlate to the contagion risk it poses to the network in this more realistic case – consistent with the use of a bank's financial interconnectedness as one metric for identifying global systemic importance. These results point to the importance of regulators using more sophisticated models and model-based metrics to monitor the evolution of contagion risk in interbank networks, where such metrics should incorporate information on banks' resilience – in absolute terms and relative to the size of interbank exposures – in addition to information on the network structure.

2 Literature review

There is a large literature investigating the interaction between contagion risk and network structure for interbank networks, across different types of contagion mechanism, different shocks, different network topologies and connectedness, and focusing on different institutional characteristics. Here, we summarise how different aspects of network and institutional structure may affect contagion risk; and how our paper contributes to the literature along these dimensions.

Existing research primarily focuses on *default* contagion – where a bank's default causes solvency and/or liquidity impacts on its counterparties.[1]

Several papers find that more connectivity can (ultimately) reduce contagion risk. Nier *et al.* (2008) and Elliott *et al.* (2014) find that as the interbank network becomes more connected, the risk of default cascades initially increases but then falls back. Increasing the connectivity increases contagion risk by allowing the shocks to propagate to more nodes; but after reaching a particular threshold, further increases reduce widespread con-

[1] Eisenberg & Noe (2001) is a classical model of default cascades and is the foundation for much of the later work. In the model there is an exogenous shock to outside assets of a bank that reduces its net worth. If the shock magnitude is sufficiently large, the bank defaults and the remaining assets are distributed to its creditors proportionally to nominal amounts of their claims. The reduction in payments to defaulting bank's creditors might in turn cause some of the creditors' default, and so on, leading to default cascades.

tagion likelihood as banks become diversified to the extent where they are not heavily exposed to a particular counterparty, and can absorb losses without going into default themselves. Freixas *et al.* (2000) meanwhile show that complete (diversified) networks are more robust to funding contagion than ring (credit chain) networks because in a more diversified system an insolvent bank can transfer a smaller fraction of its losses to each of its lenders and does not trigger further defaults. While these studies focus on default shocks, Cont *et al.* (2013) consider a common market shock affecting several banks simultaneously and then leading to contagion. They similarly find that – for a well-capitalised banking system – higher connectivity firstly leads to a higher extent of contagion, but after a certain point a further increase of connectivity reduces the propagation by diversifying across the linkages. However, in the case of an *under-capitalised* banking network, higher connectivity always implies more severe contagion.

This latter result illustrates the importance of institutional characteristics, and is consistent with the findings of Nier *et al.* (2008) that higher capitalisation reduces default contagion risk. Nier *et al.* (2008) and Halaj & Kok (2013) also find the (complementary) results that lower exposures reduce contagion risk; consistent with the findings of Georg (2013) that as interbank lending volumes increase, random networks become unstable.

The conditionality of contagion risk on bank capitalisation and exposure sizes also emphasises the wider point that contagion risk is conditional on a number of factors beyond simple network structure. This includes the type of shock being considered, and is consistent with the identification of the financial system as having a *robust-yet-fragile* property. Gai & Kapadia (2010) find that contagion may be unlikely, but *contagious* defaults – a shock at "a particular pressure point" – can lead to widespread problems. The funding runs model of Allen & Gale (2000) similarly shows, for stylized networks, that while a more interconnected network is in general more resilient to liquidity shocks, sufficiently large shocks result in contagion. Acemoglu *et al.* (2015) also differentiate between a *small shock regime* and a *large shock regime*; with contagion risk decreasing with increasing connectedness in the former, but both complete and ring networks being fragile in the latter case.

Reflecting this property, Elliott & Hazell (2016) conclude that the socially optimal network is the one in which institutions form clusters such that connections within the clusters are very strong but inter-cluster links are weak. This arrangement allows sharing the risk of small shocks within the clusters; but means that if there is a large shock to a particular bank that causes defaults of other banks within its cluster also default, the weak inter-cluster links work as firewalls and prevent propagation to other groups.

The relationship between network structure and contagion has attracted attention of empirical studies as well. Most of the early ones are based on partial information on the interbank network. For instance, Furfine (2003) only observes a segment of the market (US federal funds market). Several other studies deal with the absence of the bilateral exposures data by reconstructing the interbank network based on the balance sheet information using the *maximum entropy* (ME) approach: Upper & Worms (2004) analyse the German market, Wells (2004) and Elsinger *et al.* (2006) look at the Austrian market, and Degryse & Nguyen (2007) consider the Belgian market. Mistrulli (2011) uses actual data on bilateral exposures in the Italian interbank market and concludes that the max-

imum entropy methodology leads to underestimation of contagion risk. This estimation approach implies homogeneous exposures, but a heterogeneous degree distribution is one of the stylized facts of banking networks. It has been widely found in the literature that a core-periphery structure best describes the financial networks: Boss *et al.* (2004b) for the Austrian data, in 't Veld & van Lelyveld (2014) and Craig & Von Peter (2014) for the Dutch data, and Langfield *et al.* (2014) for the UK data.

While most of the existing papers focus on default contagion, more recent work has modelled solvency *distress* contagion (Battiston *et al.* (2012)) - which recognises that the value of interbank assets can fall, propagating distress, before banks default. Barucca *et al.* (2016) present a general framework for modelling such contagion in a network of interbank exposures; and Bardoscia *et al.* (2019) adapt this approach to produce a model that captures the risks of solvency distress contagion in a network of interbank debt claims, using a structural model to revalue those claims. In this framework, counterparty losses increase non-linearly as solvency falls, more closely reflecting reality. For this reason, their model has been used in regulatory stress testing (see e.g. Bank of England (2016)).

We contribute to the literature in several ways. First, we analyse the interaction between network structure, institutional characteristics and contagion risk for this more realistic solvency *distress* contagion mechanism, using the model of Bardoscia *et al.* (2019)). Second, we systematically analyse how contagion varies both as a function of network *connectedness* – the density of the network – and as a function of network *structure* – whether a network has a regular, random, scale-free or core-periphery structure. We do so for both idiosyncratic and system-wide shocks. Third, we consider how heterogeneity of institutional balance sheets impacts on contagion risk, by calibrating our network to historical data on the UK interbank market.

Bringing these results together, we show that for solvency distress contagion, the core-periphery network structure with a connectedness in the range of that present in the UK interbank network is particularly prone to solvency distress contagion, when considering similarly sized banks with similar capital levels. For idiosyncratic shocks, it demonstrates the same characteristic of contagion risk initially increasing and then decreasing with increasing connectedness identified by Nier *et al.* (2008) and Elliott *et al.* (2014). But we find that modelling real-world heterogeneity in banks' balance sheets significantly reduces contagion risk, under both system-wide and idiosyncratic shocks. This reflects the higher levels of capital relative to interbank exposures held by core banks, and supports the application of systemic risk buffers to highly interconnected banks.

We further confirm that for larger exposures and lower levels of capital, solvency distress contagion risk becomes less sensitive to network structure – consistent with the identification of 'large-shock regimes' and a 'robust-yet-fragile' characteristic by Acemoglu *et al.* (2015); Gai & Kapadia (2010); Cont *et al.* (2013); Allen & Gale (2000), and the findings of Nier *et al.* (2008); Georg (2013); Halaj & Kok (2013). This emphasises the importance and effectiveness of policies that increase the resilience of individual institutions, and reduce the gross magnitude of risk in the interbank network, without specifically targeting network structure.

Finally, we use the results of the contagion model for our various networks to evaluate the practical usefulness of a number of standard network metrics for monitoring solvency

contagion risk, in order to assess whether such metrics could form part of the regulatory monitoring toolkit. Several authors have developed indicators that capture risk in interbank networks, which generally incorporate information on the the resilience of the nodes within the network (rather than simply, or primarily, the network structure). Glasserman & Young (2015) for example develop a contagion index based on institutional characteristics – a bank's net worth, its outside leverage and its financial connectivity – which determines the relative likelihood that a bank will make others fail through contagion. Bardoscia *et al.* (2017b) meanwhile show that the largest eigenvalue of the interbank leverage matrix can be used to assess network resilience.

Consistent with the results of our contagion analysis, we find that for a core-periphery network of similar banks, metrics that reflect risks associated with concentration are informative for monitoring systemic risk - consistent for example with the work of Hurd *et al.* (2017), who demonstrate that assortativity has a strong effect on systemic risk measured in a default cascade setting. Specifically, for system-wide shocks we find that k-nearest neighbour and assortativity measures are statistically significant in explaining contagion risk, and the coefficients on both are positive. At low densities, higher assortativity and k-nearest neighbour reflect increasing connections within the core. The coefficients of assortativity initially increase in magnitude with connectedness before falling off as risks become less concentrated in the core; and those of k-nearest neighbour are higher for higher densities, reflecting stronger connectedness not just in the core but in the whole network.

For individual bank defaults, aggregate network metrics tell a different story: increasing connectedness on average *reduces* contagion risk, with coefficients on metrics such as the average in-degree of banks and k-nearest neighbour being statistically significant and negative, with coefficients falling in absolute magnitude as density increases. This reflects the fact that higher average connectedness means that individual exposures to the defaulting bank are smaller in magnitude and spread out across more counterparties, all else being equal, so that its default represents a smaller and more diversified shock to the system.

Coefficients on individual bank-level the metrics that measure the *connectedness of the defaulting bank* – specifically, the coefficients on the in-degree (the number of exposures to the defaulting bank) and k-nearest neighbour (number of neighbours of its nearest neighbours) of the defaulting bank – are statistically significant. Both in-degree and k-nearest neighbour exhibit decreasing sensitivity of contagion risk to connectedness as risk diversifies. This captures the fact that the more connected an individual bank, the more risk it poses to the system (hence the positive coefficients); and that this risk declines as the network becomes more connected.

However, consistent with the findings of our contagion analysis, once we introduce real-world heterogeneity among banks, our qualitative results change and the strong link between network structure and solvency contagion risk weakens. For system-wide shocks, simple network metrics no longer offer a robust way to monitor solvency contagion risk. For individual shocks, the connectedness of individual banks remains significant in explaining the risk they pose to the system, consistent with the findings of Boss *et al.* (2004a). So we conclude that more sophisticated metrics and models are indeed necessary to understand and monitor solvency contagion risk.

3 Methodology

To systematically analyse the relationship between network structure, institutional characteristics and solvency contagion risk, we use the solvency contagion model introduced by Bardoscia *et al.* (2019). We apply this model of solvency distress contagion to a wide range of simulated networks for both system-wide and idiosyncratic exogenous shocks, including networks calibrated to historical data for the UK interbank network. We generate these networks using the methodology outlined in Gandy & Veraart (2016), implemented using the software packages published by the authors.[2]

We then analyse how different standard network metrics (such as centrality, clustering, assortativity, etc.) relate to the losses generated by this contagion mechanism, and therefore which measures carry useful information about the potential vulnerability of the interbank network to this channel. The remainder of this section describes: the solvency contagion model we use; the networks we generate; how we define and initialise banks' balance sheets; the shocks we apply to the system; the regression models we use to assess the informativeness of simple network metrics; and the data we use to calibrate our heterogeneous networks.

3.1 A Model of Solvency Contagion Risk

To quantify solvency contagion risk under a set of exogenous shocks we use the model of Bardoscia *et al.* (2019), which is deployed by the Bank of England in regulatory stress testing (see e.g. Bank of England (2016)).

The model features a network of banks, connected by interbank assets and liabilities that revalue as banks' solvency positions change, due to shocks to the value of their other assets. The value of an interbank loan is determined by a modelled probability of default (PD), and an exogenously fixed recovery rate. The PD in turn is a function of bank capital, and is defined using a structural credit model – an adapted version of the Black and Cox model (Black & Cox, 1976). In this model, a bank will default when its capital falls below zero. The model can use two definitions of bank capital: the shareholder equity, or the common equity Tier 1 (CET1) capital buffer above regulatory minimum requirements. The choice of the latter is consistent with the observation that a bank may be put into resolution if their level of CET1 capital falls below the minimum requirements, and is the definition that we use in this paper (see (Bardoscia *et al.*, 2017a)).

Figure 1 illustrates how banks' probability of default increases as their capital buffer is eroded, using data for several major UK banks in 2015. The figure highlights that this relationship is non-linear and monotonically decreasing in the level of capital; and that it is a bank-specific function. As such, the impact of bank heterogeneity on solvency contagion risk is incorporated in the model.

Initial shocks to banks' capital buffers lead to increases in their PDs, and a resulting revaluation of interbank claims - which we term 'direct' losses. This revaluation in turn causes

[2] Available at `https://cran.r-project.org/web/packages/systemicrisk/index.html`.

further deterioration to bank solvency, resulting in further increases in PDs, increasing losses and amplifying the initial shock - which we term 'contagion' losses. We quantify solvency contagion risk in this paper as the contagion losses, i.e. the additional losses in the banking network following the initial shock and the direct losses.

The non-linear relationship between PD and capital illustrated in Figure 1 means that a well-capitalised banking system can absorb quite large initial shocks before there is material contagion, even when the size of interbank exposures is significant relative to the size of external assets. Conversely, a poorly-capitalised banking system will be liable to amplify even small shocks, particularly where interbank exposures are large - as was the case during the financial crisis. So, in addition to the structure of the network, the total capitalisation of individual banks and the size of exposures relative to bank capital are key factors in the degree of solvency contagion risk present in an interbank network.

Figure 1: Probability of Default as a Function of Shocks to the Initial Capital Buffer

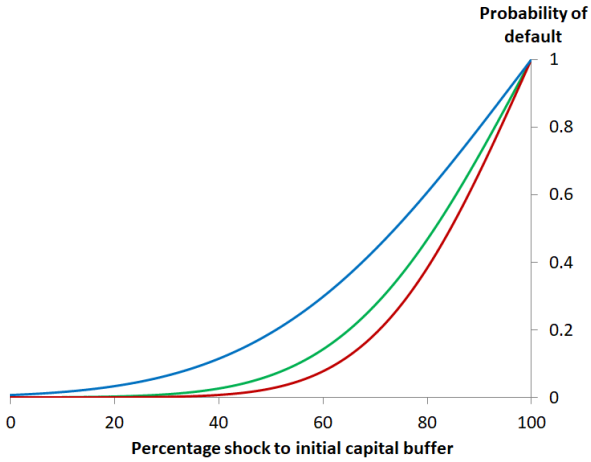

Variation in probability of default with shocks to initial capital buffer for several major UK banks. Three curves refer to three different banks: well-capitalised (red), medium-level-capitalised (green) and poorly-capitalised (blue) banks.

3.2 Network Types

We consider four types of network structure, and the effects of different levels of connectedness within each. For each of the four network topologies described, we consider varying levels of connectedness up to a density of 50%.[3] For connectedness levels higher than 50%, the networks become so dense that all the topologies become very similar. We generate 500 networks for each network topology and density, and present results averaged across these network realisations.

1. Regular networks.
We generate k-regular networks of increasing connectedness, or more specifically increasing node degree (Figure 2). These are homogeneous networks with every node having

[3]The levels are following: 2%, 3%, 4%, 5%, 6%, 7%, 8%, 9%, 10%, 15%, 20%, 25%, 30%, 35%, 40%, 45%, 50%. X% corresponds to link density of the network and means that X% of potential links are actually present

equal degree k, i.e. a flat degree distribution. Each link has the same weight. Ring and complete networks are extreme cases of regular networks.

Figure 2: Connectedness in Regular Networks

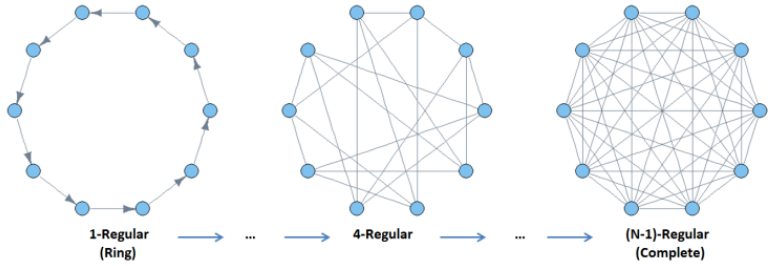

2.Erdos-Renyi networks.
An Erdos-Renyi network is a purely random network and has a Poisson degree distribution (Figure 3). We generate networks of this type (as well as scale-free and core-periphery networks that follow) in two ways: i) with homogeneous link weights; and ii) with heterogeneous link weights, based on the exposures in the actual UK interbank network. To do that, one should first extract column and row sums of the interbank exposures matrix from the data (i.e. one horizontal and one vertical vectors containing total interbank liabilities and total interbank assets of each bank in the network). Then, the matrix is randomly reconstructed by assigning a link between each pair of banks with probability X% for each connectedness level, ensuring that the total vertical and horizontal sums of the link weights for each bank are equal to the bank's total interbank liabilities and assets, respectively. The homogeneous link weights case is then obtained by equalising all the created links of this network simulation. This is repeated $M = 500$ times. We employ the Bayesian reconstruction technique developed by Gandy & Veraart (2016), implemented using the software packages published by the authors.[4]

Figure 3: Erdos-Renyi Network

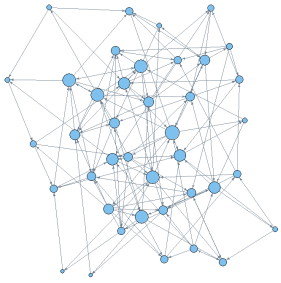

Node area reflects node centrality.

3.Scale-free networks.
Scale-free networks have a heavy-tailed degree distribution with few hub nodes that intermediate between many weakly-connected nodes - with hub nodes being a feature of real

[4]Available at https://cran.r-project.org/web/packages/systemicrisk/index.html.

interbank networks - but those hub nodes do not necessarily connect to each other (Figure 4). Scale-free networks are generated in a similar fashion as the Erdos-Renyi networks using the Bayesian reconstruction technique described above. The difference is that, in addition to the link existence probability X%, another so called intensity parameter is provided that allows for controlling how intensely the links are assigned for the hub nodes and for the spoke nodes.

Figure 4: Scale-Free Network

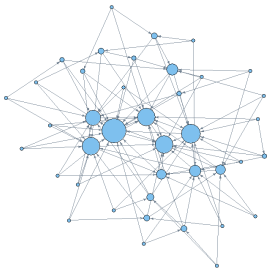

Node area reflects node centrality.

4.Core-periphery networks.
The core-periphery network structure most resembles the real interbank network. It also has a heavy-tailed degree distribution with few hubs and many weakly-connected nodes. But unlike in the scale-free case the hubs are also strongly connected to each other, forming a 'core' of the network (Figure 5). Periphery nodes are primarily connected to core nodes (and so indirectly connected to each other via the core). To achieve this structure when simulating the networks, we follow the procedure used for generation of Erdos-Renyi networks, however, this time different link existence probabilities are assigned for the core-to-core sub-matrix (high probability), core-to-periphery sub-matrix (medium probability) and periphery-to-periphery sub-matrix (low probability).[5]

Figure 5: Core-Periphery Network

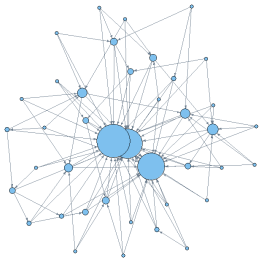

Node area reflects node centrality.

[5]Chart 19 in Appendix 6.2 illustrates the exact densities within each sub-matrix for each level of connectedness.

3.3 Balance sheets

Bardoscia *et al.* (2017a) show that decreasing exposure sizes and increasing capital levels have reduced solvency contagion risk in the UK banking system since the crisis. In this paper, we focus first on the interaction between network topology and contagion for a network of similar banks, before examining in more detail the interaction between heterogeneous capital, exposure levels and network topology. For this reason, we initially generate networks consisting of banks that are homogeneous in balance sheet size and capital ratios:

- The network consists of $N = 100$ banks.
- Banks have 'external assets' (e.g. real economy loans) and interbank assets; and external liabilities (e.g. retail deposits), interbank liabtilities, and equity.
- We set each bank's total risk-weighted assets (RWA) to 200. This determines banks' minimum capital requirements.
- We initialise all banks with CET1 capital ratios of 13.3% of RWA (the average CET1 ratio of UK banks from July 2017 (Bank of England, 2017)).
- The CET1 minimum requirement is 4.5% of RWA.
- We focus on the case in which total exposures in the network are fixed, and changes to network topology affect how such exposures are distributed among banks in the interbank network. In this setup, individual interbank exposures are all the same size - the links in a specific network all have the same weight. As more connections are added, the sizes of individual exposures fall.[6]

For regular networks - in which all banks have the same number of connections - we do not expect diversification to materially affect contagion risk for system-wide shocks. This is because the mechanism we consider is monotonically increasing in shock size; and the distress process iterates to convergence. For a regular network in which all participants are connected, a shock to any single participant will transmit to all participants (and back again). Increasing connectedness, for example, by doubling each bank's number of counterparties, reduces the absolute shock to each exposure by half, but doubles the number of exposures, such that the total amplification should be invariant to connectedness.[7] For individual bank defaults, increasing connectedness should reduce contagion risk, due to the non-linear relationship between capital losses and bank PD. Note that if the relationship between shock size and PD is linear with slope 1, such that a x% fall in capital produces a x% fall in the PD - as in the DebtRank formalism of Bardoscia *et al.* (2015) - risk diversification in a regular network for bank defaults is not possible,

[6] We have also considered cases in which individual interbank exposure sizes are fixed, such that more connections linearly increase total exposures. These represent a situation in which an increasingly large proportion of each bank's balance sheet consists of interbank exposures. We do not focus on this case for two reasons. First, the results are fairly trivial - contagion risk monotonically increases with increasing connectedness, and network structure has a limited impact on the magnitude of the risk as the total size of exposures relative to capital becomes the dominant feature. Second, to the extent that interbank lending is a market mechanism for risk-sharing, an exploration of how contagion risk evolves as interbank assets and liabilities become the dominant constituents on banks' balance sheets has limited practical relevance.

[7] This holds where all banks have at least one interbank asset and at least one interbank liability, i.e. an in-degree of one or more and an out-degree of one or more.

as shocks transmit undampened to all connected participants.

For core-periphery, scale-free and Erdos-Renyi networks, for individual bank defaults we expect to observe that contagion risk initially increases, and then decreases with higher network connectedness. This reflects an initial increase in the probability that a bank will be exposed to the defaulting bank, which is then offset by diversification benefits as the impacts of the default is dissipated among a larger number of banks. For system-wide shocks, the higher concentration of exposures relative to regular networks means that contagion risk should also be correspondingly higher. As connectedness increases, core-periphery, scale-free and Erdos-Renyi networks will start to resemble regular networks, and increasing diversification should result in a reduction in contagion risk.

The 'homogeneous balance sheet' case is useful for isolating the impact of network structure on contagion risk as much as possible. But it is not representative of reality. Real-world interbank networks do not consist of homogeneous banks, and contagion risk is sensitive to heterogeneity in the banking system. So we then consider the case of heterogeneous balance sheets. We calibrate banks' capital and interbank exposure levels to historical data on the UK banking system, to assess how these might modify our conclusions on the importance of network topology and connectedness.

Real-world banking networks tend to exhibit a core-periphery structure; and core banks tend to have larger balance sheets than periphery banks. In our homogeneous networks, core banks have larger interbank exposures, but the same level of capital as periphery banks. UK data confirms that core banks with larger interbank exposures have larger balance sheets and capital buffers (in absolute terms) than periphery banks. So, when we calibrate a network to reflect more realistic (heterogeneous) bank balance sheets, we would expect to see lower contagion risk relative to the case in which banks have homogeneous balance sheets. This again is a function of the non-linear relationship between shocks to capital and bank PD: increasing the ratio of capital to interbank exposures at a bank means that losses on those exposures have a smaller impact on the banks' PD *ceteris paribus*, and thus reduce contagion losses.

3.4 Shock scenarios

We consider two types of shock:

- System-wide shocks: these entail an equal exogenous shock to all banks, of 50% of their initial capital buffers;[8]
- Individual defaults: random defaults of 3 banks in the network.[9] For relevant heterogeneous networks, we separately consider defaults of core and periphery banks.

[8] We also considered shocks of 1%, 5%, 10%, and 20% of capital buffers, but found negligible contagion losses for these shocks, due to the non-linear relationship between capital erosion and PD increases and the relatively high starting level of bank capital.

[9] We also consider the default of 1 and 2 banks, and obtain qualitatively similar findings.

3.5 Regression Model

The impact of network structure on contagion losses in various shock scenarios is studied using the following regression model:

$$log(Losses_i) = \alpha + \beta NetworkCharacteristics_i + \gamma NodeCharacteristics_i + \\ + TopologyFE + \epsilon_i \quad (1)$$

This is a cross-sectional regression with subscript $i = 1 : NDM$, where $N = 4$ is the number of network topologies, $D = 17$ is the number of densities for each topology, and $M = 500$ is the number of simulations for each topology-density pair. The regression is run for each shock scenario and for each balance sheet case. The dependent variable $log(Losses_i)$ is natural logarithm of the contagion losses to the network for the corresponding shock scenario and balance sheet case (contagion round losses excluding the initial shock and the direct losses). The regressors include various standard network statistics.[10] These are categorised into two groups:

- Network characteristics - network-level statistics, such as average network degree centrality, density, average shortest path, etc.
- Node characteristics - node-level statistics of a defaulted node (bank), such as degree centrality, eigenvector centrality, nearest neighbour degree, etc.[11]

We also include fixed effects for the $N = 4$ network topologies.

3.6 Data

We use historical data on the UK banking system and interbank network to calibrate our heterogeneous networks. The data are from the FSA's Recovery and Resolution Plans (RRP) returns on interbank exposures, as of the second half of 2013 (the last point at which they were collected). This is the most comprehensive and granular data we have available of interbank exposures across the entire UK network: subsequent and current data collections either have a significantly lower coverage of banks, or a less complete and granular split of exposures by asset class and product type (or both). Since we are interested in calibrating our generated networks to capture as comprehensively as possible the whole (UK) network structure, while these data are a little out of date, we consider them to be the most relevant data for our purposes. These data have been comprehensively mapped in Langfield *et al.* (2014).

As part of this data collection, approximately 180 firms reported their 20 largest exposures (plus the six largest UK lenders if they are not included in those 20) for each product type, ensuring a fairly comprehensive view of the UK banking system. However, some parts of the network were still not observed due to the cut-off at 20(+6) counterparties.

Analysis of the dataset reveals two relevant features that are important for setting our results into context. First, as explored in detail by Langfield *et al.* (2014), the interbank

[10] The full list and description of all considered statistics is included in the Appendix.

[11] Note: node-level characteristics are only relevant for individual bank defaults scenario, and not for system-wide shock scenarios.

exposures network has the expected core-periphery structure.[12] Second, we find that the density of interbank exposures networks, including only those banks with at least one connection, are in the region of 2%-6%.[13]

4 Results

4.1 Contagion and Network Topology

Homogeneous bank balance sheet size and capital ratios

We first present results for networks of banks with homogeneous balance sheet sizes and capital ratios. Figure 6 shows contagion losses as a function of network connectedness - average density - for the **regular network**, for system-wide shocks and for the defaults of three banks.[14] For small system-wide shocks (not shown in the chart), there is little amplification of the shock beyond the first-round (direct) losses, as direct losses on interbank exposures are insufficient to materially change banks' PDs. For the larger shocks shown on the chart (of 50% of banks' capital buffers), there is modest additional contagion. As expected, network connectedness has no impact on contagion risk.

For individual bank defaults, we observe that increasing connectedness quickly diversifies the risk to the point where contagion losses are minimal (results for the defaults of one and two banks are qualitatively similar to those for three bank defaults). For regular networks then, network connectedness is not a key driver of contagion risk.

Figure 7 shows contagion losses for the core-periphery networks under the same shocks.

Figure 6: Regular Networks with Homogeneous Bank Balance Sheets

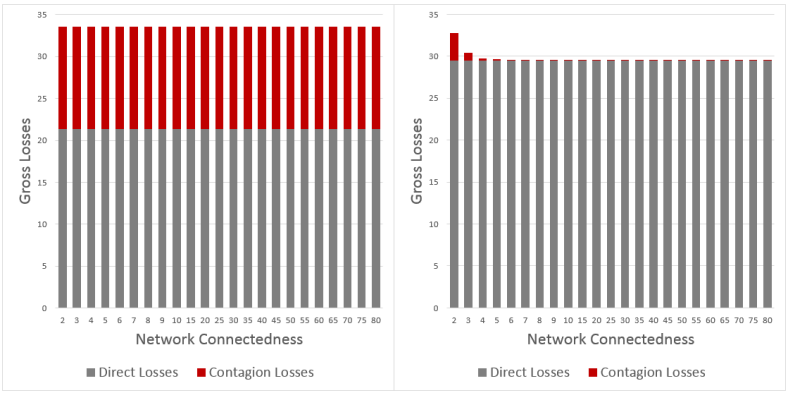

Left: system-wide shock (50% shock to capital buffer). Right: individual bank defaults (3 banks' defaults).

[12] This holds both for total exposures, and for the network of fixed income holdings which is particularly relevant for the solvency contagion risk considered in this paper. The unsecured interbank lending network is lower in closeness to a core-periphery structure.

[13] The densities when we also include banks with no connections in the statistics are in the region of 1%-4%. The actual density for networks of connected banks is likely to be a little higher, due to the data cut-off which excludes reporting of some connections for the biggest banks.

[14] In this figure, we show only the direct (first-round) losses and contagion losses; losses due to the initial shock to equity are excluded for clarity, as these are significantly larger than those due to network effects.

Figure 7: Core-Periphery Networks with Homogeneous Bank Balance Sheets

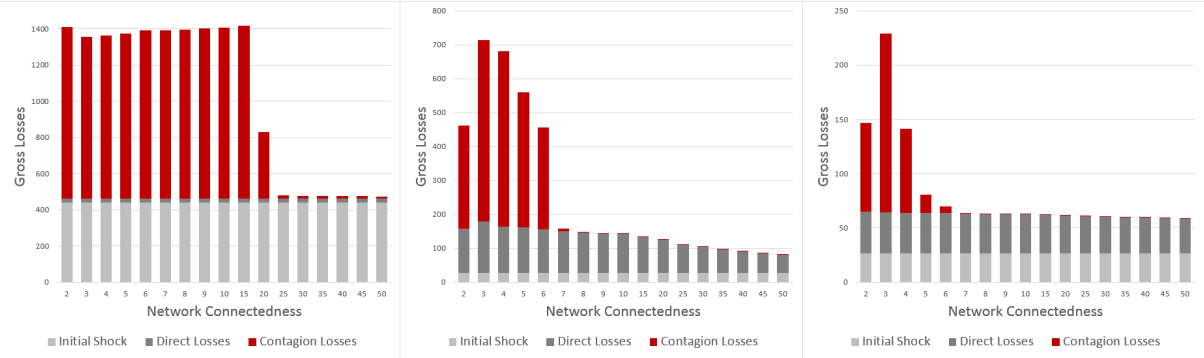

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

Contagion losses for scale-free and Erdos-Renyi networks have similar shapes, but different magnitudes (Figure 21 and Figure 22 in the Appendix 6.2). There are several features to note:

- For system-wide shocks, we observe significant contagion losses for lower levels of network connectedness. These appear to depend weakly on the level of connectedness, until it becomes high enough for risk-sharing benefits to manifest, at which point behaviour is similar to that of regular networks.
- For individual bank defaults, we observe the expected behaviour that contagion risk initially increases with increasing connectedness, before diversification benefits lead to a drop in contagion risk.
- The magnitude of contagion losses for these network topologies is much higher than for regular networks.
- We note the (obvious) result that contagion losses are much larger for individual bank defaults than for system-wide shocks of a comparable initial fraction of total capital. This is trivially due to the non-linear loss function. A 1% fall in system-wide capital due to a system-wide shock has essentially no impact on each bank's PD and so interbank exposure values are essentially unchanged; whereas a 1% fall in system-wide capital due to the default of a core bank wipes out the value of all interbank exposures to that bank, with material knock-on effects on counterparty banks' PDs and contagion risk.

The second point - that system-wide losses collapse beyond a certain average density - merits further attention. The intuition behind the result is that as connectedness increases, the networks increasingly resemble regular networks - with, for example, core banks becoming less 'core-like' and periphery banks becoming less peripheral in the core-periphery network. Figure 8 illustrates how the degree distribution for banks in the core-periphery network flattens out as average density increases.

This in turn means that the distribution of the ratio of capital to exposures is changing - as the network becomes more regular and exposures are more evenly distributed, core banks have relatively more capital against their interbank exposures (and periphery banks have

relatively less). This increase in risk-sharing means that following the exogenous shock, the losses that core banks make on their (reduced) interbank exposures have a smaller overall initial impact on PDs; and thus contagion risk falls, due to the non-linearity of the loss function.[15]

Figure 8: Degree Distributions as a Function of Network Density

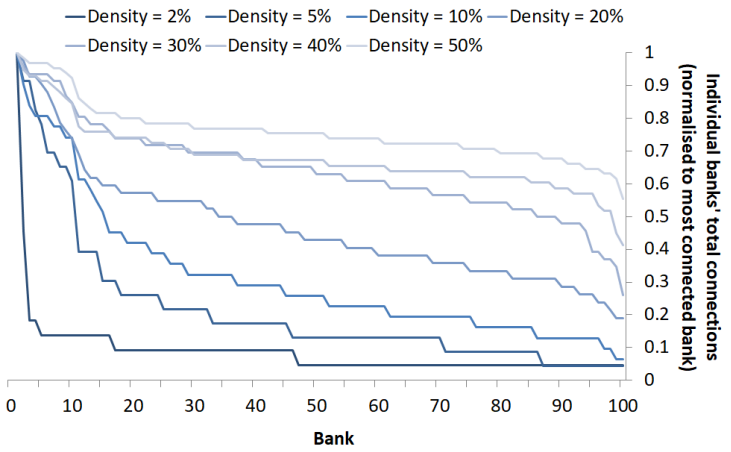

Degree distributions for banks in core-periphery network at different network densities. Total degree is normalised to the number of connections of the largest bank. As density increases, the degree distribution flattens out and networks become less core-periphery in nature.

Figure 9 plots the average k-nearest neighbour and assortativity coefficients for each level of considered network densities. While k-nearest neighbour increases monotonically with higher density, the assortativity coefficient illustrates a sharp increase around the smaller densities and then fluctuates at around -0.15 for the higher ones.[16] At the lowest levels of connectedness the network is strongly disassortative with weakly connected core and almost no connections among the peripheral banks. As connectedness increases, the net-

[15]We note that in our setup, banks are not homogeneous with respect to their *total interbank exposures* - changes in network structure also change the size and distribution of individual banks' interbank exposures. An alternative setup for focusing on the impact of network structure would be to impose homogeneity in banks' *capital/interbank exposure ratios* - so that banks with higher interbank exposures have proportionally higher capital ratios than banks with smaller interbank exposures. In this setup, for a fully connected network we find that for system-wide shocks, contagion risk becomes invariant to average density, due to the nature of our contagion mechanism. This is because the same relative initial shock to capital has the same impact on all banks' PDs; all interbank exposures are equally written-down; and since the capital/interbank exposures ratio is equal for all banks, subsequent rounds of losses are also invariant to the number and size of exposures on any node. This illustrates the importance of disentangling the roles of network structure and node properties when considering contagion mechanisms that interact with both. One could in principle attempt to generate a network that allowed variable individual exposure weights, in order to have bank balance sheets homogeneous in all three of total interbank assets, capital and total RWAs: this would require for example that core banks have many small exposures, and periphery banks have one or two large ones. But it is hard to envisage how we could define a reasonable and robust approach in practice to generating such a network across all of the dimensions we've considered - core banks after all connect to periphery banks, so if the latter have a single large exposure, there is no balance sheet capacity left for the connected core bank to connect to other banks, unless we start adding in periphery-periphery exposures (which would change the network topology).

[16]There is material variation in k-nearest neighbour and assortativity metrics across the 500 simulated networks at each density, supporting our econometric analysis. At a density of 5% for example, k-nearest neighbour varies between around 28 to 36 while assortativity varies in the range -0.18 to -0.01.

work becomes less disassortative further confirming the above-mentioned observation on core (peripheral) banks becoming less 'core-like' ('periphery-like').

Figure 9: k-nearest Neighbour and Assortativity Coefficient

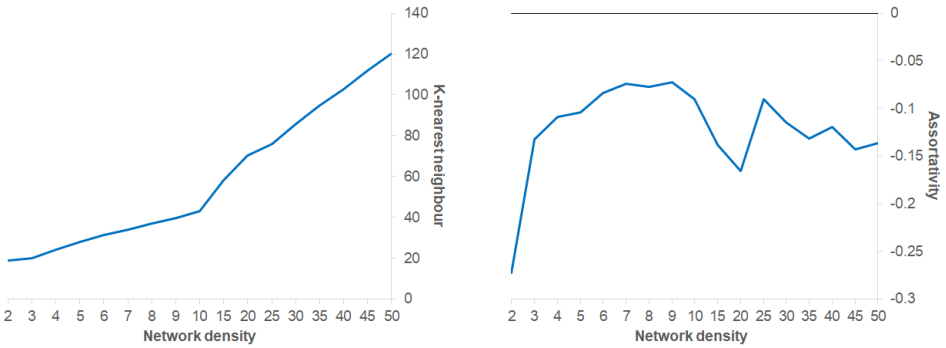

Network metrics of core-periphery networks: a) k-nearest-neighbour and b) assortativity coefficient. Averages across 500 simulations for each density level are presented.

Figure 10: Density of Connections Between the Largest Banks in the Network

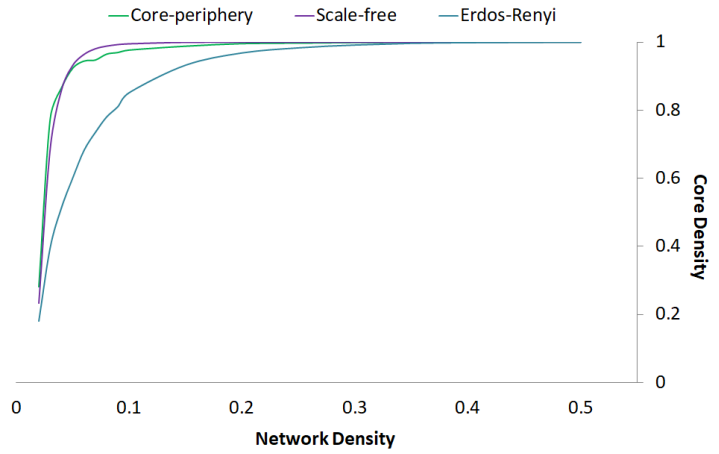

Core density is highest in the scale-free network for average network densities above 4; and much higher in scale-free and core-periphery networks than in the Erdos-Renyi network. Combined with the results in Figures 7, 21 and 22, this shows that higher density of interconnectedness among the largest banks in the network translates to more persistent contagion risk.

We also note that the average density at which contagion risk collapses varies among the networks. Diversification benefits manifest at lower levels of connectedness for the Erdos-Renyi network than the core-periphery and scale-free networks; and at lower levels for core-periphery than for scale-free. This finding is consistent with the fact that the Erdos-Renyi network is the least heterogeneous. Analysing the density of connections among the largest or 'core' banks similarly reveals that this density is higher in the scale-free case than in the core-periphery case for all but the lowest average network densities. So we find that the point at which diversification benefits dominate occurs at higher average densities when the largest, most connected banks are more interconnected. This is illustrated by Figure 10.

Figure 11: Contagion Losses as a Function of Interbank Exposure Size and Network Density

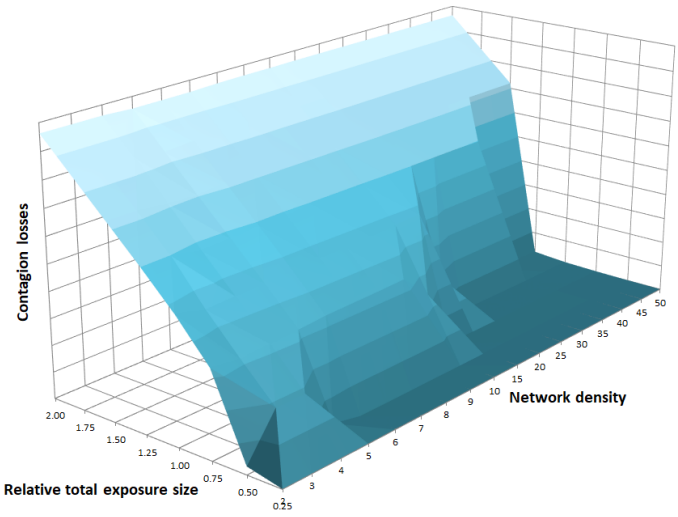

Contagion losses as a function of interbank exposures and network density, where interbank exposures are expressed as a fraction/multiple of the baseline level of exposures used for our experiments. As exposures increase in size, contagion risk persists to higher network densities and becomes less sensitive to network connectedness.

Given the evident sensitivity of contagion risk to capital/exposure ratios, we next explore the impact of varying the size of interbank exposures, focusing on the core-periphery network.[17] Figure 11 illustrates, for the core-periphery network, how contagion losses for a system-wide shock of 50% of banks' capital buffers vary as a function of total exposures.

When interbank exposures are reduced to half the level in our base case, the magnitude of solvency contagion risk falls significantly; and we recover a hump-like profile, in which increasing connectedness initially increases contagion risk, before diversification benefits dominate. As the size of interbank exposures *increases*, the point at which contagion risk collapses is pushed to higher levels of connectedness. For increasing exposure sizes then, contagion risk in heterogeneous networks becomes less sensitive to the level of connectedness of the network.

These findings illustrate that both the type of network structure, and the level of connectedness within it, can have a significant impact on solvency contagion risk. The most risky networks are those that exhibit heterogeneity in network structure, and in particular those with relatively low average connectedness. These characteristics are typical of interbank networks, which tend to have a core-periphery structure and low density, as discussed above. At first glance, this suggests that at least from the perspective of solvency contagion risk, interbank exposure networks have evolved in a manner that makes them particularly risky. However, as discussed above, in real banking networks, heterogeneity is a feature not just of the network, but of the banks within it; and the size of interbank exposures as a fraction of bank capital is an important driver of contagion risk, and its sensitivity to network structure. To assess whether our pessimistic conclusion holds for

[17]The effect on core-periphery networks is representative of the effects on scale-free and Erdos-Renyi networks; and the impact on regular networks is simply that contagion losses for all levels of connectedness increase (decrease) in magnitude as exposures increase (decrease) in size.

networks with more realistic banks, we turn our consideration to networks of banks with varying exposure and capital levels calibrated using historical data for the UK banking system.

Heterogeneous bank size and capital ratios

Figure 12 illustrates that once we allow for observed heterogeneity in banks, the dependence of solvency contagion risk on network structure changes. For individual bank defaults, we lose the characteristic hump in contagion risk as connectedness increases. Strikingly, the magnitude of potential losses also falls significantly.[18]

These results are explained by the fact that the most connected banks - for example, the

Figure 12: Core-Periphery Networks with Heterogeneous Bank Balance Sheets

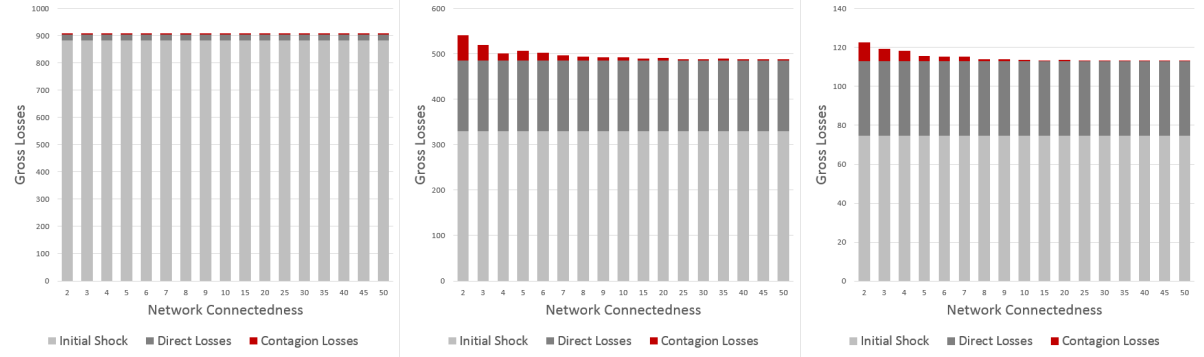

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

core banks in the core-periphery network - are larger and hold more capital relative to interbank exposures than their less connected siblings. Figure 13 compares the ratio of capital to total interbank exposures for groups of banks in our generated networks. It shows that where we impose homogeneous balance sheets on our banks, the core banks have less capital against their interbank exposures than when we calibrate bank balance sheets to real data. The periphery banks meanwhile have smaller capital against their interbank exposures on average when calibrated to real data.

So for core banks, a given shock to interbank exposures represents a smaller shock relative to assets and capital than in the case where we had banks with homogeneous balance sheets. This in turn explains why the increasing connectedness no longer initially increases risk; with smaller exposures relative to capital and a non-linear loss function for core banks, any diversification of exposures is sufficient to dissipate the impact of the initial shock. The increase in the exposure/capital ratio for periphery banks is insufficient to counterbalance this effect, as the core is better able to absorb losses on those exposures.

These findings emphasise the importance both of policies that limit the size of interbank exposures; and that reinforce the resilience of large, highly interconnected banks - the core of the banking system. We do not attempt however to identify a 'critical condition' for when the capital/exposure ratio in the core becomes critical for financial stability, as this

[18]Contagion losses for scale-free and Erdos-Renyi networks are presented in Figure 23 and Figure 24 in the Appendix 6.2.

Figure 13: Capital/ Interbank Exposure Ratios for Homogeneous and Heterogeneous Bank Balance Sheets

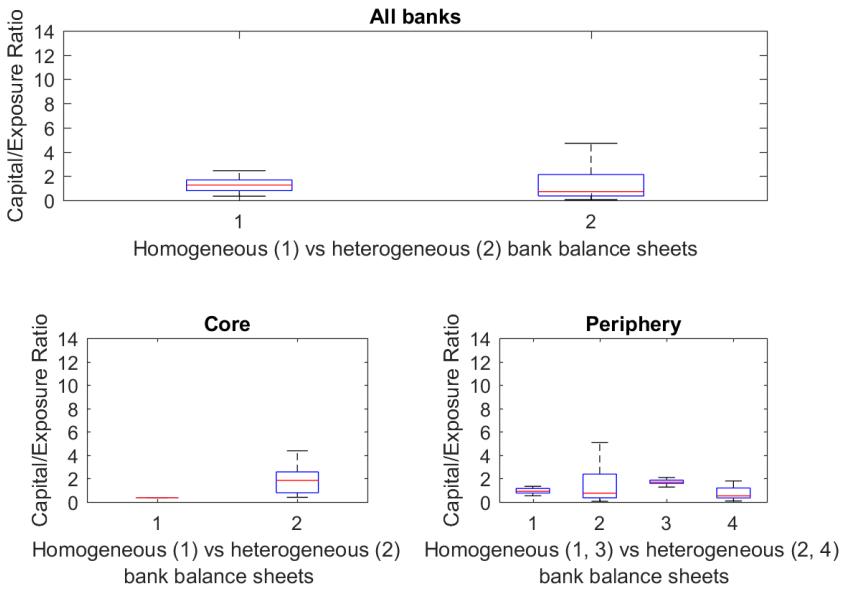

Distribution of capital to interbank exposure ratios for core and periphery banks for network density of 10%. Periphery banks are split further into two groups of equal number, containing the larger and smaller 'periphery' banks respectively. Outliers have been removed.

depends both on the size and distribution of the initial shock we are concerned with, and on the shapes of the probability of default functions of all of the banks in the network. For small initial shocks to equity, interbank exposures would have to be many times larger than capital in order to produce meaningful amplification in a well-capitalised system (with a probability of default function illustrated by the red line in Figure 1); whereas for large enough initial shocks, only minimal further losses via solvency contagion would be required to tip the system into widespread insolvency. Rather, we emphasise the value of simulating solvency distress contagion in real interbank networks, and considering how the size of exogenous shock the system can absorb without significant amplification varies with exposure size.

4.2 Network Metrics and Their Informativeness

We have demonstrated that network structure and the heterogeneity of the banks within it are important factors in explaining the risk of solvency contagion. We now ask whether there are network metrics we can use that can efficiently capture these dependencies, and so provide a simple way to monitor solvency contagion risk across different banking networks and through time - or whether model-based approaches and more sophisticated network metrics for measuring and tracking solvency contagion risks are required.

We focus on core-periphery networks, as these are the most relevant for the banking system. First, we run a set of global regressions for the case of homogeneous banks and exposures - examining whether any network metrics can successfully explain contagion

risk across different connectedness levels and shock types. The short answer is no. While some metrics can explain contagion with statistical significance (see Table 2 in the Appendix 6.3), the signs attached to these metrics are often unintuitive and inconsistent between similar metrics. This is unsurprising - we have seen for example that contagion risk first increases then decreases with increasing density in the case of idiosyncratic shocks in a core-periphery network of homogeneous banks; and that with system-wide shocks, the system transitions from a high-contagion risk to a low-contagion risk regime, with little variation in risk within each regime. So the dependence of contagion risk on structure is qualitatively different in different regimes and under different shock types, and simple metrics cannot capture the changes in solvency contagion risk across them all.

So we run regressions for each density, and examine how the coefficients on each metric vary with density, distinguishing between system-wide shocks and individual bank defaults. Table 1 presents a closer look at the regression results for the core-periphery network with 5% density, which resembles the actual UK interbank network both in terms of topology and level of connectedness.[19] The top panel contains estimation results for the network-level features that characterise the overall network, and the bottom panel contains results for the node-level features that characterise the node that has defaulted as an initial shock to the system. Note that node-level features are only relevant for the *defaults in the core* shock scenario and not for the *system-wide shock* scenario, where all nodes get a common percentage initial shock to their capital buffer but no node defaults initially. We can see from this table that, for the homogeneous banks case, higher average in-degree leads to lower contagion losses at the network level. Thus, the diversification benefits already manifest at the density of 5%. However, the same network statistic (in-degree) but at a node-level has positive sign. This implies that even though the contagion losses decrease in overall network connectedness, they still increase in the connectedness of the node that initially defaults.

Five figures below (14, 15, 16, 17, and 18) visualise the regression results for some of the regressor coefficients across all 17 densities. We find that several of the network metrics are informative under this set-up (for example, see Table 1 for density 5%). For system-wide shocks, the k-nearest neighbour and assortativity measures are statistically significant in explaining contagion risk and the coefficients on both are positive. At low densities, higher assortativity and k-nearest neighbour reflect increasing connections within the core. The coefficients of assortativity initially increase in magnitude with connectedness before falling off as risks become less concentrated in the core; and those of k-nearest neighbour are higher for higher densities, reflecting stronger connectedness not just in the core but in the whole network (Figure 14).

For individual bank defaults, aggregate network metrics tell a different story: increasing connectedness on average *reduces* contagion risk, with coefficients on metrics such as the average in-degree of banks and k-nearest neighbour being statistically significant and negative, and falling in absolute magnitude as density increases (Figure 15). This reflects the fact that higher average connectedness means that individual exposures to a defaulting bank are smaller in magnitude and spread across more counterparties all else being equal, so its default represents a smaller and more diversified shock to the system.

[19] The detailed regression results for other 16 densities are available on request.

Figure 14: Homogeneous Bank Balance Sheets: Coefficients on k-nearest-Neighbour and Assortativity by Density for System-Wide Shocks

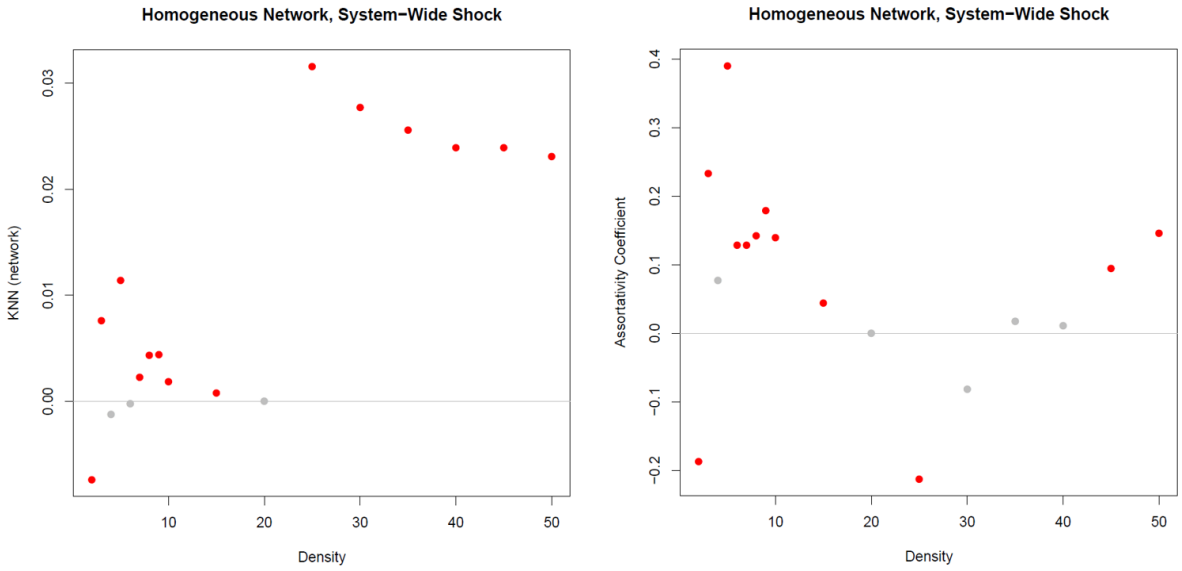

Regression coefficients of contagion losses on a) k-nearest-neighbour and b) assortativity metrics at a network level by density for core-periphery networks of homogeneous banks under system-wide shocks. The coefficients statistically significant at 10% level are coloured in red.

Turning to bank-level measures, we find that coefficients on the metrics that measure the *connectedness of the defaulting bank*, specifically the in-degree (the number of exposures to the defaulting bank) and k-nearest neighbour (number of neighbours of its nearest neighbours) of the defaulting bank, are statistically significant for all density levels considered (Figure 16). Both in-degree and k-nearest neighbour exhibit decreasing sensitivity of contagion risk to connectedness as risk diversifies. This captures the fact that the more connected an individual bank, the more risk it poses to the system (hence the positive coefficients); and that this risk declines as the network becomes more connected.

For many of the other standard network metrics, coefficients are not consistently statistically significant, and so they are not suitable for use in monitoring contagion risk. But for a core-periphery network of banks with homogeneous capital levels, we find that the aggregate network metrics identified above *can* provide a simple way to monitor the relative solvency contagion risk in the banking system; and that a few bank-level network metrics are useful indicators of the solvency contagion risk posed by individual banks (relative to each other and over time) to the system.

However, as we saw above, once we introduce real-world heterogeneity among banks, our qualitative results change and the strong link between network structure and solvency contagion risk weakens. We find that this result also holds for the power of network metrics in telling us something about contagion risk. For system-wide shocks, those network metrics that did offer a robust way to monitor solvency contagion risk in the homogeneous case no longer do so (Figure 17). For individual shocks, the k-nearest neighbour measure is no longer informative; while the in-degree measure does still contain some information but is less robust (Figure 18).

Table 1: Regression Results for Core-Periphery Networks with 5% Density

Regressor	System-wide shock		Defaults in Core	
	Homogeneous banks	Heterogeneous banks	Homogeneous banks	Heterogeneous banks
Intercept	**7.023*** **	**1.058*** **	**2.138*** **	-14.950
	(0.000)	(0.000)	(0.000)	(0.128)
In-degree	**-0.099*** **	**0.062*** **	**-1.206*** **	1.543
	(0.000)	(0.002)	(0.000)	(0.304)
Betweenness	0.019	**-0.066***	0.153	-1.303
	(0.364)	(0.054)	(0.101)	(0.613)
KNN	**0.011**	**0.006**	-0.007	-0.038
	(0.000)	(0.003)	(0.193)	(0.814)
Eigenvector Centrality	0.055	-0.233	0.614	-5.491
	(0.695)	(0.310)	(0.309)	(0.742)
Clustering	0.099	**-0.368*** **	**-0.946*** **	-14.967
	(0.222)	(0.006)	(0.007)	(0.120)
Assortativity	**0.390*** **	0.103	0.073	-4.965
	(0.000)	(0.203)	(0.727)	(0.394)
Shortest Path	-0.026	**0.147*** **	**-0.150***	1.594
	(0.182)	(0.000)	(0.080)	(0.503)
	Node-Level Metrics			
In-degree			**0.048*** **	**0.240*** **
			(0.000)	(0.000)
Out-degree			**0.016*** **	-0.015
			(0.000)	(0.823)
Eigenvector Centrality			**-0.006** **	0.033
			(0.016)	(0.605)
KNN			**0.082*** **	0.074
			(0.000)	(0.474)
Betweenness			**0.011*** **	-0.014
			(0.000)	(0.832)
R^2	24.26%	14.63%	79.28%	24.85%
N	500			

Parameter estimates' p-values are given in parentheses. $^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$.

Figure 15: Homogeneous Bank Balance Sheets: Coefficients on k-nearest-Neighbour and In-Degree by Density for Core Bank Defaults

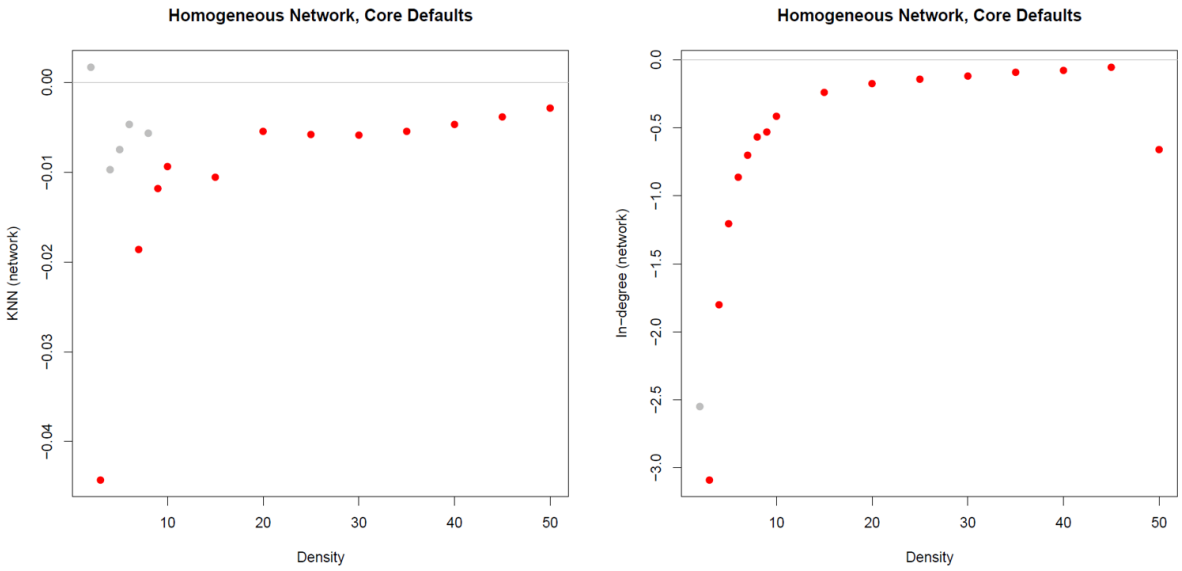

Regression coefficients of contagion losses on a) k-nearest-neighbour and b) in-degree metrics at a network level by density for core-periphery networks of homogeneous banks under individual core bank defaults. The coefficients statistically significant at 10% level are coloured in red.

Figure 16: Homogeneous Bank Balance Sheets: Coefficients on node-level k-nearest-Neighbour and In-Degree by Density for Core Bank Defaults

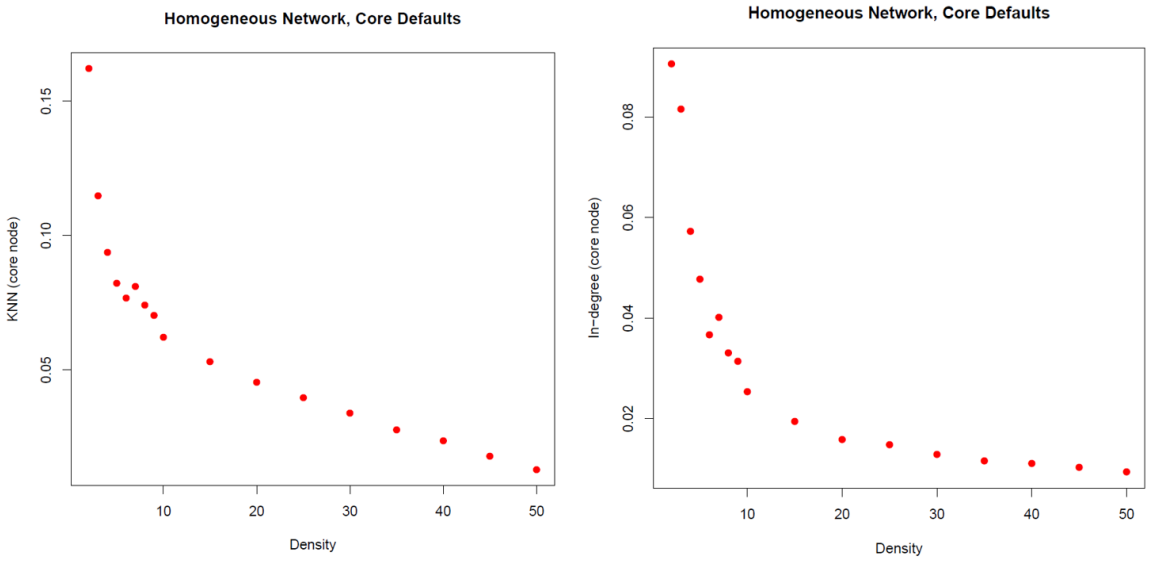

Regression coefficients of contagion losses on a) k-nearest-neighbour and b) in-degree node-level metrics by density for core-periphery networks of homogeneous banks under individual bank defaults. The coefficients statistically significant at 10% level are coloured in red.

Figure 17: Heterogeneous Bank Balance Sheets: Coefficients on k-nearest-Neighbour and Assortativity for System-Wide Shocks

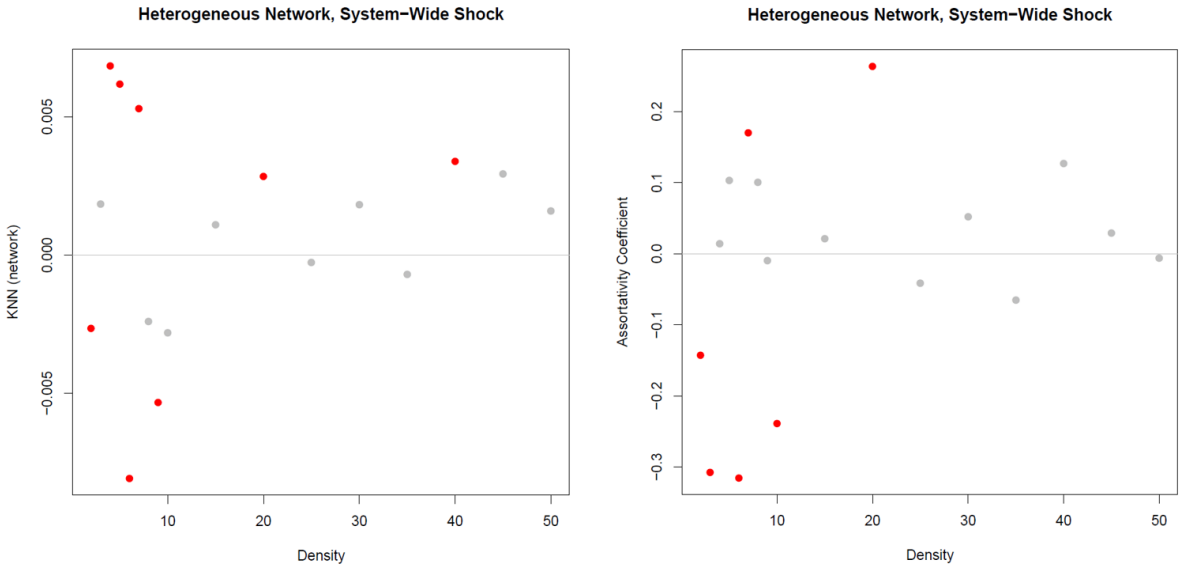

Regression coefficients of contagion losses on a) k-nearest-neighbour and b) assortativity metrics at a network level by density for core-periphery networks of heterogeneous banks under system-wide shocks. The coefficients statistically significant at 10% level are coloured in red.

Figure 18: Heterogeneous Bank Balance Sheets: Coefficients on Node-Level k-nearest-Neighbour and In-Degree for Core Bank Defaults

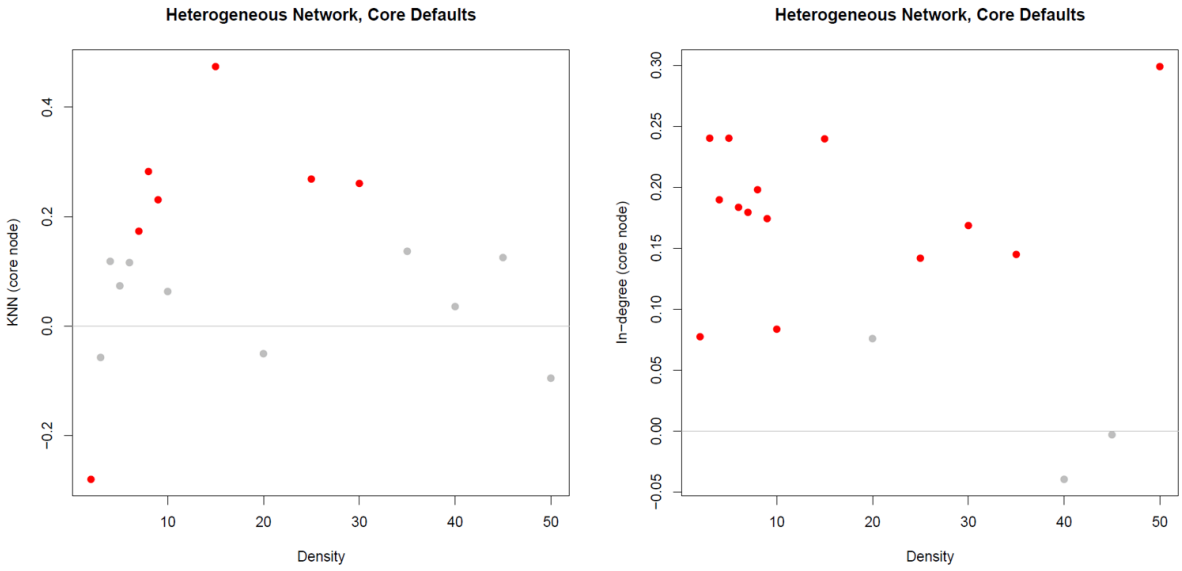

Regression coefficients of contagion losses on a) k-nearest-neighbour and b) in-degree node-level metrics by density for core-periphery networks of heterogeneous banks under individual bank defaults. The coefficients statistically significant at 10% level are coloured in red.

So, once we enter the realm of real banking systems, most simple network metrics are unreliable both for monitoring contagion risks over time, and for differentiating between banks, even when we know the structure and density of the network.

5 Conclusion

In this paper, we demonstrate that solvency distress contagion risk in interbank networks – quantified using the existing model of Bardoscia *et al.* (2019) – is sensitive both to network structure, and to the degree of connectedness within that structure. Heterogeneous network structures, such as the core-periphery topology that typifies real interbank networks, are significantly more prone to solvency contagion risk than random and regular networks.

The sensitivity of contagion risk to network connectedness depends on the size of interbank exposures relative to bank capital. Larger interbank exposures result in larger contagion risk and less sensitivity to the precise degree of connectedness in the network; while for smaller exposures, contagion risk is lower, and material only in a narrow window of low average connectedness, typical of a highly concentrated interbank network.

These findings are consistent with the literature on contagion risks associated with default cascades. For intermediate regimes in which shocks result in some contagion, contagion risk initially increases then decreases with increasing connectedness for interbank networks (as found by e.g. Nier *et al.* (2008); Elliott *et al.* (2014)). For larger shocks, or where exposures are large relative to capital, the system is more vulnerable to widespread and damaging contagion, a property identified by authors including Acemoglu *et al.* (2015); Gai & Kapadia (2010); Cont *et al.* (2013); Allen & Gale (2000); and consistent with the findings of Nier *et al.* (2008); Georg (2013); Halaj & Kok (2013).

We further find that when bank balance sheets are calibrated to historical data on the UK interbank network - introducing heterogeneity of banks as well as network structure - contagion risk falls dramatically and network structure becomes a less important factor. This is due to the fact that the highly connected core banks have higher levels of capital relative to their interbank exposures, and so are better able to absorb contagion risk, enhancing the stability of the network.

These findings lend support to the policy of imposing higher capital buffers on systemically important banks - whose systemic importance is in part determined by their degree of interconnectedness (BCBS, 2013). They are also supportive of the limiting and monitoring of interbank exposures, both to individual counterparties and in totality.

Consistent with the results of our contagion analysis, the results of our econometric analysis show that for systems of banks with similar capital levels and balance sheets, simple network metrics such as connectedness, assortativity and k-nearest neighbour are statistically significant in explaining contagion risks in core-periphery networks. For more realistic networks however – featuring real-world heterogeneity in bank balance sheets – these simple network metrics become unreliable indicators of solvency contagion risk. At an individual bank level, a simple measure of a bank's connectedness to the network (its

degree) does however correlate to the contagion risk it poses to the network in this more realistic case – consistent with the use of a bank's financial interconnectedness as one metric for identifying global systemic importance. These results suggest that policymakers should continue efforts to model contagion risks explicitly, rather than relying on simple aggregate indicators.

One aspect of heterogeneity that we have not considered in this paper is heterogeneity in the risks facing the banks - heterogeneity in balance sheets that would motivate applying different shocks to different banks (as results for example from regulatory stress tests). Analysis of the interaction between heterogeneous balance sheet risks and network structure and their impact on solvency contagion risk - including for example an exploration of how similarities and/or differences in balance sheet risks between core banks impact the contagion risk - represent potential avenues for further research.

References

Acemoglu, Daron, Ozdaglar, Asuman, & Tahbaz-Salehi, Alireza. 2015. Systemic risk and stability in financial networks. *American Economic Review*, **105**(2), 564–608.

Allen, Franklin, & Gale, Douglas. 2000. Financial contagion. *Journal of Political Economy*, **108**(1), 1–33.

Bank of England. 2016. *Stress testing the UK banking system: 2016 results*. Tech. rept. Bank of England.

Bank of England. 2017. *Bank of England Financial Stability Report, July 2017*. Tech. rept. Bank of England.

Bardoscia, Marco, Battiston, Stefano, Caccioli, Fabio, & Caldarelli, Guido. 2015. DebtRank: A microscopic foundation for shock propagation. *PloS ONE*, **10**(6), e0130406.

Bardoscia, Marco, Barucca, Paolo, Brinley Codd, Adam, & Hill, John. 2017a. The decline of solvency contagion risk. *Bank of England Staff Working Paper*.

Bardoscia, Marco, Battiston, Stefano, Caccioli, Fabio, & Caldarelli, Guido. 2017b. Pathways towards instability in financial networks. *Nature Communications*, **8**(14416).

Bardoscia, Marco, Barucca, Paolo, Codd, Adam Brinley, & Hill, John. 2019. Forward-looking solvency contagion. *Journal of Economic Dynamics and Control*, **108**(C).

Barucca, Paolo, Bardoscia, Marco, Caccioli, Fabio, D'Errico, Marco, Visentin, Gabriele, Battiston, Stefano, & Caldarelli, Guido. 2016. Network valuation in financial systems. *Available at SSRN 2795583*.

Battiston, Stefano, Puliga, Michelangelo, Kaushik, Rahul, Tasca, Paolo, & Caldarelli, Guido. 2012. Debtrank: Too central to fail? Financial networks, the Fed and systemic risk. *Scientific Reports*, **2**, 541.

BCBS. 2013. *Global systemically important banks: Updated assessment methodology and the higher loss absorbency requirement*. Tech. rept.

Black, Fischer, & Cox, John C. 1976. Valuing corporate securities: Some effects of bond indenture provisions. *Journal of Finance*, **31**(2), 351–367.

Boss, Michael, Summer, Martin, & Thurner, Stefan. 2004a. Contagion Flow through Banking Networks. *Pages 1070–1077 of:* Bubak, Marian, van Albada, Geert Dick, Sloot, Peter M. A., & Dongarra, Jack (eds), *Computational Science - ICCS 2004*. Berlin, Heidelberg: Springer Berlin Heidelberg.

Boss, Michael, Elsinger, Helmut, Summer, Martin, & Thurner, Stefan. 2004b. Network topology of the interbank market. *Quantitative Finance*, **4**(6), 677–684.

Cont, Rama, Moussa, Amal, & Santos, Edson B. 2013. *Network structure and systemic risk in banking systems*. Cambridge University Press. Pages 327–368.

Craig, Ben, & Von Peter, Goetz. 2014. Interbank tiering and money center banks. *Journal of Financial Intermediation*, **23**(3), 322–347.

Degryse, Hans, & Nguyen, Gregory. 2007. Interbank exposures: An empirical examination of contagion risk in the Belgian banking system. *International Journal of Central Banking*, **3**(2), 123–171.

Eisenberg, Larry, & Noe, Thomas H. 2001. Systemic risk in financial systems. *Management Science*, **47**(2), 236–249.

Elliott, Matthew, & Hazell, Jonathon. 2016. *Endogenous financial networks: Efficient modularity and why shareholders prevent it.* 2016 Meeting Papers 235,. Society for Economic Dynamics.

Elliott, Matthew, Golub, Benjamin, & Jackson, Matthew O. 2014. Financial networks and contagion. *American Economic Review*, **104**(10), 3115–53.

Elsinger, Helmut, Lehar, Alfred, & Summer, Martin. 2006. Risk assessment for banking systems. *Management Science*, **52**(9), 1301–1314.

Freixas, Xavier, Parigi, Bruno M, & Rochet, Jean Charles. 2000. Systemic risk, interbank relations and liquidity provision by the central bank. *Journal of Money, Credit and Banking*, **32**(2), 611–638.

Furfine, Craig H. 2003. Interbank exposures: Quantifying the risk of contagion. *Journal of Money, Credit and Banking*, 111–128.

Gai, Prasanna, & Kapadia, Sujit. 2010. Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **466**(2120), 2401–2423.

Gandy, Axel, & Veraart, Luitgard A M. 2016. A bayesian methodology for systemic risk sssessment in financial networks. *Management Science*, **63**(12), 3999–4446.

Georg, Co-Pierre. 2013. The effect of the interbank network structure on contagion and common shocks. *Journal of Banking and Finance*, **37**, 2216–2228.

Glasserman, Paul, & Young, H Peyton. 2015. How likely is contagion in financial networks? *Journal of Banking and Finance*, **50**, 383–399.

Halaj, Grzegorz, & Kok, Christoffer. 2013. Assessing interbank contagion using simulated networks. *ECB Working Paper*.

Hurd, Thomas, Gleeson, James, & Melnik, Sergey. 2017. A framework for analyzing contagion in assortative banking networks. *PLoS ONE*, **12**(2).

in 't Veld, Daan, & van Lelyveld, Iman. 2014. Finding the core: Network structure in interbank markets. *Journal of Banking and Finance*, **49**, 27–40.

Langfield, Sam, Liu, Zijun, & Ota, Tomohiro. 2014. Mapping the UK interbank system. *Journal of Banking and Finance*, **45**, 288–303.

Mistrulli, Paolo Emilio. 2011. Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance*, **35**(5), 1114–1127.

Nier, Erlend, Yang, Jing, Yorulmazer, Tanju, & Alentorn, Amadeo. 2008. Network models and financial stability. *Bank of England Staff Working Paper*.

Upper, Christian, & Worms, Andreas. 2004. Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review*, **48**(4), 827–849.

Wells, Simon. 2004. Financial interlinkages in the United Kingdom's interbank market and the risk of contagion. *Bank of England Staff Working Paper*.

6 Appendix

6.1 Standard Network Measures

Definition A1.1. ***Graph G = (V,E)*** *(network)* is a mathematical structure that consists of a set V of *vertices* (or nodes) and a set E of *edges* (or links), where elements of E are unordered pairs of vertices $\{u, v\} \in V$. The number of vertices and of edges is referred to as the network order and size, respectively.

Definition A1.2. A network can be represented by an ***adjacency matrix A*** with elements $a_{uv} \in \{0, 1\}$, where one means there is a link between nodes u and v and zero means these two nodes are not directly connected.

One of the main characteristics of an individual node in a network is centrality, i.e. its importance. The commonly used centrality measures include degree, closeness, betweenness and eigenvalue centrality.

Definition A1.3. ***Degree centrality*** measures importance of the node in terms of the number of links incident upon (connected to) it: more central node has higher degree. In a digraph node can have two degrees: in-degree and out-degree, which are number of incoming and outgoing links, respectively. And in case of weighted network one can also consider total weight of incident links (called node strength).

Definition A1.4. ***Closeness centrality*** measures node importance by capturing how close the node is to other nodes in the network. Closeness centrality of node v is computed as inverse mean of all shortest path distances from node v (more central node has higher value):

$$c_c(v) = \frac{V}{\sum_{u \in V} d_{sp}(u, v)}, \tag{2}$$

where d_{sp} is shortest path distance.

Definition A1.5. ***Betweenness centrality*** is the number of shortest paths in the network that pass through the given node:

$$c_b(v) = \sum_{j \neq u \neq v \in V} \frac{g_{sp}(j, u|v)}{g_{sp}(j, u)}, \tag{3}$$

where $g_{sp}(j, u|v)$ is the number of shortest paths between j and u that intersect with node v, and $g_{sp}(j, u)$ is the total number of shortest paths.

Definition A1.6. ***Eigenvalue centrality*** builds on the idea that the node is central if its neighbours are central, i.e. it is proportional to centrality of its neighbours:

$$c_e(v) = \alpha \sum_{u \in V} a_{uv} c_e(u) \tag{4}$$

This can be rewritten in a matrix form: $\mathbf{A}\mathbf{c}_e = \lambda \mathbf{c}_e$, where $\mathbf{c}_e = (c_e(1), c_e(2), ...)$ is vector of centralities and λ is a constant which is equal to $\frac{1}{\alpha}$.

Definition 1.7. ***Path*** $\gamma_{v_n v_0}$ is an ordered sequence of $n(\gamma_{v_n v_0})+1$ nodes $(v_0, v_1, ..., v_n)$ and $n(\gamma_{v_n v_0})$ links between them $((v_0, v_1), (v_1, v_2),..., (v_{n-1}, v_n))$ that form a connection between v_0 and v_n. No node is repeated more than once.

In a connected network there is a path between any pair of nodes.

Definition A1.8. ***Shortest path*** $\gamma_{u,v}$ is a path between two nodes such that there exists no other path that is shorter. And the *length of the shortest path* $n_{sp}(u,v)$ is the number of links comprising it.Shortest paths have no loops and do not have to be unique.

Definition A1.9. ***Shortest path distance*** $d_{sp}(u,v)$ minimizes the sum of link costs $c_{u,v}$ along the path $\gamma_{u,v}$, where the link costs can be represented as inverse link weights:

$$d_{sp}(u,v) = \min_{\gamma_{u,v}} \sum_{(u_i,u_{i-1})\in E_\gamma} c_{u_i u_{i-1}} \tag{5}$$

Using shortest path distance provides additional insights into the analysis of financial network, as it accounts for inverse link weights (unlike shortest path length). Heavier links, i.e. larger exposures, would mean shorter distances between nodes. This in turn would mean that if a given bank defaults, its counterparties with higher exposures to it will incur larger losses and will be more likely be affected.

Definition A1.10. ***Density*** of a network *Graph* $G = (V, E)$ is the fraction of existing links that are actually present, i.e.

$$density = \frac{|E|}{\binom{N}{2}} = \frac{2|E|}{N(N-1)}, \tag{6}$$

where N is number of nodes in the network. Network with N nodes is called *complete* if it has N(N-1) links, i.e. $density = 1$ each node is directly connected to every other node in the network. If $density << 1$, the network is called *sparse*.

Definition A1.11. ***Clustering coefficient*** measures the probability that two nodes connected to a third are also connected to each other. In other words it is a ratio of number of complete triplets to the number of connected triplets.

Definition A1.12. ***Degree-degree correlation (assortativity)*** is an important feature describing the manner by which the nodes connect to each other. In particular, whether nodes tend to connect to other nodes with similar degree or not.

Positive degree correlation (assortative network) implies the tendency of the high (low) degree nodes to connect to nodes with similar degree, i.e. to other highly (weakly) connected nodes. *Negative* degree correlation (disassortative network) implies that the hubs are likely to avoid other hubs, and on contrary to link with low-degree nodes. Finally, in case of a neutral network it is difficult to detect any pattern as association seems to be rather random.

6.2 Additional Figures

Figure 19: Core-Periphery Network Density

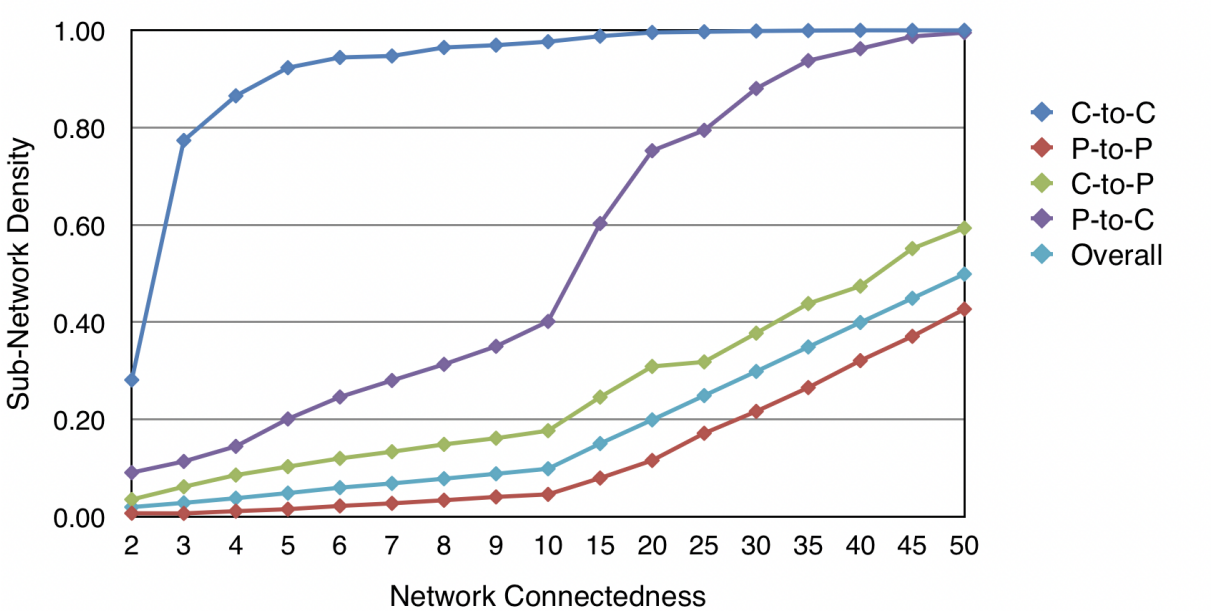

Densities of various sub-networks (sub-matrices) within core-periphery network along with the overall network densities for different levels of connectedness. The sub-networks are as per schematic illustration in Figure 20.

Figure 20: Core-Periphery Matrix Schema

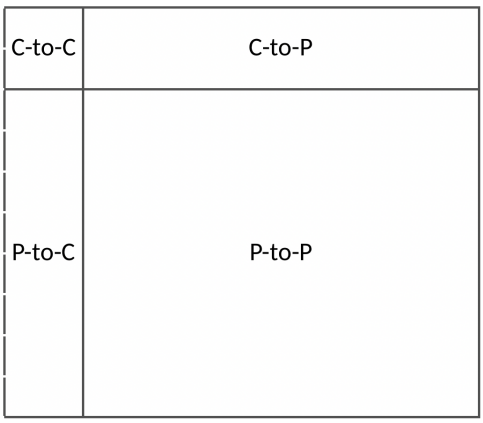

The schematic illustration of core-periphery network exposures matrix. The four areas indicate sub-networks of exposures: C-to-C (core-to-core), P-to-P (periphery-to-periphery), C-to-P (core-to-periphery) and P-to-C (periphery-to-core). For instance, C-to-P contains exposures of core banks to periphery banks, i.e. liabilities of periphery banks to core banks.

Figure 21: Contagion Losses in Scale-Free Networks with Homogeneous Bank Balance Sheets

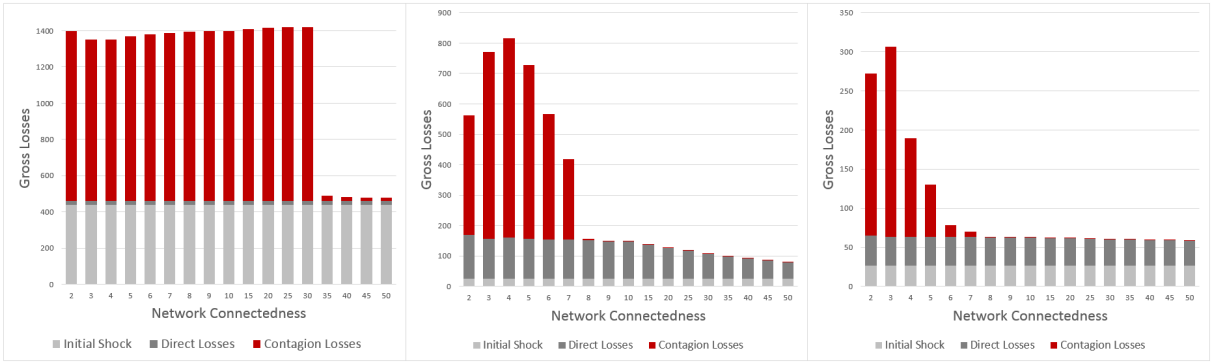

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

Figure 22: Contagion Losses in Erdos-Renyi Networks with Homogeneous Bank Balance Sheets

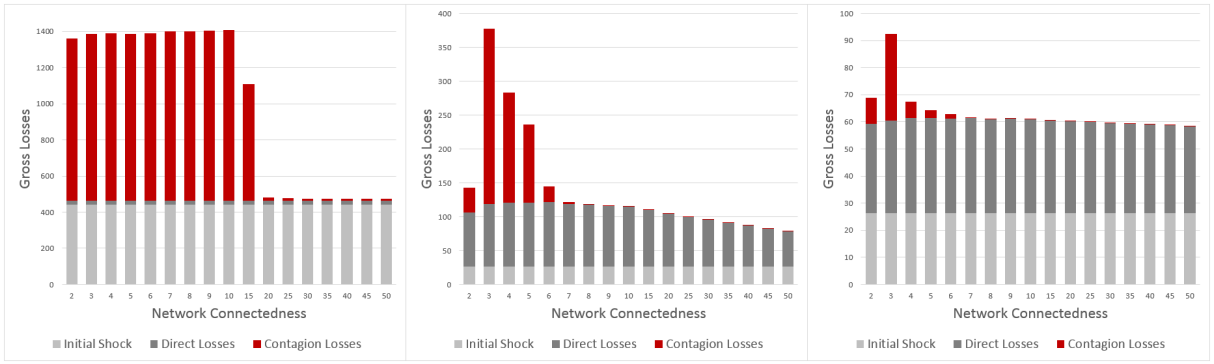

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

Figure 23: Contagion Losses in Scale-Free Networks with Heterogeneous Bank Balance Sheets

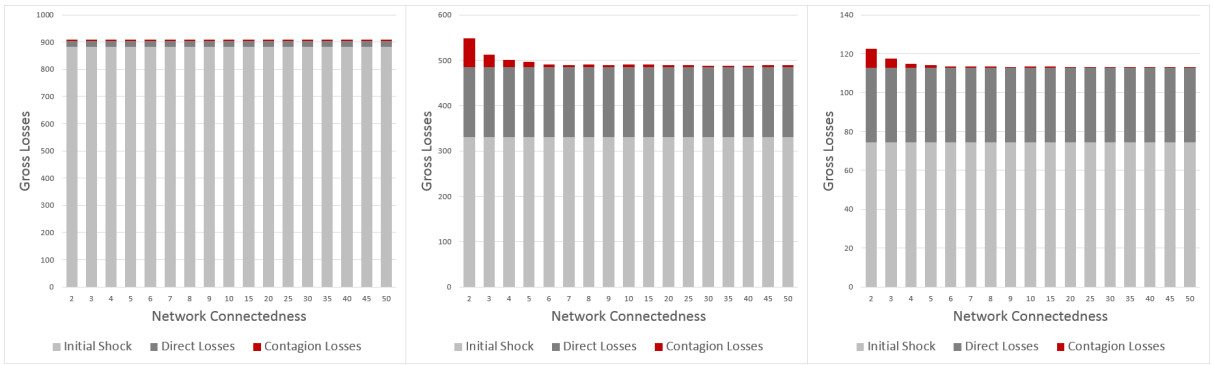

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

Figure 24: Contagion Losses in Erdos-Renyi Networks with Heterogeneous Balance Sheet Banks.

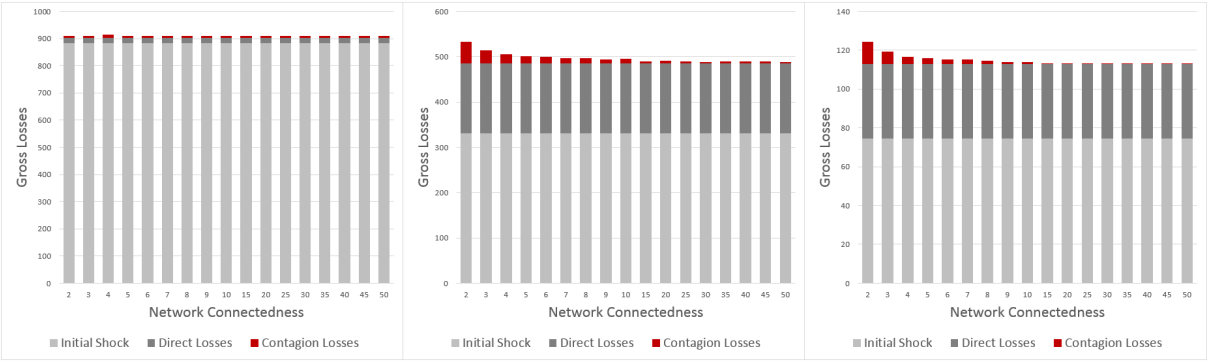

Left: system-wide shock (50% shock to capital buffer). Middle: individual bank defaults in the core (3 core banks' defaults). Right: individual bank defaults in the periphery (3 periphery banks' defaults).

6.3 Regression Results for Core-Periphery Networks with All Densities

Table 2: Regression Results for Core-Periphery Networks with All Densities

Regressor	System-wide shock		Defaults in Core	
	Homogeneous banks	Heterogeneous banks	Homogeneous banks	Heterogeneous banks
Intercept	**20.654*** **	**1.465*** **	0.128	**10.106*** **
	(0.000)	(0.000)	(0.666)	(0.000)
In-degree	**0.165*** **	**-0.003*** **	**0.296*** **	**0.524*** **
	(0.000)	(0.000)	(0.000)	(0.000)
Betweenness	**0.035*** **	**-0.001*** **	**-0.810*** **	**3.357*** **
	(0.000)	(0.000)	(0.000)	(0.000)
KNN	**-0.075*** **	**-0.001*** **	**-0.131*** **	**-0.087*** **
	(0.000)	(0.000)	(0.000)	(0.000)
Eigenvector Centrality	**-10.810*** **	**0.691*** **	**-18.031*** **	**-28.362*** **
	(0.000)	(0.000)	(0.000)	(0.000)
Clustering	**-8.084*** **	**-0.083*** **	**0.670***	**-5.465*** **
	(0.000)	(0.000)	(0.077)	(0.002)
Assortativity	**-1.129*** **	**-0.059*** **	**-7.638*** **	-1.181
	(0.002)	(0.000)	(0.000)	(0.244)
Shortest Path	**-4.237*** **	**0.081*** **	**1.257*** **	**-3.093*** **
	(0.000)	(0.000)	(0.000)	(0.000)
	Node-Level Metrics			
In-degree			**0.021*** **	**0.053*** **
			(0.000)	(0.000)
Out-degree			**-0.036*** **	**-0.058*** **
			(0.000)	(0.000)
Eigenvector Centrality			0.001	**0.030*** **
			(0.445)	(0.000)
KNN			**0.035*** **	**-0.071*** **
			(0.000)	(0.000)
Betweenness			**0.069*** **	**0.062*** **
			(0.000)	(0.000)
R^2	86.29%	14.41%	97.19%	33.48%
N	8500			

Parameter estimates' p-values are given in parentheses. $^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$.