Bank of England

Network analysis of the UK reinsurance market

Staff Working Paper No. 1,000
October 2022

Artur Kotlicki, Andrea Austin, David Humphry, Hannah Burnett, Philip Ridgill and Sam Smith

Bank of England

Staff Working Paper No. 1,000

Network analysis of the UK reinsurance market

Artur Kotlicki,(1) Andrea Austin,(2) David Humphry,(3) Hannah Burnett,(4) Philip Ridgill(5) and Sam Smith(6)

Abstract

We provide an empirical analysis of the network structure of the UK reinsurance sector based on 2016 Solvency II regulatory data. We examine counterparty credit risk originating from reinsurance contracts as a source of financial contagion in the insurance industry. The granularity of the Solvency II data provides a new opportunity for detailed analysis of the actual exposures in the system, detection of potential systemic vulnerabilities, and reinsurance spirals. In our multi-layered network approach, we incorporate information on reinsurance contract risk types and ownership structure for both life and non-life insurers.

Our findings suggest that the UK reinsurance sector exhibits the 'small-world' property with a scale-free, core-periphery structure and topological characteristics common to other financial networks. These characteristics of risk dispersion from the periphery to the core make the network 'robust-yet-fragile' to financial shocks. We explore the robustness of the network to adverse shocks through a stress-simulation exercise, where we find it robust to system wide shocks affecting the value of total investments, and to idiosyncratic shocks applied to large, highly interconnected reinsurers.

Key words: Reinsurance, systemic risk, financial contagion, scale-free network.

JEL classification: D85, G01, G22, G28.

---

(1) Bank of England. Email: artur.kotlicki@bankofengland.co.uk
(2) Australian Energy Regulator. Email: andrea.morris@aer.gov.au
(3) Bank of England. Email: david.humphry@bankofengland.co.uk
(4) Bank of England. Email: hannah.burnett@bankofengland.co.uk
(5) Bank of England. Email: philip.ridgill@bankofengland.co.uk
(6) Bank of England. Email: sam.smith@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. The authors would like to thank Rama Cont, Carol Alexander, Hanqing Jin, Gerardo Ferrara, Arjun Mahalingam, Dimitris Papachristou, and Peter Sondhelm for valuable comments, and Georgia Papamarkidou for data preparation and analysis. The authors would also like to thank participants at Bank of England internal seminars and participants at EIOPA's Seminar on Data for Risk-Based Supervision, hosted by the Central Bank of Ireland, 2019, for helpful discussion and comments. All errors remain those of the authors.

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

ISSN 1749-9135 (on-line)

1 Introduction

The aim of this study is to use network analysis to assess empirically the view that reinsurers do not present a risk of counterparty credit contagion from their default. We examine the network structure of the UK insurers' reinsurance contracts, reviewing the nature of the interconnections, whether risk is dispersed or concentrated, and the implications for financial stability.

Our study builds on the growing literature of network analysis applied to the financial sector which gained impetus following the 2008 financial crisis. Network models provide an adequate framework for assessing the potential for losses to spread in a financial system following counterparty default (Cont et al., 2013). In particular, the topology of networks affects their vulnerability to the risk of contagion (Allen and Gale, 2000; Caccioli et al., 2012; Roukny et al., 2013). A nonlinear relationship between the interconnectedness and the stability of financial markets, characterised as the 'robust-yet-fragile' property, has been well documented in the literature on financial contagion (Gai and Kapadia, 2010; Gai, 2013; Acemoglu et al., 2015; Caccioli et al., 2015) – while increased network connectivity and diversification of exposures can reduce the likelihood of contagion, it can also amplify losses when contagion occurs. In other words, financial networks exhibit a phase transition with respect to adverse shocks: below a certain threshold, losses are attenuated through risk-sharing and diversification practices, while beyond this tipping point, interconnectedness provides the means for the contagion to spread, amplifying the financial stress experienced in the system.

The literature on financial networks highlights the importance of key network characteristics for network stability and the global level of systemic risk (Amini et al., 2012; Roukny et al., 2013; Nier et al., 2007; Battiston et al., 2012a). Notably, financial networks can often be characterised by a heavy-tailed power law degree distribution of connections (Boss et al., 2004; Cont et al., 2013; Caccioli et al., 2015), and a core-periphery structure (Chen et al., 2020; Fricke and Lux, 2015; Craig and von Peter, 2014; Caccioli et al., 2018). As such, the network structure is far from being random, where strong hierarchical relations are indicative of firms' preferences when choosing their counterparties to conduct business. In particular, a core-periphery structure is characterised by a set of highly connected hub nodes in the core that intermediate connections between peripheral nodes, where nodes in periphery tend not to connect with each other. Barabási and Albert (1999) show that the emergence of a core-periphery structure in real-world networks can be attributed to preferential attachment, where new nodes in the network are more likely to form a connection with already well-connected connected nodes in the system.

Nodes in the core of the network play an important part from the systemic risk perspective: through their diversification and risk-sharing practices, hub nodes aid in the mitigation of losses following a counterparty default. However, they also render the system more vulnerable to their own default, as the resulting financial distress can propagate to a large part of the system, triggering systemic defaults. Consequently, networks exhibiting the core-periphery structure are shown to be resilient to the failure of a random node, but are vulnerable to defaults of systemically important hub nodes (Caccioli et al., 2012; Roukny et al., 2013). In contrast to the 'too-big-to-fail' theory, which considers only the absolute size of an institution, network theory is more suited for an identification of systemically important nodes that are 'too-central-to-fail' (Battiston et al., 2012b). In particular, simulation studies provide useful tools in assessing the robustness and stability in financial markets, as well as identifying systemically important nodes whose default can trigger a cascade of losses.[1]

Attention has only turned in recent years to the role of the insurance sector in creating or propagating systemic risks. Insurance risk is generally considered 'uncorrelated' by design. This means that the occurrence of an insured event does not increase the likelihood of a separate insured event happening. This is in strong contrast to issues related to credit risk, for instance in bond insurance[2] and CDS contracts, where the inability of an entity to pay their liabilities makes it more likely that other entities within the system will also be unable to pay theirs. In consequence, resulting in a default being much more correlated across the financial system, and therefore rendering the system more susceptible to a systemic event. Furthermore, insurers enable risk transfer in the economy through accepting and pooling risks, and through providing long-term savings products. The subsequent investment of premiums by insurers in financial assets can make them susceptible to systemic risk – risk management practices such as balance sheet diversification are thus a core activity of insurers. An example of which involves purchasing reinsurance, where the reinsurer accepts part of the insurance losses of the insurer in return for a premium.[3]

In 2016 the International Monetary Fund described two channels through which systemic risk might spread through the insurance sector. One channel is through counterparty default – a 'domino' view, where the failure of one insurer triggers the failure of others (International Monetary Fund, 2016). Factors that affect how this channel operates in practice include the size of the insurer, its interconnectedness, whether its activities can be substituted by other insurers, its leverage, its funding

[1]Refer to Upper (2011) for a detailed survey on the work in this topic.

[2]Insurance of credit risk, such as the risk of municipal bonds is a notable exception. Prior to the financial crisis in 2018 several municipal bond insurers also provided credit guarantees for mortgage-backed securities, creating an exposure to the housing market; see Saporta (2016).

[3]For instance, Bäuerle and Glauner (2018) present a model in which a single reinsurer providing excess of loss reinsurance contracts is able to reduce the value at risk for a group of insurers.

liquidity risk, and the complexity of its operations. An often cited example, albeit caused by products not traditionally associated with insurance, is that of the near-failure of American International Group (AIG), which received support from the US government due to the potential default of its CDS counterparties (McDonald and Paulson, 2015). Insurers face counterparty credit risk from their reinsurers, creating the potential for one 'domino' – a reinsurer – to knock over several others if it is unable to pay its reinsurance claims.

The other channel for systemic risk, which does not rely on firm failure, is the 'tsunami' view – capital weaknesses can stop insurers from performing the role of taking on risk during crises. This could impact economic activities such as the availability of insurance cover and the availability of funding to the economy through investments,[4] which could further contribute to financial market turbulence through pro-cyclical investment behaviours (Ellul et al., 2015). In particular, Cont et al. (2020) quantify the impact of funding liquidity availability in the market on the amplification of equity losses.

Insurers carrying out the traditional role of risk transfer from perils or mortality tend to be regarded as not posing systemic risk, at least from the perspective of the domino view. However, reinsurance stands out within the insurance business model as a source of counterparty credit risk. This type of transaction creates the risk of direct contagion – a domino effect – between insurers. The generally accepted view in the literature is that the use of reinsurance in practice does not, however, give rise to systemic risk (IAIS, 2011, 2012; French et al., 2015). This is because of five main factors:

1. Insurers only cede a small proportion of their total reinsured risks to each reinsurer, so losses from individual failures are more likely be absorbed by the insurer.

2. Collateral may be posted by reinsurers to mitigate the effect of their failure.

3. Insurers make conscious choices to cede certain risks that allow for greater diversification benefits and improved capital management (therefore not contributing to the build-up of risk in the system).

4. The links between insurers are hierarchical and do not give rise to feedback loops. For example, the International Association of Insurance Supervisors noted in 2011 that 'while primary insurers link to reinsurers, interlinkages among primary insurers are comparatively limited. In other words, links between entities in the insurance market are almost entirely hierarchical,

[4]In a recent study, Malik and Xu (2017) examine the interconnectedness among global systemically important banks and global systemically important insurers for US, European and Asian regions in the period of 2007 to 2016.

and there is no network-like inter-insurance market similar to the interbank market. [...] As a result, there are fewer feedback mechanisms to create non-linearity and a potential for systemic risk within the insurance sector.'

5. The timing difference for liquidity between insurance and banking makes counterparty exposures less problematic. Insurance claims are paid out over a much longer period of time, possibly years, than, say, the interbank loan market, which may require settlement overnight. This gives reinsurers in distress more opportunity to recover before its ceding insurers are affected.

This consensus is supported by historical experience, where less than 4% of all impaired insurers failed as a result of the failure of a reinsurer (IAIS, 2012). There is, however, one notable exception – the London Market Excess of Loss spiral, which took place in the late 1980s and inadvertently saw larger losses arise than anticipated because of retrocession (Bain, 1999). In this case, the reinsurance contracts between insurers and syndicates within the Lloyd's of London Market saw syndicates cede risk only to accept that same risk through another contract. Such retrocession spirals aggregate losses from multiple parts of the market in an opaque manner, potentially leaving the dampening reinsurer with extreme losses and thus destabilising the entire system (Klages-Mundt and Minca, 2020). Since then, the use of reinsurance and retrocessions by syndicates has changed a great deal, with lower proportions of reinsurance, higher levels of retained risk, and other measures put in place by the Society of Lloyd's to identify these risks.

Empirical research on financial contagion from reinsurance, whilst limited in scope,[5] finds no evidence of systemic risk due to counterparty defaults in the insurance market (that is, the 'domino' view). Given the opacity of the reinsurance industry and scarcity of data, most recent studies focus on the property-causality reinsurance market in the United States (Park and Xie, 2014; Chen et al., 2020; Klages-Mundt and Minca, 2020).

In particular, Park and Xie (2014) consider two potential contagion mechanisms in the period of 2003–2009: the direct counterparty risk due to failure of top reinsurers, and an indirect information-based effect of a reinsurer's downgrade. The latter effect contributes to loss spill-overs even to insurers with no direct exposure to the downgraded reinsurer. However, the worst-case scenario in the simulation study, comprising of a failure of the top reinsurer group, had a minor effect on the solvency of insurance market participants. Therefore, despite the

[5]Research in this field is often hindered by data limitations: parts of the market remain opaque to both regulators and market participants (Davison et al., 2016). In particular, disclosures on retrocession agreements are scarce in contract details, making estimation of counterparty exposures difficult in practice. We refer to Appendix D for a discussion on Solvency II data limitations in the context of identification of network cycles.

strong interconnectedness of the system, the likelihood of systemic risk triggered by reinsurance contracts is found to be small.

Chen et al. (2020) extend the study of interconnectedness within the US property-casualty reinsurance market with more extensive and granular data for both insurers and reinsurers. The contagion mechanism includes both counterparty exposures due to reinsurance premiums paid and reinsurance recoverable amounts outstanding. In a simulation analysis, the authors consider both the failure of individual reinsurers – the domino view – and the potential amplification of systemic stresses emanating from asset markets via the default of reinsurers because of deterioration in their financial assets. Similar to the previous studies, they do not observe widespread insolvencies in the industry even during very severe stress scenarios.

An empirical analysis by Lelyveld et al. (2011) examines the resilience of Dutch insurers to failing reinsurance covers in the period of 2003–2005. On the individual institution level, no default cascade occurred in their scenario analysis. However, there was a risk to entities within a group from intra-group reinsurance contracts.

Kanno (2016) assesses the systemic importance of insurers in the global non-life insurance market during 2006–2013. Using the Eisenberg and Noe (2001) approach for the allocation of losses, the author finds systemic risk to be relatively restricted: only a few contagious defaults were observed under a severe initial shock to the economy. The network analysis of contagious defaults was limited, however, due to the use of aggregate data on reinsurance transactions.

A noteworthy exception is the work of Klages-Mundt and Minca (2020), in which the authors caution that the risk of contagion may be underestimated because of the potential for unexpectedly large losses from non-proportional reinsurance. They emphasize the existence of complex interactions between insurance losses and counterparty default, which are not captured adequately by the existing contagion models (Eisenberg and Noe, 2001; Acemoglu et al., 2015; Elliot et al., 2014). In particular, non-linearities from excess of loss contracts are shown to obfuscate risks, rendering the system vulnerable to excess costs from network effects. However, the methodology of Klages-Mundt and Minca (2020) is extremely sensitive to model parameters. In their simulations on the 2012 US property-casualty reinsurance data, small perturbations in the estimated parameter set are shown to affect key players of the market in a substantial manner. As a result, this methodology is not well suited to the current data regime, where the required contract details are not available and instead need to be estimated using common rules of thumb.

Similarly, Davison et al. (2016) emphasise the role of market opacity in a possible emergence of reinsurance spirals that can increase the risk of contagion due to concentration of losses. However, their simulations of the frequency, severity, and patterns of exposure to loss find that reinsurance networks are robust to plausibly large losses. That said, different patterns of retrocession and large haircuts to recoverable amounts can influence the levels of contagion. These results should be taken with caution however, as the data on size and connections of the network is not available to the authors, who resort to testing several plausible networks based on a sample of large global reinsurers.

Our paper is closest to the work of Chen et al. (2020), however the differentiating feature of our network analysis of the insurance market is our ability to measure the interconnectedness directly through the use of Solvency II regulatory returns that contain details on reinsurance contracts between insurers, including premiums, sums insured, and amounts recoverable. It is the first study of the UK insurance sector of this kind, and unlike previous studies that have focused on property-causality underwriters, our work incorporates both life and non-life insurance contracts. In addition, our analysis includes syndicate-level data from the Lloyd's of London market to give a detailed view of the interconnectedness of the London insurance market. We use line of business data to identify the prevalence of cycles of risk transfer within the network.

Our simulation-based stress test analysis focuses on a direct contagion through exposures via reinsurance contracts, and incorporates bankrupcy costs. Similarly to Chen et al. (2020), we consider both the contagion impact of a network wide shock and the impact of default of the most connected insurers. Our results show the UK reinsurance market to be robust even under extreme stress scenarios. We also show that the network topology shares common characteristics with other financial networks, and in particular can be characterised by a scale-free, core-periphery structure.

The remainder of this chapter is organised as follows: in Section 2 we introduce data sets considered in our analysis, we describe the adjustments made to them, and provide detailed description of the methodology used in our network analysis. Section 3 then presents the main results of our the network analysis, including the topological characterisation of the UK reinsurance market, and observed nuances which affect its vulnerabilities to contagion risk. Section 4 presents simulation analysis to assess the degree of systemic risk and potential for network contagion due to direct counterparty risk. We present our conclusions in Section 5.

2 Datasets and Methodology

We use for our study Solvency II data on reinsurance contracts submitted to the Prudential Regulation Authority (PRA) by authorised insurance companies and Lloyd's of London syndicates, covering a single year, recorded at year-end 2016.

We note that for simplicity, throughout the study we use the term *insurer* to refer to the company that is ceding risk, but also to refer to insurance companies generally when we do not want to specifically refer to reinsurers. We use *reinsurer* to refer to the company that is accepting the risk. The same company may be referred to as either an insurer or a reinsurer within this study depending on the context of the discussion. We also do not make an explicit distinction between reinsurance and retrocession (that is, reinsurance of reinsurance business).

The data submitted to the PRA is non-public, and thus in this paper we only present anonymised or aggregated results. We identify an insurer with the following characteristics where relevant for the analysis: whether the insurer is an insurance company or a Lloyd's of London syndicate, where the insurer is a life or non-life insurer, and whether the insurer is an insurance group.

2.1 Data Overview

For the purpose of network analysis, we construct two data sets for year-end 2016: one on the nature of contracts in place (referred to as the *treaty and facultative data*), and the other on the amounts recoverable (referred to as the *recoverables data*), with metadata about line of business and risk description added to allow for layering of the network and more in-depth analysis of the results.

Treaty and facultative data The first set of data includes information about facultative and outgoing treaty reinsurance programmes by contract as of year-end 2016. *Facultative reinsurance* is transacted on an individual risk basis, where the ceding company has the option to offer individual risks to the reinsurer and the reinsurer retains the right to accept or reject the risk. *Treaty reinsurance* on the other hand is a transaction encompassing a block of the ceding company's book of business. The reinsurer must accept all business included within the terms of the reinsurance contract.[6]

Our data on facultative and treaty reinsurance contracts (henceforth 'treaty and facultative data') captures reinsurance contracts with individual insurers, where a UK insurer has ceded risk to another (not necessarily UK-based) insurer. It covers the identity of the reinsurer, the sum insured, the premium, the line of business, and a description of the risk insured. Furthermore, the data records whether the contract is proportional, such as a quota share, where the reinsurer has a liability for a percentage of the loss, or non-proportional, such as excess of loss, where the reinsurer has a liability for losses that exceed a certain amount. Other information in the data includes the geographic location of the reinsurer, and whether the reinsurer is external to the insurer's group of companies, or internal to it, such as a subsidiary or a captive reinsurer.

[6] See Munich RE (2010) for a concise introduction to facultative and treaty reinsurance.

Concerning the line of business, Solvency II requires insurers to identify contracts as one of eight lines of business for life and 28 lines of business for non-life. We aggregate the treaty and facultative data into fewer lines of business based on similarity of the risks they cover.[7] 'Fire and damage to property' is the most common type of non-life reinsurance, while 'multiline' contracts, which cover multiple risks, are the ones with the highest premium ceded on average (in terms of median value). The 'other life' category, which includes reinsurance contracts against longevity risk, is the most frequent type of life reinsurance contract. 'Unit-linked or index-linked' reinsurance, which combines insurance and investment into a single integrated plan, and 'other life' reinsurance are the lines of business that have the highest premium ceded for the life market. Figure 1 provides a quantitative summary of premiums ceded by each line of business in the treaty and facultative data set. Reinsurance contracts are seen to be highly heterogeneous, with a small number of contracts reporting extreme values of exposures. Such heterogeneity in counterparty exposures is a common feature of financial networks, as emphasised in the past studies on banking systems (Caccioli et al., 2012; Cont et al., 2013; Boss et al., 2004).

The treaty and facultative data set includes 799 reinsurers and 41,883 contracts (82% of UK contracts). We do not have sight of all facultative contacts because as part of Solvency II reporting insurers are only required to disclose the 10 largest contacts by exposure for each line of business.[8] Consequently, over 93% of the treaty and facultative data come from treaty contracts (by number of contracts). Treaty contracts are in general more valuable than facultative contracts in terms of premiums ceded, with more than 99% of UK premiums attributed to treaty reinsurance.

Whilst there are fewer life contracts than non-life contracts, they account for a much higher proportion of premiums ceded – out of the observed £107 billion premiums ceded, £69 billion is life and £38 billion is non-life. This is because unit-linked life insurance, which is mainly a type of investment product, can be organised as a reinsurance contract by the insurer. Non-proportional contracts, such as excess of loss, are more common than proportional contracts, accounting for more than 81% of all contracts in the data set. This feature tends to follow the life and non-life distinction – life contracts tend to be proportional (87% of contracts), while non-life contracts tend to be non-proportional (84% of contracts). As a result, although non-proportional contracts are more common, they account

[7]The aggregated lines of business are: fire and other damage to property; marine, aviation and transport, general liability, motor, non-life annuities; credit and suretyship, health, medical expense; multiline; other non-life; life; with-profit participation; unit-linked or indexed-linked; other life.

[8]The number of disclosures is sometimes lower than 10 per insurer, indicating that the insurer in this case has less than 10 contracts in total per line of business.

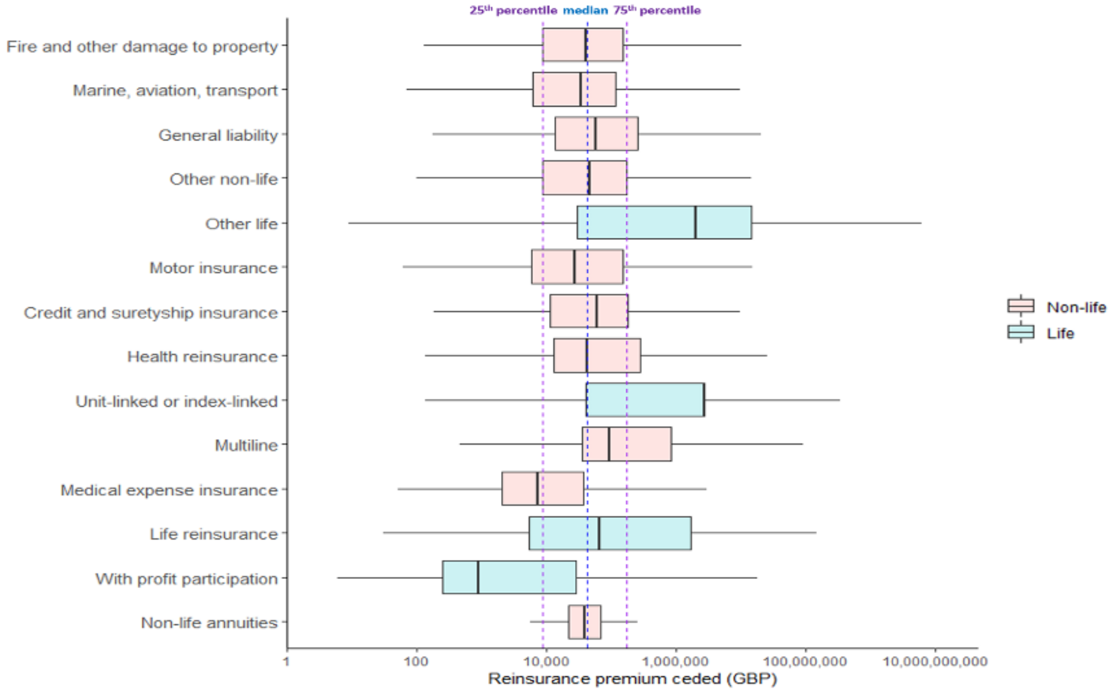

Figure 1: Box-plot of treaty and facultative premiums by line of business and by contract type. Blue dashed lines reflect distribution of all UK reinsurance premiums ceded. Values in GBP on a $\log_{10}$ scale.

for a lower share of premiums ceded. Most reinsurance contracts are external to the insurance group (93% of the insurance contracts) but internal contracts account for a large share of premiums ceded (81% of UK premiums). In particular, our data contains a few intra-group contracts with exceptionally large premiums ceded that significantly skews the data. Notably, values of these contracts, although unusually large, are confirmed on an individual contract basis to be valid and hence are not excluded from the analysis.

We refer to Appendix A for a detailed breakdown of the number and value of reinsurance contracts by their type and line of business.

In general, premiums ceded by contract follow a positively skewed distribution. The mean is greater than the median: that is, there is a minority of high-value contracts that pull-up the mean. This characteristic is maintained whether we look at all contracts together, or other categorisations, such as life and non-life contracts separately (see Table 1).

	All	Life	Non-life	Groups
Mean	2,571,399	34,816,994	951,842	2,153,463
Standard deviation	62,452,612	254,020,332	28,320,323	51,695,191
25^{th} *percentile*	8,634	17,589	8,437	8,696
Median	43,125	1,383,006	40,893	42,000
75^{th} *percentile*	174,137	7,950,774	152,164	164,801

Table 1: Descriptive statistics for treaty and facultative contracts based on premium ceded. Values in GBP.

Recoverables data The second set of data, referred to as the 'recoverables data', includes current levels of reinsurance amounts recoverable from individual reinsurers by UK insurers as of year-end 2016. In particular, the data contains the total amount of recoverables per individual reinsurer, and may encompass the recoverables from multiple contracts and lines of business. By usual convention, we classify a recoverable amount from a reinsurer under life insurance if more than 50% of the technical provisions (that is, the liability amount) is for the life insurance business, and we classify as non-life otherwise.

This data set covers 22,713 separate recoveries by UK insurers from their reinsurers (88% of the UK total recoveries) and 3,560 reinsurers. The total technical provisions recoverable comes to £240 billion in 2016, of which life technical provisions account for 86% of this total. Netting off collateral, total recoverables amounted to £116 billion. As in the case of the treaty and facultative reinsurance data set, the distribution of recoverables is skewed, with some very high values of recoverables, mainly relating to life insurance (see Table 2).

	All	Life	Non-life	Groups
Mean	5,746,898	133,882,490	1,729,200	4,109,034
Standard deviation	159,434,728	865,804,686	30,237,584	118,965,451
25^{th} *percentile*	289	48,097	1,834	273
Median	13,761	1,287,026	26,567	13,186
75^{th} *percentile*	214,800	18,104,041	275,338	203,139

Table 2: Descriptive statistics from individual insurers of amounts of reinsurance recoverable per reinsurer (net of collateral). Values in GBP.

Data limitations Our data corroborates the view that reinsurance is an international industry. According to our treaty and facultative data, eight countries accounted for 95% of the reinsurance premiums ceded. The life market was more concentrated geographically with six countries accounting for 95% of the premiums ceded, while the non-life market was less concentrated (13 countries).

This brings to light a limitation of using these data sets for network analysis – each insurer submission only contains the reinsurance ceded by UK insurers, and does not identify where non-UK insurers have ceded risks to UK (re)insurers.

Quantifying bilateral exposures Instead of considering counterparty risk on an individual reinsurance contract basis, we are interested in quantifying the total bilateral exposures in the UK reinsurance market. In light of this, we aggregate multiple instances of unique reinsurance contracts between an insurer and the same counterparty. In particular, we quantify the total bilateral exposure using the available additional data on the reinsurance contracts and reinsurance recoverables as follows.

- For the treaty and facultative data set, we use the information on the outgoing premium ceded and the amount of exposure ceded or sum reinsured to compute the aggregate value of risk transfer from an insurer to its specific reinsurer as of year-end 2016. When discussing insurance risks for a specific line of business (a particular network layer), we quantify the aggregate exposure using only the relevant contracts, and disregard other reinsurance contracts.

- For the recoverables data set, we use the value of the total reinsurance recoverable net of collateral to compute the aggregate counterparty default exposure from a reinsurer to its primary insurer. In particular, this value thus represents the maximal (short-term) loss to the insurer in case of an immediate default of its reinsurer counterparty as of year-end 2016.

We perform adjustments to the data prior to assigning counterparty exposure size to each market participant. We exclude missing weight values in the treaty and facultative data set and in the recoverables data (that is, blank fields). Contracts with a specified zero value may be contracts that start in a neutral position and whose value can change over time,[9] or may be the result of reporting error. Such contracts are excluded from our main network analysis as they do not represent a current exposure.[10]

The reporting templates for the treaty and facultative data allow insurers to record the sum reinsured or exposure ceded as -1 to represent unlimited liability, which are a potential source of high losses incurred by the reinsurer. In our study,

[9]For instance, some longevity reinsurance contracts have this feature.

[10]In our network analysis, we also computed unweighted topology statistics of the network that includes contracts with a value of £0, and find that these statistics are not materially different to those for the weighted network.

we convert unlimited contracts amounts to a right-tail value (97.5%) of the reinsurance recoverables distribution, representing an extreme but plausible scenario. We also remove contracts with premiums or sums insured that are significantly negative from the treaty and facultative data set, as we were unable to verify the validity of these contracts.

We do not exclude negative values from the recoverables data set as these are likely to represent reinsurance contracts containing additional performance clauses, which indicate that the insurer owes the reinsurer a payment. From the perspective of counterparty risk contagion, which relates to the failure of payment made by a reinsurer to its counterparty, it is important to include these contracts in the network. When computing the aggregate exposure, we include the absolute value of these contracts, and for any reinsurance contract with negative values we reverse the role of the insurer and its counterparty.

In total we include 5,349 unique contracts with a negative recoverable value, of which 159 are life contracts and 5,076 are non-life contracts.[11] We include further details on the number of contracts and performed data adjustments in Appendix C.

Figures 2a and 2b display the respective density plots of the aggregate counterparty exposures in the treaty and facultative data set,[12] and the recoverables data set. In both data sets there is a small proportion of very high value counterparty exposures, highlighting the skew to the data seen previously in the descriptive statistics.

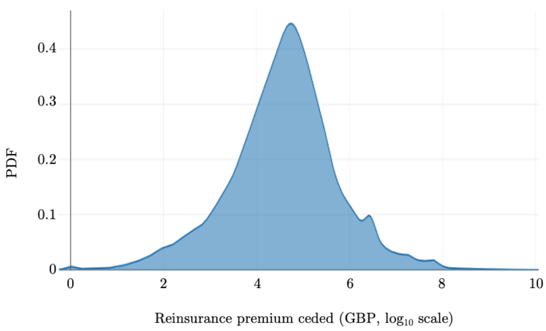

(a) Counterparty exposures in the treaty and facultative data set.

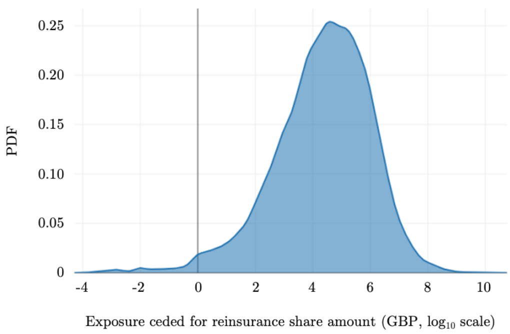

(b) Counterparty exposures in the recoverables data set.

Figure 2: Density plots of the aggregate counterparty exposures in the reinsurance recoverables (left) and treaty and facultative (right) data sets. Values in GBP on a $\log_{10}$ scale.

[11] In 114 cases it not possible to identify if the contracts are life or non-life.

[12] The density plot excludes approximately 3,000 treaty and facultative contracts for which the sum insured is unlimited.

Figure 3 provides further detail for the treaty and facultative data: the complementary cumulative distribution of counterparty exposures exhibits a linear decay on a logarithmic scale, suggesting a Pareto tail – emphasising there are few very large exposures, and many smaller ones. These results corroborate the view that the UK reinsurance market exhibits strong heterogeneity in reinsurance exposures; and are consistent with the results for other financial networks studied in literature (see, for example, Caccioli et al. (2015) and Cont et al. (2013)).

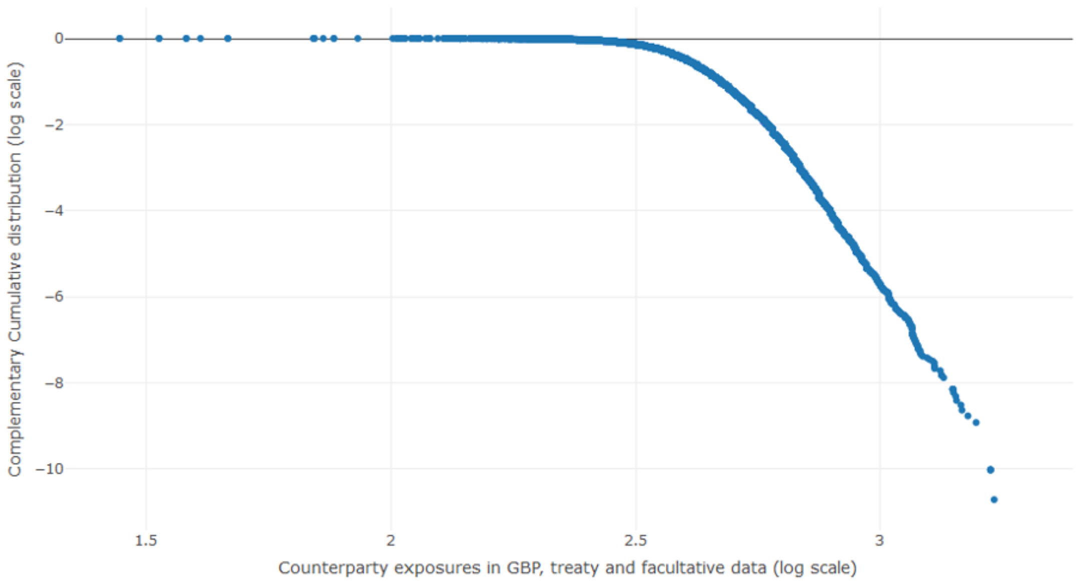

Figure 3: Complementary cumulative distribution plot of the aggregate counterparty exposures in the treaty and facultative data set. Values in GBP on a $\log_{10}$-$\log_{10}$ scale. Reinsurance exposures are strongly heterogeneous and exhibit a Pareto tail.

2.2 Identification of Insurance Market Participants

Under the Solvency II data regime, insurers are required to identify their reinsurance counterparty using both their name and a code – either the Legal Entity Identifier (LEI), or a Specific Code (SC) created by the insurer.[13] Across the treaty and facultative data and recoverables data there were approximately 14,900 unique names and code identifiers submitted. Of these, 3,100 had LEIs whilst the remaining 11,800 were submitted with SCs. Through a process of standardising

[13] Although Specific Codes are often used where an LEI is not available, their use is permitted more generally.

insurer names – using LEIs to verify names, and names to verify, or identify, LEIs – we are able to reduce the final number of reinsurance counterparties identified to approximately 5,100. If not for this standardisation of the list of reinsurance counterparties, three variants of the same name, for instance, would be counted as three separate reinsurers.

Insurer group structure Some insurers were also allocated to an insurance group where one could be identified. For UK insurers this was done using other information on group structure reported under Solvency II. For other insurers this was carried out primarily by name, combined with the use of public information, where possible. Where an insurance group could not be identified, the solo legal entity was treated as an insurance group. In total 1,276 solo insurers with non-zero reinsurance contracts were allocated to an insurance group in the recoverables data set, and 500 were identified for the treaty and facultative data set.

Treatment of Lloyd's of London One unique aspect of the UK insurance market is Lloyd's of London. Although often spoken of as a single entity, in reality it is a marketplace made up of many investors (members), which can include insurance groups, who bear the insurance risk that is allocated through syndicates underwriting insurance (see Figure 4), and which fall under centralised oversight by Lloyd's. Typically, Lloyd's syndicates subscribe to underwrite risks jointly. Syndicates of underwriters are managed on behalf of the members by managing agents, which employ the underwriters, and provide other services essential to the running of syndicates.

If the capital of individual investors is insufficient to bear losses, then there are sources of mutual capital to protect policyholders in the form of capital called from other investors at Lloyd's as well as a Central Fund, subject to approval by the Council of Lloyd's. Risk underwritten and capital are overseen by the Corporation of Lloyd's.[14]

Our network analysis identifies individual syndicates as individual entities in the market and, using supervisory and public information, we allocate them to the insurance groups that own them. This approach enables us to capture risks within an insurance group related to both the risk transfer through their insurance companies and the business they conduct through the Lloyd's market.

2.3 Reinsurance Networks

We represent the interconnectedness of the reinsurance market using a weighted directed network, where market participants are represented as nodes and their

[14]For further information see: `https://www.lloyds.com/about-lloyds/what-is-lloyds`.

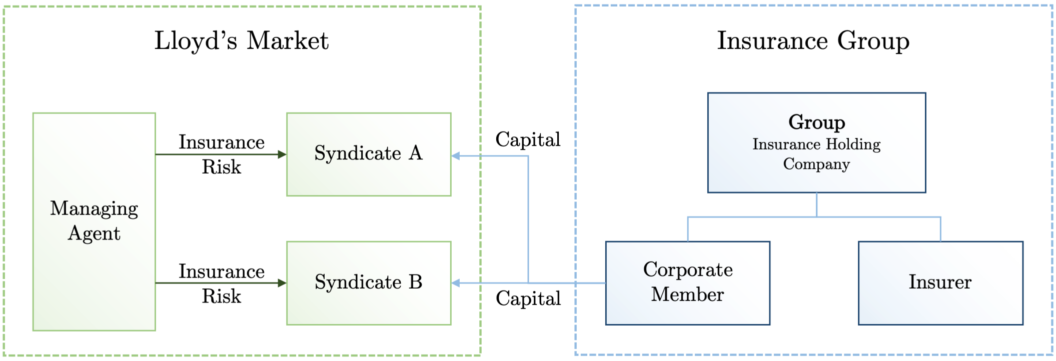

Figure 4: Simple representation of the Lloyd's of London market, and an example how it may be connected to an insurance company. Members of Lloyd's provide capital to underwrite policies through syndicates – formed by one or more members joining together to accept insurance risks. Corporate members include insurance groups that provide the majority of the capital for the Lloyd's market. A managing agent is a company set up to manage and oversee daily operations of one or more syndicates on behalf of the members.

interactions are defined by edges between them. We refer to Appendix B for a formal introduction of basic network theory concepts.[15]

Network layers: distinguishing reinsurance contracts For some types of insurance risk, the pattern of counterparty exposures may differ depending on the type of risk or its line of business. To explore these differences, we incorporate a multi-layered network approach in certain parts of the analysis, where we assign network layers to describe different types of interactions between the insurer and reinsurer. For example, at the level of lines of business, while the set of nodes for each layer remains unchanged, a given layer is defined to contain only a subset of links for a particular business line. Other works have found taking a multi-layer approach useful to illustrate certain characteristics observed in real world systems; for example, see Kivelä et al. (2014).

This approach allows us to assess the systemic risk beyond counterparty exposures; for example, by looking at the insurance risk transfer and conducting an analysis of potential retrocession spirals. Similarly, this approach enables us to perform an analysis of sub-networks that only consist of certain type of nodes, such as life or non-life insurance networks. In principle, a multi-layered framework

[15]The network analysis has been conducted using an interactive network analysis tool developed in R Shiny for the purpose of this study. Source code is made available at `https://github.com/bank-of-england/NetworkApp`; see Appendix E for more information.

allows for an added level of granularity, and enables testing of the resilience of the system to particular types of shock scenarios and risks.

3 The Network Structure of the UK Insurance Market

Characterising the structure of the network using topological measures provides insights about the interactions among the market participants, and important economic information about contagion risk and the stability of the system (Gai and Kapadia, 2010; Roukny et al., 2013; Caccioli et al., 2012; Amini et al., 2012).

Our empirical findings on the structure of the UK insurance network are in general quantitatively and qualitatively consistent for the networks that arise from both the treaty and facultative data set (see Table 3), and the recoverables data set (see Table 4). Therefore, we often do not make an explicit distinction between the data source but instead discuss topological properties of 'the reinsurance network' more generally. Furthermore, when discussing the network characterisation of the UK insurance market, we implicitly refer to the single (monoplex) network that captures all aggregate risk exposures. We also note that in general, network characteristics and topology remain consistent across different layers.[16]

3.1 Core-Periphery Structure

We find core-periphery relationships in the reinsurance data sets considered, with a small core of densely connected reinsurers, and other reinsurers in the periphery dispersing risk (see Figure 5 for a visual representation of our networks). Some UK insurers have a significant number of reinsurance contracts with reinsurers that no other UK insurer is using, which gives rise to a strong hierarchical relationship between market participants. These UK insurers act to extend the network by interacting with nodes outside the UK insurance market and disperse risk further.[17] This is in contrast to other financial networks where the periphery tends to be connected to a single common set of hubs (Fricke and Lux, 2015; Barucca and Lillo, 2016).

[16] We omit presentation of more detailed results so as to not disclose information that could lead to identification of particular firms in the case of smaller network layers.

[17] The nature of the scope of the data set, by not including the reinsurance contracts of these reinsurers, and therefore not knowing if risk is then ceded back into the UK insurance sector means we cannot draw too strong a conclusion here. However, at least for the purposes of understanding how UK reinsurance risk is ceded, these insurers do appear to extend the network of UK insurance risks and increase the dispersion of risk.

Network	All	Life	Non-life	Group
Density	1.08%	1.17%	1.23%	2.40%
Diameter	9	3	8	9
Average path length	2.752	1.501	2.69	2.552
Average clustering	0.302	0.03	0.306	0.482
Assortativity	-0.093	-0.216	-0.111	-0.208
Average betweenness	0.019	0.003	0.021	0.024
Average degree	19.18	3.77	20.21	18.42
Degree: power law exponent estimate	1.6804	2.2942	5.5594	1.6989
In-degree: power law exponent estimate	1.7114	2.9643	1.697	8.0265
Out-degree: power law exponent estimate	7.7807	2.0633	7.7807	5.0975
Link fraction (weighted) top 5% connected nodes	60.6% (81.4%)	49.8% (67.5%)	59.5% (66.0%)	59.5% (84.5%)

Table 3: Summary of network statistics based on the treaty and facultative dataset. Results shown for the full network including all reinsurance contracts, as well as network layers considering life contracts, non-life contracts, and contracts on a group-level only.

We find that life and non-life networks share the same characteristics as the network for the UK market as a whole – and in particular, a core and periphery structure. However, life reinsurance contracts tend to lead to larger exposures than non-life contracts. The non-life reinsurance network is denser than the life network, suggesting greater risk-sharing by non-life insurers. Moreover, risk-sharing by life insurers tends to be conducted via central reinsurers to a greater extent than for non-life insurers. We refer to Appendix D for further detail.

The degree distribution sheds more light on the heterogeneous structure of exposures and the network's resilience to financial contagion. Indeed, visual inspection of the network plots in Figure 5 reveals that some reinsurers have very many connections, while others have very few, suggesting the presence of hierarchical relationships. This observation is further supported by the double logarithmic plots of the empirical complementary cumulative distributions of degree and strength for the treaty and facultative data set, shown in Figure 6. We observe similar results for the recoverables network.

Moreover, the reinsurance network is characterised by heavy tailed degree distributions and negative degree correlations. We observe linear decay in the tails of the in-degree, out-degree and total degree distributions suggesting a heavy Pareto

Network	All	Life	Non-life	Group
Density	0.12%	0.45%	0.14%	0.13%
Diameter	9	9	8	6
Average path length	3.271	3.684	3.304	2.842
Average clustering	0.141	0.051	0.163	0.11
Assortativity	-0.346	-0.291	-0.381	-0.399
Average betweenness	0.11	0.018	0.109	0.117
Average degree	10.38	3.85	10.39	7.89
Degree: power law exponent estimate	1.7977	3.6299	11.761	1.9137
In-degree: power law exponent estimate	1.673	2.0865	2.0479	2.2482
Out-degree: power law exponent estimate	1.956	3.3398	1.9096	2.0775
Link fraction (weighted) top 5% connected nodes	98.0% (96.0%)	68.0% (90.0%)	97.8% (90.0%)	100.0% (99.0%)

Table 4: Summary of network statistics based on the recoverables dataset. Results shown for the full network including all reinsurance contracts, as well as network layers considering life contracts, non-life contracts, and contracts on a group-level only.

tail. In the treaty and facultative network for example, the maximum likelihood estimates of the exponent parameter are 1.7, 1.7 and 7.8 for the distribution of total degree, in-degree and out-degree respectively.[18] We recall that a network whose degree distribution $P(k)$ follows a power law, that is $P(k) \sim k^{-\gamma}$, is known as a scale-free network. A smaller estimate of exponent γ relates to a heavier tail in the distribution of connectivity. Our results are indicative of a structure that includes a relatively small number of hubs – that is, nodes that are very highly connected in the network – that form the main core of the network.

To assess the goodness-of-fit of our power law estimates, we use a one-sample Kolmogorov-Smirnov test at a 1% significance level. We do not find enough evidence to reject the null hypothesis that the original data could have been drawn from the fitted power-law distribution,[19] and hence conclude the UK insurance

[18]We implement the method of Clauset et al. (2009) to calculate the parameters of the fitted distribution. A summary of maximum likelihood estimates in the case of the recoverables data set and other network layers across both data sets is presented in Tables 3 and 4. In general, we find estimates in the usual range of 2–3.

[19]For the treaty and facultative data for example, the reported p-values for the test statistic are 0.32, 0.39 and 0.96 for total, in- and out-degree respectively.

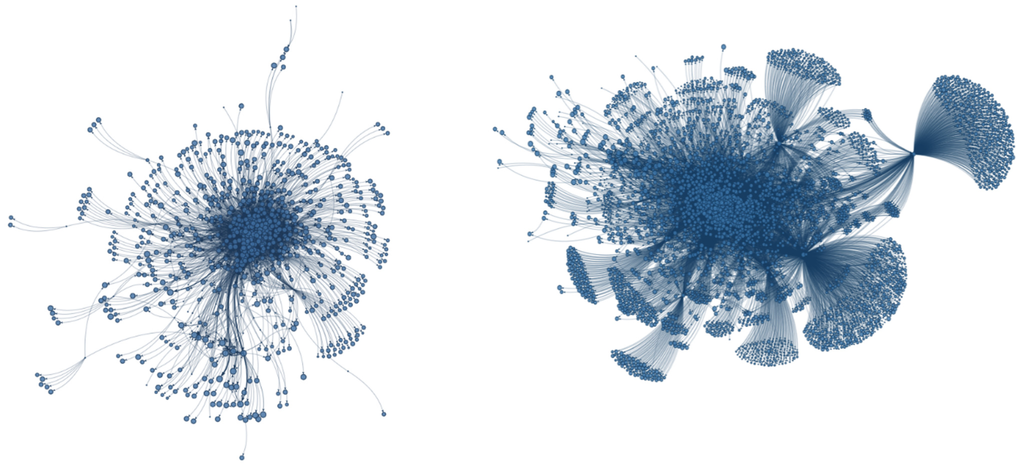

(a) Treaty and facultative data set. (b) Reinsurance recoverables data set.

Figure 5: Visualisation of UK reinsurance networks, year-end 2016.

market forms a scale-free network. Our results are consistent with the previous study of Chen et al. (2020) on the US property-casualty insurance, as well as with degree distributions observed in other financial networks (Boss et al., 2004; Degryse and Nguyen, 2004; Iazzetta and Manna, 2009; Cont et al., 2013; Caccioli et al., 2015).

We also observe that highly connected insurers tend to have larger exposures, as evidenced by scatter plot in Figure 7. The statistical dependence is supported formally by the Kendall tau test at a 1% significance level. For the treaty and facultative data we have a tau statistic (and the corresponding p-value) of 0.57 (< 0.01) for total degree against strength, 0.61 (< 0.01) for in-degree against in-strength, and 0.94 (< 0.01) for out-degree against out-strength.

Our findings of a heterogeneous degree distribution and negative degree correlations reinforce our view that the network has a core-periphery structure. Highly connected reinsurers at the core of the graph play the role of central hubs that mediate risk transfer between lesser connected insurers in the market. Our empirical analysis finds a small core (less than 5% of all insurers in the network), consisting of hub reinsurers, interacting with a large part of the network and mediating a significant portion of risk transfer activity in the market.[20] For the facultative and treaty data set, the top 5% of connected insurers hold 60% of links in the network, amounting to 81% of the total value of reinsurance in the market; while for the recoverables data set, the hub insurers hold 98% of links, representing 96% of the

[20]Due to data confidentiality, we omit detailed presentation of results on the core reinsurers in the network.

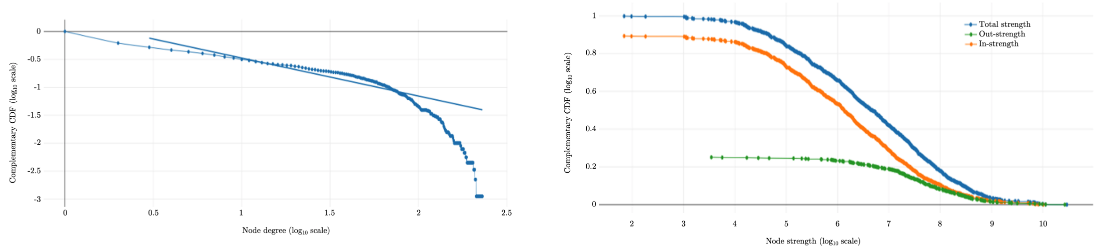

(a) Total degree distribution with a power law fit on a $\log_{10}$-$\log_{10}$ scale.

(b) Node strength distribution plot on a $\log_{10}$-$\log_{10}$ scale.

Figure 6: Degree (left) and strength (right) distribution for network of facultative and treaty reinsurance contracts.

total reinsurance amount.

The distribution of connectivity has important implications on systemic vulnerability (Allen and Gale, 2000; Battiston et al., 2012b). In particular, Caccioli et al. (2012) show that conditional on observing a contagion event[21] the size of the loss cascade increases with the average degree: higher connectivity implies that losses can be transmitted faster to a larger number of market participants. On the other hand, the probability of observing a contagion event is a concave function of the average degree: while highly connected networks allow for higher diversification of exposures and thus increase the robustness of the system to reinsurer default, poorly connected networks do not have a sufficient number of links to trigger a cascade. Consequently, networks with a medium level of connectivity tend to have the highest probability of observing a default cascade, as there are enough links for the losses to spread but insurers do not benefit from much diversification of exposures. Our stress-simulation analysis in Section 4 provides further insights into the network's resilience.

3.2 Small-World Graph

Links in the UK insurance market are sparse, and with network density of only around 1%, the vast majority of insurers are not connected to one another (see Tables 3 and 4). The majority of insurers in the periphery of the network have a very small number of connections, and only a small portion of reinsurers in the core are connected to a large number of market participants. Interestingly, with the exception of a few nodes, all insurers belong to the same connected component. In other words, almost all market participants belong to one risk sharing network.

[21] A network contagion defines an event in which a default of a single node in the system leads to a consecutive failure of other institutions in the network.

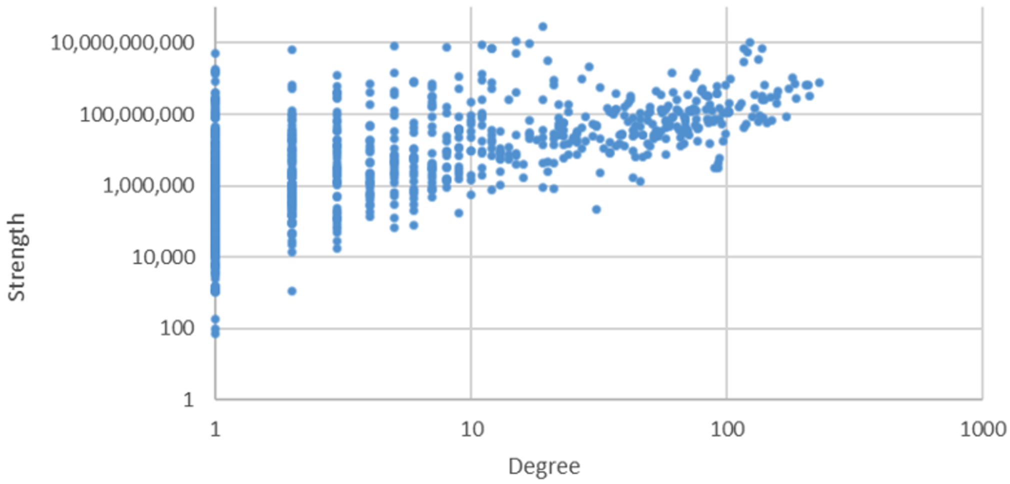

Figure 7: A positive relation between node degree and strength distribution for the facultative and treaty network.

Despite the fact that most insurers belong to a single risk sharing network, we do observe clustering of insurers forming around reinsurers in the core and find high local clustering as measured by the clustering coefficient. In other words, there is a high probability that two insurers with a reinsurance contract linking them also have another common counterparty. This relates to the counterparty risk externality, as described by Acharya and Bisin (2014). That is, the action of a reinsurer underwriting contracts with other insurers increases its default probability, and hence leads to an increased default risk faced by its counterparties. The externality arises in an opaque market, as insurers do not have specific information on the risk taken by its reinsurer.

As a result of the high local clustering, we observe local community structures in the network. These are small sub-networks characterised by a relatively high density of links. In those communities, the likelihood of a connection between its member insurers significantly exceeds the average probability that any two insurers in the network are connected.

Hub reinsurers serve as the common connections for communities, which results in a short path between any two insurers in the system. Consequently, the network diameter, defined as the shortest path length between the two most distant insurers, tends to be low for scale-free networks with a core-periphery structure.

The presence of these hubs, commonly observed in financial networks, distinguishes scale-free networks from random networks (Barabási and Pósfai, 2016; Barabási and Albert, 1999). As suggested by Barabási and Albert (1999), real networks often exhibit a preferential connectivity: the emergence of scale-free topology can be described by the preferential attachment process in which a new node in the system has a higher probability of forming a link with already well connected nodes than with nodes that have a small number of connections. Consequently, insurers are more likely to transfer part of their risk to a large reinsurer in the core of the network than a reinsurer found in the periphery.

We observe that insurers with a smaller number of connections tend to have counterparties with densely connected neighbours (that is, they form a community), while insurers that act as hubs connect insurers that tend not to link with each other. This result is corroborated by Figure 8 which displays the relationship between the local clustering coefficient and total degree for the treaty and facultative network. Analogous results were found for the recoverables data set.

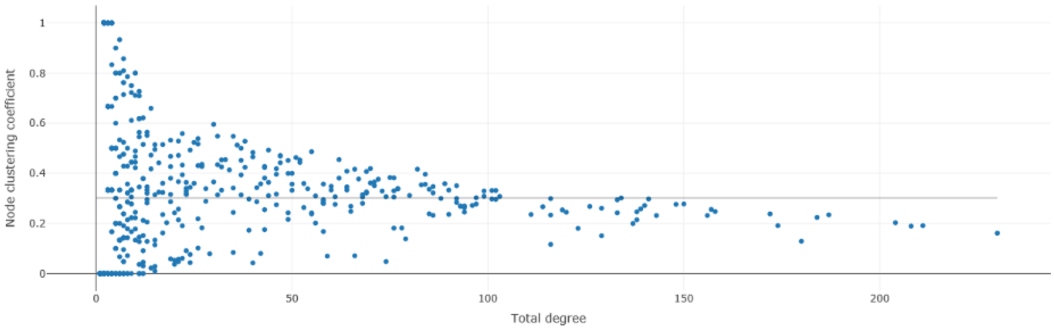

Figure 8: The relationship between the local clustering coefficient and node's degree for the treaty and facultative network. The clustering coefficient remains bounded away from zero, which is characteristic of small-world graphs.

This hierarchical relationship between insurers and reinsurers is captured by the assortativity coefficient: a negative value shows the tendency of highly connected insurers (hubs) to enter in a reinsurance contract with insurers that have a small degree (see Tables 3 and 4). The negative relationship between an insurer's degree and the clustering coefficient is a feature of the core-periphery structure.

Although most of the insurers in the network are not directly connected with one another, they can be indirectly connected by a small number of links through a small set of hubs. Despite the scarcity of links, the average path length between insurers remains small: a pair of insurers is linked by a path with a length orders of magnitude lower that the network size. Shorter paths create the conditions for losses to spread more quickly.

Link sparsity, high local clustering and short average path lengths often distinguish real networks from simple theoretical models (Cont and Tanimura, 2008): an empirical observation that inspired the development of a new class of random networks, termed *small-world* graphs (Watts and Strogatz, 1998; Watts, 1999;

Newman, 2003).

Our analysis supports the view that the UK reinsurance market forms a small-world graph. A notable feature of this class of graph is that small-world networks are more robust to random shocks in the system than other network architectures: since hubs mediate most connections, the default of a small reinsurer should have minimal effect on systemic stability. On the other hand, default of a large central reinsurer can have a significant adverse effect on the system. Consequently, we expect the commonly described 'robust-yet-fragile' property of financial networks to hold for the UK insurance market as well (Gai and Kapadia, 2010; Gai, 2013; Caccioli et al., 2012). This reinforces our earlier findings on the core-periphery nature of connectivity, and emphasises the importance of stress simulation, conducted in Section 4, for assessing the level and source of vulnerability in the UK reinsurance sector.

We present a further in-depth network analysis of the UK insurance market in Appendix D.

4 Simulation of Stress Scenarios

The analysis of UK reinsurance market structure reveals potential sources of systemic vulnerability. In this section, we use a simulation-based approach to assess the risk of counterparty credit contagion. In particular, to test the stability of the reinsurance recoverables network, we examine the effect of a shock to insurers' and reinsurers' financial investments, excluding unit-linked investments,[22] as well as the default of individual large, highly interconnected reinsurer groups within the UK network.[23]

Balance sheet data In our simulation-based approach, we consider the network of reinsurance recoverables consisting of 4,378 individual insurers. To assess the impact of bilateral counterparty exposures on systemic risk, we consider a simple representation of a balance sheet of each market participant.

For the UK insurers, we build their balance sheets using the readily available data from Solvency II regulatory returns. For the remaining non-UK insurers, the balance sheet information is obtained manually from SNL, Capital IQ, and US Statutory Insurance Financial information (all from Standard & Poor's). We then

[22]Unit-linked investments are made on behalf of the policyholder, who bears the investment risk (similarly as in the case of mutual funds).

[23]Further layers could be added to the stress simulation exercise, based on retrocession, where reinsurers cede exposure to other reinsurers. That said, our investigations reveal that the data does not contain standardised, precise risk descriptions to enable this analysis, and hence we focus only on the counterparty default risk instead.

construct an approximately equivalent balance sheet to that of Solvency II, for example by removing intangible assets such as goodwill and deferred acquisition costs. Due to time limitations, balance sheet information could be obtained only for the largest non-UK reinsurers, as characterised by the size of their exposures in the UK insurance market.

As a consequence of this data limitation, we possess detailed balance sheet information only for approximately 500 insurers in the network. From the perspective of stress testing, it is however crucial that other insurers are retained in the network. Their removal could significantly alter the graph topology, and thus the way in which default contagion can spread through the system. Therefore, in order to capture the network effects adequately, we retain the remaining insurers in our simulation approach by endowing them with artificial balance sheet data. In particular, we build total assets and liabilities around their observed reinsurance assets and liabilities, assuming a proportional relationship between these items. Our approach is motivated by common reinsurance market risk practices. We calibrate the proportion using the aforementioned commercial data as well as data taken from the 2016 EIOPA Insurance Stress Test (EIOPA, 2016). Notably, we apply a capital buffer of assets over liabilities of 10% for the artificial balance sheet data, in line with the average from the 2016 EIOPA Insurance Stress Test (EIOPA, 2016).

Stress test methodology For consistency with the EIOPA's exercise, we calibrate the stress to a 4% loss in the value of non-unit-linked investments; we also include much more severe stresses of 10% and 15% of total investments. However, it should be noted that EIOPA's stress test changes the value of assets and liabilities to result in a change in equity. Given that we do not have data on the underlying cash flows and hedging behaviour, we are not able to fully replicate the mechanics of EIOPA's stress scenario on the liability side. Instead, we observe the net impact of the stress scenario on equity and replicate it with a suitable shock on the asset side of the balance sheet.

In our framework, any recoverable amounts that could not be paid were allocated amongst creditors in proportion to their claim. In particular, we use an extension of the Eisenberg and Noe (2001) model, where the obligation of all firms within the system are determined simultaneously and consistently with the priority of debt claims and the limited liability of equity. This extension, introduced by Rogers and Veraart (2013), includes bankruptcy costs through the Greatest Clearing Vector Algorithm (GA) that allows insurers to fail in succession until only solvent firms remain. These bankruptcy costs include, for instance, the costs of insolvency specialists or inefficiencies in the ability to realise asset values. In line with the findings of Oxera (2007), we consider in our stress simulation scenarios bankruptcy costs of 10% and 20%. In contrast to other financial sectors, these relatively small bankruptcy costs are motivated by the fact that insurance claims are paid over long time periods (typically years). This allows for a more orderly recovery of claims by counterparties of the defaulted reinsurer. Moreover, we find our results on loss contagion in general to be robust to varying bankruptcy costs.

Summary of results Applying an adverse scenario in line with the 2016 EIOPA Insurance Stress Test, defined by a loss to total investments of 4% and bankruptcy cost of 10%, we find that whilst 25% of the capital in the system was lost, out of 4,378 firms we observe only 14 insurers initially defaulting due to the adverse shock and then a single default due to network effects. Turning to a much more severe stress scenario with a 15% loss to total investments and bankruptcy cost of 20%, we observe 136 defaults in total, where only 18 (13%) of these stem from the contagion effects. This relatively small number of defaults is despite a severe loss of 93% of capital in the system.

Table 5 presents the results based on a range of macroeconomic shocks to total investments (excluding assets backing unit-linked policies) and the assumed bankruptcy cost of 10%. Similar results are observed under a 20% bankruptcy cost assumption, as outlined in Table 6.

Shock on total investments	**4%**	**10%**	**15%**
Solvent nodes	4356	4307	4248
Shock defaults	14	62	118
Contagion defaults	1	2	5
Total losses	347 (24.7%)	864 (61.7%)	1297 (92.7%)
Shock losses	344 (99.0%)	859 (99.4%)	1290 (99.5%)
Contagion losses	3.5 (1.0%)	5.0 (0.6%)	7.1 (0.5%)

Table 5: Stress simulation results with 90% recovery rate assumed: shock to total investments (excluding assets backing unit-linked policies). Values given in billions of GBP. Results for the UK insurance market consisting of 4,378 nodes and a total initial capital of 1,400 billion GBP. Shock losses result directly from the initial shock of the scenario, whilst contagion losses are a result of network effects.

Our results show a limited impact of default contagion on the stability of the UK reinsurance market. In particular, risk of default remains small even in the presence of large macroeconomic shocks that erase most of the capital held by the market participants. In contrast to other financial networks, our results suggest that the connectivity induced by reinsurance contracts does not introduce fragility to the system. Secondary losses from counterparty defaults constitute

Shock on total investments	**4%**	**10%**	**15%**
Solvent nodes	4356	4307	4235
Shock defaults	14	62	118
Contagion defaults	1	2	18
Total losses	348 (24.9%)	867 (61.9%)	1302 (93.0%)
Shock losses	344 (98.8%)	859 (99.1%)	1290 (99.0%)
Contagion losses	4.3 (1.2%)	8.1 (0.9%)	12.4 (1.0%)

Table 6: Stress simulation results with 80% recovery rate assumed: shock to total investments (excluding assets backing unit-linked policies). Values given in billions of GBP. Results for the UK insurance market consisting of 4,378 nodes and a total initial capital of 1,400 billion GBP. Shock losses result directly from the initial shock of the scenario, whilst contagion losses are a result of network effects.

only approximately 1% of total capital lost under the considered range of stress scenarios.

Whilst these results provide a view on how the topology of the reinsurance network affects its stability, they should not be regarded as definitive. In particular, we have had to make important assumptions about entities in the network for whom we lack data. Consequently, our exercise may fail to capture certain idiosyncrasies amongst insurers, for example ones that coincide with a highly connected node in the network. Since hub reinsurers found in the core of the network tend to have a higher level of capitalisation than an average insurer, our market-wide stress scenario does not induce their immediate default. Therefore, even with a reduced capital buffer under a stress scenario, strong risk diversification allows these hub reinsurers to maintain systemic stability by absorbing losses from contagious defaults and thus shielding the firms in the network periphery.

In light of this, we are interested in quantifying the impact of the failure of a given hub node (or a set of nodes related to a single insurance group) on the stability of the system. In our approach, we consider an idiosyncratic shock to highly connected reinsurers[24] – that belong to a specific large insurance group – resulting in their immediate default.[25] We consider the worst-case scenario in which the shocked reinsurers see the value of their total assets reduced to zero; that is, the immediate counterparties have an assumed 0% recovery rate on their reinsurance exposure. For the remaining contagious defaults, as previously, we

[24]We use the in-strength centrality measure (that is, the number of contracts weighted by size) to choose the potential candidates.

[25]Due to data confidentiality, we omit presenting details on these core reinsurer groups.

assume bankruptcy costs of 20%.

Table 7 presents the resulting systemic impact following the default of the most prominent entities within a single insurance group. Overall, we find the UK reinsurance market to be robust to default of the largest, most connected insurers. In general, losses are contained within a small part of the network formed by the defaulted entities and their direct counterparties. In our simulations, we observe at most four additional contagion-induced defaults, while typically only a single default by contagion is seen. Interestingly, this single default is attributed to the same insurer in the network, which rather highlights its inadequate capitalisation to withstand a major shock than the systemic impact of its reinsurers.

Defaulted group	**I**	**II**	**III**	**IV**	**V**
Solvent nodes	4376	4377	4377	4372	4376
Shock defaults	2	1	1	3	2
Contagion defaults	1	1	1	4	1
Total losses	10.1 (0.7%)	10.4 (0.7%)	4.3 (0.3%)	21.6 (1.5%)	5.7 (0.4%)
Shock losses	5.9 (58%)	6.2 (60%)	0.07 (2%)	16.2 (75%)	1.5 (26%)
Contagion losses	4.2 (42%)	4.2 (40%)	4.2 (98%)	5.4 (25%)	4.2 (74%)

Table 7: Stress simulation results with 80% recovery rate assumed: shock from a default of individual entities within a single large insurance group. Values given in billions of GBP. Results for the UK insurance market consisting of 4,378 nodes and a total initial capital of 1,400 billion GBP. Shock losses result directly from the initial shock of the scenario, whilst contagion losses are a result of network effects.

Our results are in line with the previous conclusions of IAIS (2012) and Geneva Association (2010), which find reinsurance unlikely to contribute or amplify systemic risk due to the relatively small size of premiums ceded and retroceded as compared to the primary insurance market. This view is corroborated by the results from our stress test simulations, both in the case of a macroeconomic shock affecting insurers' portfolios and an idiosyncratic shock leading to a default of prominent individual entities within a single insurance group. We find that the robustness of the system can be attributed to a combination of two factors: the structure and the size of counterparty default risk exposures. In particular, the observed network characteristics, such as the scale-free degree distribution and the core-periphery structure, help to attenuate and disperse disturbances within the system. At the same time, since counterparty exposures due to reinsurance contracts form a relatively small fraction of assets and liabilities on the balance sheet, even large macroeconomic shocks are insufficient to reach the tipping point beyond which interconnectedness intermediates spread of the contagion and thus aids shock amplification. This is in a stark contrast to banking networks, where interbank exposures form a significant part of their business, which renders them robust-yet-fragile to the risk of contagion (Gai and Kapadia, 2010; Acemoglu et al., 2015; Caccioli et al., 2015).

5 Conclusions

Our research has sought to review the evidence for the commonly held view that, in and of itself, reinsurance does not pose contagion risk. The hierarchical relationships and limited ceding of insurance risk are believed to provide natural brakes for the spread of counterparty losses experienced by reinsurers.

Using 2016 UK data on reinsurance contracts and amounts recoverable, our findings do not contradict this view. We find the reinsurance network consists of densely connected insurers at the core, playing the role of 'hubs' for the UK insurance market, while there are other sparsely connected reinsurers in the periphery. Reinsurers in the core tend not to connect with one another, that is, relationships tend to be strongly hierarchical. The UK reinsurance network exhibits topological properties similar to theoretical scale-free and small-world networks. Relatively few connections are needed to reach each insurer in the network, which gives rise to the 'robust-yet-fragile' property.

We find that life and non-life networks both exhibit the core-periphery structure. In comparison, life reinsurance contracts tend to lead to larger exposures than non-life contracts, and the non-life reinsurance network is denser than the life network, suggesting greater risk-sharing by non-life insurers. Risk-sharing by life insurers tends to be conducted via central reinsurers to a greater extent than for non-life insurers.

Although there is a small core of hubs that are highly connected in the UK insurance market, these hubs tend to connect with each other in a lesser degree than is the case for other financial networks. Moreover, community structures do not appear to form around lines of business, which indicates a strong diversification of insurance risk.

Finally, we use stress-simulation analysis to gain insight into whether the 'robust-yet-fragile' property of the reinsurance network could lead to instability. Using the EIOPA (2016) stress test as reference, we calibrate a shock to the insurers' financial investments excluding assets held for unit-linked policyholders, and find defaults from contagion to be relatively low. Similarly, to review the risk posed by the failure of hubs, we look at the effects of the individual default of highly connected reinsurers within the UK insurance market. Again, we find little contagion risk from these shocks. Our results, however, are based on severe balance sheet assumptions due to data limitations. In particular, our exercise ignores any connections between entities outside of the UK reinsurance market that are not reported under Solvency II. Moreover, the results are based on a large number of non-UK insurers for which we do not have readily available balance sheet data, and thus for which we impose artificial assumptions about their assets and liabilities.

Future work Solvency II creates a rich set of data for analysis, the potential of which we have only partly been able to explore. While we have data on reinsurance contracts held by UK insurers there are further data that would be useful in building a global perspective:

- Group-level data that report transactions by non-UK insurers, including intra-group transactions.
- Other group-level data from EU insurers that are shared within supervisory colleges.
- Data on the insurers' use of special purpose vehicles (SPVs) and insurance linked securities.
- Publicly available data on reinsurance contracts ceded by US insurers.

Besides expanding the global scope of the reinsurance network analysis, we could also introduce a temporal dimension by incorporating reinsurance data sets for a wider range of reporting dates. A longer time series of data would help identify whether contractual relationships are stable over time. For example, for reinsurance recoverables in particular there could be some interesting changes between year-end 2016 and year-end 2017 given the numerous catastrophic events that occurred during 2017.

Network effects (and potential retrocession spirals) could be present due to the presence of feedback loops in the system. Further work could take place to standardise risk descriptions in Solvency II reporting to allow a finer identification of risks. A future analysis could also prioritise the identification of more important network cycles – for example, a feedback loop of losses is only possible if links in the cycle correspond to sufficiently large exposures with respect to the insurer's reserve levels. Further work could consider a sub-network that only contains large exposures, where the expected size of losses exceeds a threshold for losses as a proportion of initial capital.

We believe network analysis can be a useful addition to insurance stress tests. For example, it allows supervisors to explore second round effects from contagion following initial shocks. In addition, in order to extend the stress-simulation analysis presented in this paper, it would be useful to:

- Include more data on non-UK insurers, including more detailed and heterogeneous balance sheets.

- Review systemic risk due to an insurance specific shock, for example a major catastrophe. This would create an increase in the value of the reinsurers' liabilities, potentially giving rise to contagion for sufficiently large shocks. It would be particularly interesting to see whether non-proportional, excess of loss contracts can propagate losses as layers of cover are exhausted by losses.

References

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

Viral Acharya and Alberto Bisin. Counterparty risk externality: Centralized versus over-the-counter markets. *Journal of Economic Theory*, 149(C):153–182, 2014.

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

Hamed Amini, Rama Cont, and Andreea Minca. Stress Testing the Resilience of Financial Networks. *International Journal of Theoretical and Applied Finance (IJTAF)*, 15(1):1–20, 2012.

Andrew Bain. Insurance spirals and the Lloyd's market. *The Geneva Papers on risk and insurance*, 24(2):228–242, 1999.

Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.

Albert-László Barabási and Márton Pósfai. *Network science*. Cambridge University Press, Cambridge, 2016.

Paolo Barucca and Fabrizio Lillo. Disentangling bipartite and core-periphery structure in financial networks. *Chaos, Solitons & Fractals*, 88:244–253, 2016.

Stefano Battiston, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, and Joseph Stiglitz. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121–1141, 2012a.

Stefano Battiston, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli. DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk. *Scientific Reports*, 2(541), 2012b.

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

Nicole Bäuerle and Alexander Glauner. Optimal risk allocation in reinsurance networks. *Insurance: Mathematics and Economics*, 82:37–47, 2018.

Fabio Caccioli, Thomas A. Catanach, and J. Doyne Farmer. Heterogeneity, correlations and financial contagion. *Advances in Complex Systems*, 15(2):1250058, 2012.

Fabio Caccioli, J. Doyne Farmer, Nick Foti, and Daniel Rockmore. Overlapping portfolios, contagion, and financial stability. *Journal of Economic Dynamics and Control*, 51:50–63, 2015.

Fabio Caccioli, Paolo Barucca, and Teruyoshi Kobayashi. Network models of financial systemic risk: a review. *Journal of Computational Social Science*, 1: 81–114, 2018.

Hua Chen, J. David Cummins, Tao Sun, and Mary A. Weiss. The Reinsurance Network Among U.S. Property–Casualty Insurers: Microstructure, Insolvency Risk, and Contagion. *Journal of Risk and Insurance*, 87(2):253–284, 2020.

Aaron Clauset, Cosma R. Shalizi, and Mark E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.

Rama Cont and Emily Tanimura. Small-world graphs: Characterization and alternative constructions. *Advances in Applied Probability*, 40(4):939–965, December 2008.

Rama Cont, Amal Moussa, and Edson B Santos. Network structure and systemic risk in banking systems. In Jean-Pierre Fouque & Joseph A. Langsam, editor, *Handbook of Systemic Risk*, pages 327–368. Cambridge University Press, July 2013.

Rama Cont, Artur Kotlicki, and Laura Valderrama. Liquidity at risk: Joint stress testing of solvency and liquidity. *Journal of Banking & Finance*, 118:105871, 2020.

Ben Craig and Goetz von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23:322–347, 2014.

Matt Davison, Darrell Leadbetter, Bin Lu, and Jane Voll. Are Counterparty Arrangements in Reinsurance a Threat to Financial Stability? Staff Working Papers 39, Bank of Canada, 2016.

Hans Degryse and Gregory Nguyen. Interbank Exposures: An Empirical Examination of Systemic Risk in the Belgian Banking System. Working Paper No. 43, National Bank of Belgium, 2004.

Reinhard Diestel. *Graph Theory.* Springer-Verlag Berlin Heidelberg, Germany, 2017.

EIOPA. 2016 EIOPA Insurance Stress Test Report. Report EIOPA 16/302, European Insurance and Occupational Pensions Authority, 2016.

Larry Eisenberg and Thomas Noe. Systemic Risk in Financial Systems. *Management Science*, 47(2):236–249, 2001.

Matthew Elliot, Benjamin Golub, and Matthew O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–3153, 2014.

Andrew Ellul, Chotibhak Jotikasthira, Christian T. Lundblad, and Yihui Wang. Is Historical Cost Accounting a Panacea? Market Stress, Incentive Distortions, and Gains Trading. *Journal of Finance*, 70(6):2489–2538, 2015.

Andrea French, Mathieu Vital, and Dean Minot. Insurance and Financial Stability. *Bank of England Quarterly Bulletin*, Q3 2015.

Daniel Fricke and Thomas Lux. Core–Periphery Structure in the Overnight Money Market: Evidence from the e-MID Trading Platform. *Computational Economics*, 45(3):359–395, 2015.

Prasanna Gai. The Robust-Yet-Fragile Nature of Financial Systems. In Prasanna Gai, editor, *Systemic Risk: The Dynamics of Modern Financial Systems*, pages 9–28. Oxford University Press, Oxford, 2013.

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

Geneva Association. Systemic Risk in Insurance: An Analysis of Insurance and Financial Stability. Special Report of The Geneva Association Systemic Risk Working Group, The Geneva Association, March 2010.

IAIS. Insurance and Financial Stability. Technical report, International Association of Insurance Supervisors, 2011.

IAIS. Reinsurance and Financial Stability. Technical report, International Association of Insurance Supervisors, 2012.

Carmela Iazzetta and Michele Manna. The topology of the interbank market: developments in Italy since 1990. Temi di Discussione (Working Paper) No. 711, Bank of Italy, Economic Research and International Relations Area, May 2009.

International Monetary Fund. Chapter 3. The Insurance Sector – Trends and Systemic Risk Implications. In *Global Financial Stability Report: Potent Policies for a Successful Normalization*. International Monetary Fund, April 2016.

Matthew O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, 2008.

Masayasu Kanno. The network structure and systemic risk in the global non-life insurance market. *Insurance: Mathematics and Economics*, 67(C):38–53, 2016.

Mikko Kivelä, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. Multilayer networks. *Journal of Complex Networks*, 2(3): 203–271, September 2014.

Ariah Klages-Mundt and Andreea Minca. Cascading Losses in Reinsurance Networks. *Management Science*, 66(9):4246–4268, September 2020.

Iman Lelyveld, Franka Liedorp, and Manuel Kampman. An empirical assessment of reinsurance risk. *Journal of Financial Stability*, 7(4):191–203, 2011.

Sheheryar Malik and TengTeng Xu. Interconnectedness of Global Systemically-Important Banks and Insurers. Working Paper No. 17/210, International Monetary Fund, September 2017.

Robert McDonald and Anna Paulson. AIG in Hindsight. *Journal of Economic Perspectives*, 29(2):81–106, 2015.

Munich RE. Reinsurance: A Basic Guide to Facultative and Treaty Reinsurance, 2010.

Mark Newman. The Structure and Function of Complex Networks. *SIAM Review*, 45(2):167–256, January 2003.

Mark Newman. *Networks: An Introduction*. Oxford University Press, Oxford, United Kingdom, 2010.

Erlend Nier, Jing Yang, Tanju Yorulmazer, and Amadeo Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033–2060, 2007.

Oxera. Insurance guarantee schemes in the EU: Comparative analysis of existing schemes, analysis of problems and evaluation of options, 2007. Final report prepared for European Commission (DG Internal Market and Services).

Sojung Carol Park and Xiaoying Xie. Reinsurance and systemic risk: The impact of reinsurer downgrading on property-casualty insurers. *Journal of Risk and Insurance*, 81(3):587–622, 2014.

L. C. G. Rogers and Luitgard A. M. Veraart. Failure and Rescue in an Interbank Network. *Management Science*, 59(4):882–898, April 2013.

Tarik Roukny, Hugues Bersini, Hugues Pirotte, Guido Caldarelli, and Stefano Battiston. Default Cascades in Complex Networks: Topology and Systemic Risk. *Scientific Reports*, 3(2759), 2013.

Vicky Saporta. Macroprudential Policy for Insurers. Speech at the association of british insurers annual conference, Bank of England, 2016.

Christian Upper. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111–125, 2011.

Duncan J. Watts. *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton University Press, Princeton, 1999.

Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

A Treaty and Facultative Statistics

Table 8 presents a quantitative summary of premiums ceded by contract type and by line of business in the treaty and facultative data set.

Type of contracts	**Frequency** (percent of total)	**Value in £m** (percent of total)
Life, external	1905 (4.5%)	62181 (57.7%)
Life, internal	81 (0.2%)	7557 (7%)
Non-life, external	37048 (88.5%)	25286 (23.5%)
Non-life, internal	1042 (2.5%)	11833 (11%)
Life, proportional	1721 (4.1%)	69160 (64.2%)
Life, non-proportional	258 (0.6%)	115 (0.1%)
Non-life, proportional	6177 (14.7%)	27621 (25.6%)
Non-life, non-proportional	33700 (80.5%)	10306 (9.6%)
Fire and other damage to property	15941 (38.1%)	12717 (11.8%)
Marine, Aviation, Transport	12253 (29.3%)	3288 (3.1%)
General liability	5556 (13.3%)	4973 (4.6%)
Other non-life	1689 (4%)	1566 (1.5%)
Other life	1428 (3.4%)	52003 (48.3%)
Motor insurance	1315 (3.1%)	2642 (2.5%)
Credit and suretyship insurance	1311 (3.1%)	685 (0.6%)
Health reinsurance	1162 (2.8%)	4316 (4%)
Unit-linked or index-linked	318 (0.8%)	16681 (15.5%)
Multiline	317 (0.8%)	7326 (6.8%)
Medical expense insurance	294 (0.7%)	335 (0.3%)
Life reinsurance	201 (0.5%)	916 (0.9%)
With profit participation	56 (0.1%)	139 (0.1%)
Non-life annuities	42 (0.1%)	110 (0.1%)

Table 8: Summary of premiums ceded in the treaty and facultative data set for the 2016 year-end.

B Network Theory

Here we introduce basic network theory concepts.[26] In the context of reinsurance market, we define a network or a graph, denoted by $\mathcal{G} = (V, E)$, through a non-empty set of vertices (or nodes) $V = \{v_1, \ldots, v_N\}$ representing insurance market participants, and the set of edges $E = \{(v_i, v_j)\}$, with $v_i, v_j \in V$, representing bilateral exposures originating from reinsurance contracts. In particular, if (v_i, v_j) is in the set E, then we say that v_i is *adjacent* to v_j, and vertices v_i, v_j are said to be neighbours. As a result, every graph may be characterised by its *adjacency matrix*, defined as follows.

Definition (Adjacency matrix)**.** *The adjacency matrix of a graph $\mathcal{G}$ is the $N \times N$ matrix $A = A(\mathcal{G})$ whose entries a_{ij} are given by*

$$a_{ij} = \begin{cases} 1, & \textit{if } v_i \textit{ is adjacent to } v_j; \\ 0, & \textit{otherwise.} \end{cases} \tag{1}$$

To represent the direction of risk transfer in the insurance network, we represent bilateral exposures in a form of a directed graph. The graph is said to be *directed* if the set of edges E is defined by ordered pairs of nodes. By convention, the adjacency matrix of a directed network has element $a_{ij} = 1$ if there is an edge from j to i (and zero otherwise). In that case, the adjacency matrix does not need to be symmetric. Consequently, in a directed network each edge (v_i, v_j) has associated a source node (the initial vertex), $v_i \in V$, and a target node (the terminal vertex), $v_j \in V$, specifying the nature of the interaction. By convention, for risks ceded under the treaty and facultative contracts, we draw an edge from the insurer to a reinsurer to highlight the insurance risk transfer from the primary insurer to its counterparty. On the other hand, when discussing the bilateral exposures due to counterparty default risk from existing reinsurance recoverables, we consider an edge originating from an reinsurer to the primary insurer that is owed money in the form of reinsurance recoverables.

When studying the reinsurance network, we are also interested in quantifying the size of counterparty exposures. We achieve this by considering a weighted version of the adjacency matrix A in (1), denoted by $W = W(\mathcal{G})$, where we give the elements of the adjacency matrix values equal to the weights of the corresponding connections. In particular, for the recoverables network $\mathcal{G}^r$, we replace each non-zero entry a_{ij} of the adjacency matrix A with the respective size of the aggregate counterparty default exposure from node v_j to node v_i. Similarly, for the treaty and facultative network $\mathcal{G}^s$, non-zero entries a_{ij} are replaced with the aggregate value of risk ceded from v_i to node v_j. We remark that by construction matrices $W(\mathcal{G}^r)$ and $W(\mathcal{G}^s)$ are non-negative.

[26]See, for example, Diestel (2017); Newman (2010) and Jackson (2008).

Basic concepts in network theory Often an indirect connection between two nodes in the network is of interest. The concept of a (directed) path from node v_i to node v_j relates to an ordered sequence of nodes, starting from v_i and finishing at v_j, such that there exist a (directed) edge between each consecutive pair of nodes. Formally, a path is a non-empty graph $P = (V^*, E^*)$ of the form

$$V^* = \{v_0, v_1, \ldots, v_k\}, \quad E^* = \{(v_0, v_1), (v_1, v_2), \ldots, (v_{k-1}, v_k)\}, \tag{2}$$

where the v_i are all distinct. For simplicity, we may refer to a path by the natural sequence of its nodes, where in the above example we would write $P = v_0 v_1 \ldots v_k$. The vertices v_0 and v_k are said to be *linked* by P, and the number of edges of a path is its *length*, k. Two nodes may have more than one distinct path connecting them, and the shortest path is said to be the *geodesic* path. The length of the shortest path defines the distance between the two nodes. The *diameter* of a graph can then be defined as the length of the longest geodesic path between any pair of nodes in the network for which a path actually exists. Studying these properties proves useful in identifying how risk is transferred through the reinsurance market and how losses could spread (or be contained).

A related idea to geodesic paths in a network is the one of betweenness centrality. In network analysis, centrality measures aim to characterise important nodes in the network.[27] In particular, the betweenness centrality of a node v_l considers the number of shortest paths linking any two nodes in the network that pass through the given node v_l. Formally, let b_{ij}^l be 1 if node v_l is in the geodesic path from v_i to v_j, and 0 otherwise (including when such path does not exist). Then, define the betweenness centrality B_l of vertex v_l as

$$B_l = \sum_{i,j} b_{ij}^l. \tag{3}$$

Retrocession spirals, such as the London Market Excess Loss observed in the late 1980s, reveal how global interconnectedness can make the reinsurance market vulnerable to contagion. We can formalise the notion of a spiral, in which insurers cede risk only to accept the same risk through another contract, with a concept of a network *cycle*. Consider a path $P = v_0 v_1 \ldots v_{k-1}$ such that $k \geq 3$. Then, we define a cycle by the graph $C := P + v_{k-1} v_0$, and its length k is given by the number of edges (or vertices).

[27] As the notion of node importance is often ambiguous in a general context, there are many definitions of centrality and the associated measures. A notable example includes the DebtRank measure of systemic impact of a financial intuition, introduced in Battiston et al. (2012b). In our analysis, however we do not attempt to assess the systemic importance of insurers based on some of these measures. Instead, we use a simulation-based approach to assess the systemic vulnerabilities in different stress scenarios, including a default of the prominent insurers.

The number of paths of a given length k on a network can be easily computed using the adjacency matrix. Noticing that the product $a_{il}a_{lj}$ is 1 if a path of length 2 from v_j to v_i via v_l exists (and zero otherwise); the total number of paths of length 2 from v_j to v_i is given by $N_{ij}^{(2)} = \sum_{l=1}^{N} a_{il}a_{lj} = [A^2]_{ij}$. This concept can be generalised to an arbitrary path length k, $N_{ij}^{(k)} = [A^k]_{ij}$. Similarly, the total number of cycles of length k is given by $L^{(k)} = \sum_{i=1}^{N}[A^k]_{ii} = \mathrm{Tr}(A^k)$.

If there exists a path between each node in the graph, the said network is *connected*. Otherwise, it is a disconnected network, which can be partitioned into distinct connected components. A *connected component* of an undirected network is defined by a subgraph in which any two nodes are connected to each other by paths, and also is such that no further nodes can be added to the subgraph preserving the above property. In a directed graph, components can be classified into two types: a strongly connected component in which every node needs to be connected with other nodes by a directed path, and a weakly connected component in which existence of an undirected path between every node in the subgraph is sufficient. Component analysis provides an alternative measure of interconnectedness of the network; and in the case of the reinsurance network, strongly connected component may be used to identify the active risk sharing community subjected to a potential contagion risk (Chen et al., 2020).

An important property of a node is its degree. We define the *degree* (or valency) $d(v_i)$ of a vertex v_i to be the number $|E(v_i)|$ of edges at v_i. In a directed graph, we can extend the notion of a degree to an in-degree of a vertex v_i, denoted by $d^{in}(v_i)$, to be the number of edges with v_i as the terminating vertex. Similarly, the out-degree of a vertex v_i, denoted by $d^{out}(v_i)$, is the number of edges with v_i as the initial vertex. Recalling that the adjacency matrix of a directed network has element $a_{ij} = 1$ if there is an edge from j to i, in- and out-degrees can be written as

$$d^{in}(v_i) = \sum_{j=1}^{N} a_{ij}, \qquad d^{out}(v_j) = \sum_{i=1}^{N} a_{ij}. \tag{4}$$

By definition, we have that $d(v_i) = d^{in}(v_i) + d^{out}(v_i)$. In a reinsurance network, in-degree is the number of incoming links to the reinsurer, and thus corresponds to the total number of counterparties that are exposed to its default risk. Similarly, the out-degree computes the number of outgoing links from the insurer, which represents the total number of reinsurance contracts held by the insurer. The degree is the sum of an insurer's in- and out-degree components, and measures the connectivity of the insurer. Analogously, the *strength* of an insurer in the network incorporates the weight of each link when computing this measure. That is, we can write the respective in- and out-strength as

$$s^{in}(v_i) = \sum_{j=1}^{N} w_{ij}, \qquad s^{out}(v_j) = \sum_{i=1}^{N} w_{ij}, \tag{5}$$

where W is the weighted adjacency matrix with entries w_{ij}. Similarly, we define the strength of node v_i by $s(v_i) = s^{in}(v_i) + s^{out}(v_i)$.

Naturally, the number m of links present in the network can be related to the sum of the degree of all vertices, that is $m = \frac{1}{2}\sum_{i=1}^{N} d(v_i) = \frac{1}{2}\sum_{i,j} a_{ij}$. Consequently, we may be interested in expressing the global level of network connectivity through its density measure. In particular, we define the connectance or *density* ρ of a graph as a ratio of links that are present and the maximum possible number of edges, that is

$$\rho = \frac{2m}{N(N-1)}. \tag{6}$$

When considering the maximum possible number of edges we do not allow for multi-edges (that is, at most one edge between any pair of vertices) or self-edges (that is, no edges between a vertex and itself) in the network; in which case, we can have at most $\binom{N}{2} = \frac{1}{2}N(N-1)$ links. A network with no links has density 0, and a complete network has network density 1.

Another measure often of interest in financial networks is the *assortativity* of a network, which relates to the Pearson correlation coefficient of degree between pairs of linked nodes. Positive values indicate a relationship between nodes of similar degree (that is, highly connected nodes are connected to other highly connected nodes and vice versa), while negative values indicate relationships between nodes of different degree. Financial networks are often found to be disassortative (Caccioli et al., 2015; Cont et al., 2013), where market participants exhibit specific preferences when selecting their business counterparty.

Similarly, the nature of interconnectedness may be studied though the notion of a clique. Formally, in an undirected version of a graph, a *clique* is defined as a maximal subset of the vertices in which every member of the set is connected by an edge to every other. Here, a maximal subset means that no other node can be added to the subset while preserving the above property. A related concept is that of a local clustering coefficient, which measures the probability of neighbours of node v_i being themselves neighbours. In particular, we define a clustering coefficient c_i for a vertex v_i as the ratio of the number of pairs of neighbours of v_i that are connected and the number of pairs of neighbours of v_i (Newman, 2010). Following Watts and Strogatz (1998), we can define a clustering coefficient c for an entire network as the mean of the local clustering coefficients for each vertex, that is $c = \frac{1}{N}\sum_{i=1}^{N} c_i$.

The definition of a clique, requiring that every possible edge between its nodes is present, is a very stringent one, and often found to be impractical when studying real-world financial networks. Instead, it may be useful to study a *community* structure in these networks. In particular, we may be interested in dividing the graph into groups, clusters, or communities according to the pattern of edges. Commonly, the practice of community detection attempts to partition the vertices in such a way that there are many edges inside each group and only a few edges between groups. As the number of groups is not fixed beforehand, community detection includes a wide variety of different algorithms that can provide general insight on the intricate nature of interconnectedness within a network, which is not easily captured through the raw network topology.

C Overview of Original Data and Adjustments

Table 9 presents a detailed overview of original data and performed adjustments for the following networks and layers:

- (TF-A) Treaty and Facultative, all contracts.
- (TF-L) Treaty and Facultative, only life contracts.
- (TF-NL) Treaty and Facultative, only non-life contracts.
- (TF-G) Treaty and Facultative, group contracts.
- (R-A) Recoverables (net of collateral), all contracts.
- (R-L) Recoverables (net of collateral), only life contracts.
- (R-NL) Recoverables (net of collateral), only non-life contracts.
- (R-G) Recoverables (net of collateral), group contracts.

Network		**TR-A**	**TR-L**	**TR-NL**	**TR-G**	**R-A**	**R-L**	**R-NL**	**R-G**
Nodes	Total	893	162	825	385	4378	427	3682	2954
	Incoming links	799	125	743	345	1494	156	1354	920
	Outgoing links	225	49	202	117	3940	351	3246	2640
	Both	131	12	120	77	1056	80	918	606
Links (edges)	All contracts	41883	2003	39880	38190	25891	834	21346	25432
	Links	8564	305	8338	3545	22713	822	19125	11649
Adjust-ments	Initial contracts	50771	3913	46858	46281	26269	834	21374	26269
	Weight threshold above £0	1359	276	1083	1292	0	0	0	0
	Self-links	14	0	14	768	34	0	27	0
	Blanks	7515	1634	5881	6031	344	337	0	340
	Reverse negatives	0	0	0	0	5349	159	5076	5349
	Final number of contracts	41883	2003	39880	38190	25891	834	21346	25432

Table 9: Overview of data adjustments and quantitative statistics. Networks based on the treaty and facultative data set (TR) and the recoverables data set (R). Information shown for all contracts (A), life contracts (L), non-life contracts (NL), and contracts between distinct insurance groups (G).

D UK Insurance Market: a Heterogeneous System

Our network statistics (see Tables 3 and 4 for details) highlight that while core-periphery relationships exist in the reinsurance network there are nuances in how life insurers and non-life insurers share risks. The life reinsurance network is more hierarchical than the non-life network. In particular, it is characterised by a much lower diameter, a lower average path length and a lower betweenness centrality than the non-life network. This suggests greater connectivity, mediated by hubs, between life (re)insurers than is the case for general (re)insurers. Life insurers also tend to cede reinsurance to fewer counterparties than non-life insurers, as can be seen in Figure 9.

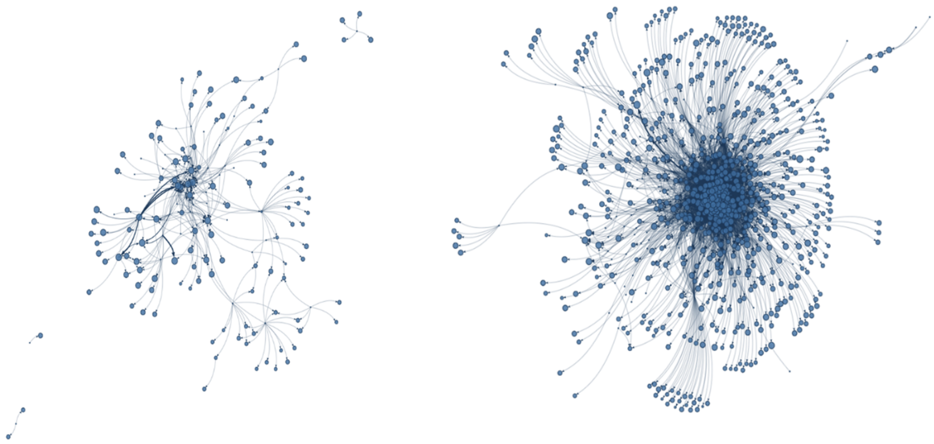

(a) Life reinsurance network. (b) Non-life reinsurance network.

Figure 9: Visualisation of sub-networks with life (left) and non-life (right) contracts using the treaty and facultative data set.

Figure 9 also reveals that the network of non-life reinsurance is denser and more connected than its life reinsurance counterpart, suggesting higher levels of insurance risk diversification. There are relatively few reinsurers of non-life insurance that have many connections, as corroborated by estimates of the degree distribution exponents in Tables 3 and 4, and they are proportionately fewer than comparable reinsurers of life insurance. However, the largest non-life reinsurers account for a higher proportion of contracts, by value, than the largest life insurers.

To examine further sources of heterogeneity within the overall network, we apply community analysis using clustering algorithms to identify cohesive groups known as communities. As explained earlier, these communities consist of densely connected insurers within the wider network, and the number of links between insurers in different communities remains relatively sparse (see Figure 10).

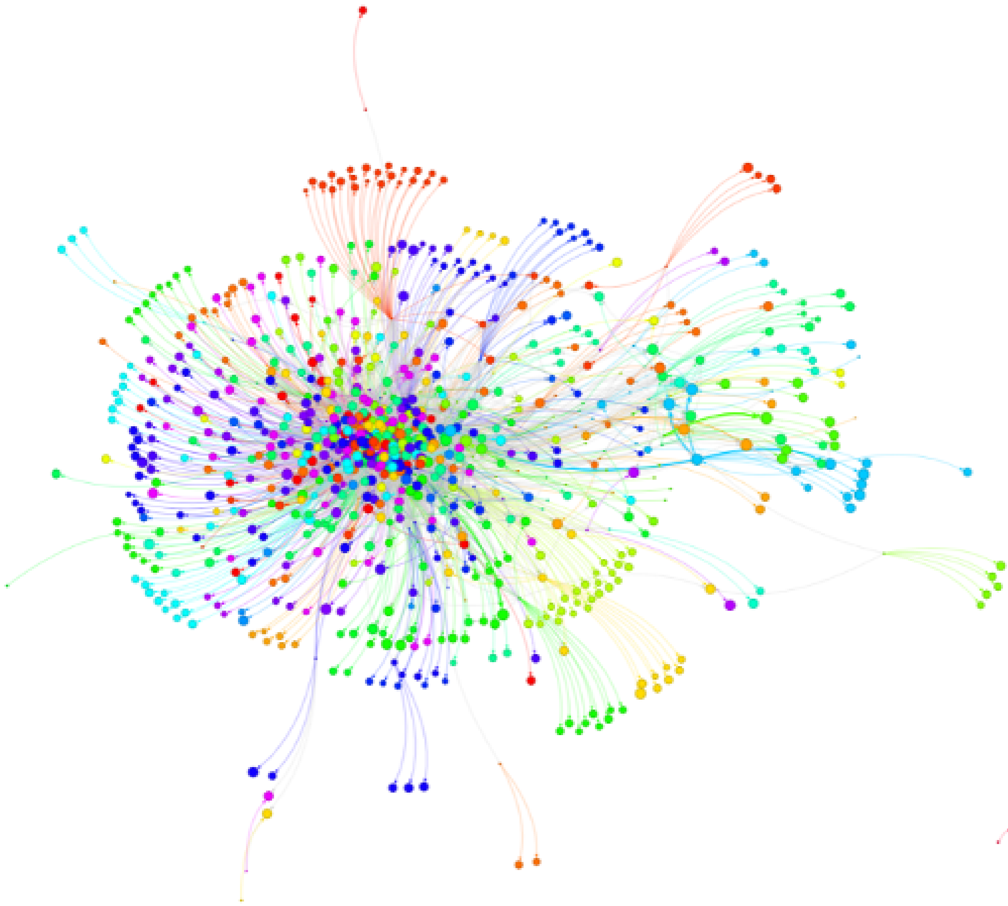

Figure 10: Community structure in the treaty and facultative network, detected using a Spinglass algorithm.

We find that clusters of densely connected reinsurers tend not to be organised around individual lines of business. Although we found some particular examples of clusters that were more specialised towards life insurance (see Figure 11 for the life and non-life network), most of the communities contained insurers with reinsurance contracts under multiple lines of business. This could simply be a reflection of diversification, particularly for non-life insurers, where insurers underwrite multiple lines of business and have multiple reinsurance contracts in place.

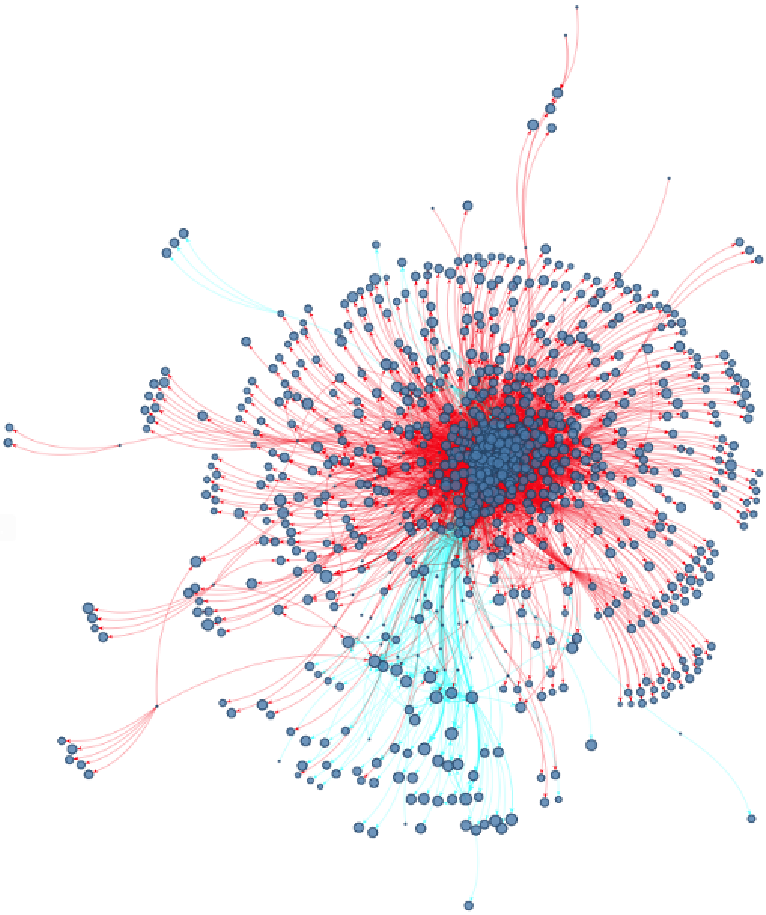

Figure 11: Treaty and facultative network visualisation with life edges in blue and non-life edges in red.

Retrocession spirals The London Market Excess of Loss spiral that affected the Lloyd's of London market participants in the late 1980s is a prime example of how global interconnectedness in the reinsurance market can cause contagion to spread (Bain, 1999). Despite the belief at the time that all parties are properly insured, the intricate structure of retrocession contracts resulted in an unusually large concentration of losses. In particular, retrocession spirals lead to a counter-intuitive non-linear behaviour of losses in the system, where a disproportionate amount of excess liability can be left with a single reinsurer (Klages-Mundt and Minca, 2020). In the presence of these network effects that are often invisible to the market participants, insurers face a high level of uncertainty about their risk, and thus may suffer from misspecification of internal risk models. In particular, as recently shown by Klages-Mundt and Minca (2020), the presence of retrocession spirals can have a detrimental effect on financial stability, making the reinsurance system vulnerable to systemic risk.

Network analysis can therefore help us to identify the dangerous retrocession cycles that can result in a severe underestimation of risk by the insurers. In particular, we use the treaty and facultative data to examine the prevalence of network cycles, and hence identify contracts that can lead to such retrocession spirals: that is, identify where risks are potentially retroceded back to the original insurer. Moreover, to focus on the nature of the risk, we layer the network by the different lines of business and identify network cycles of different sizes at each layer. Figure 12 shows a network plot for the entire UK insurance market as well as particular layers for the three most common lines of business.

Table 10 reports the number of identified network cycles. Note that due to computational complexity, only cycles of up to size of five are considered.

Line of Business	*Cycle Size*			
	2	*3*	*4*	*5*
Full network	335	5488	98831	1849710
Fire and other damage to property	104	888	8747	87286
Marine, aviation, transport	107	929	8915	87454
General liability	28	95	432	1768

Table 10: Summary of potential retrocession spirals (network cycles) for major lines of insurance business.

The data, however, imposes limitations on the network cycle detection analysis. Firstly, as we do not have world-wide data on reinsurance contracts, there could be important cycles left undetected due to lack of data for non-UK insurers and their ceded risk. Furthermore, although our results shed light on the potential for retrocession cycles in the UK insurance market, we do not have fine enough detail about the nature of the risk to infer that a loss event could be amplified because of the existence of a network cycle. The lines of business are broad descriptions of risk, and, for instance, damage to property could include residential housing in London and commercial property in Lagos. Consequently, more analysis is needed to detect whether the same risks exist, and given the available data granularity this has to be done on an individual basis. Our analysis only highlights the potential for retrocession cycles to exist, but we cannot be conclusive.

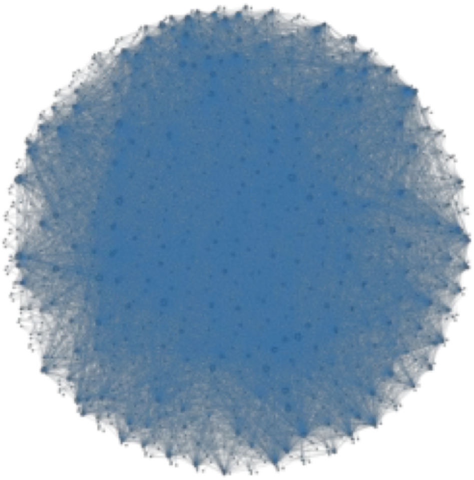

(a) Full network.

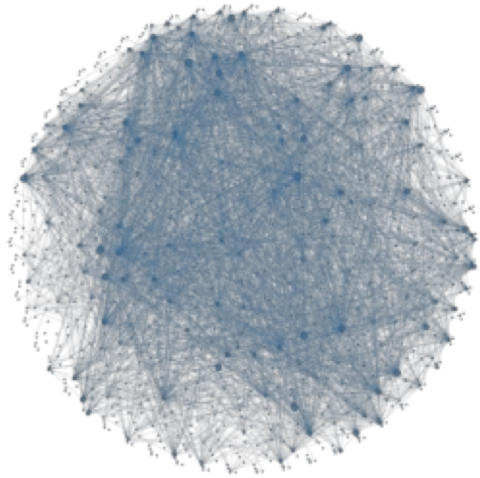

(b) Marine, aviation, transport.

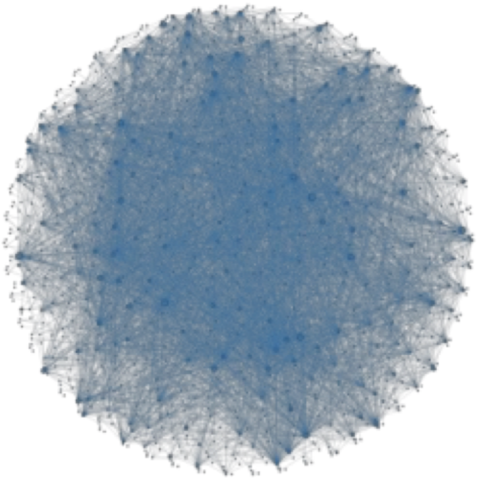

(c) Fire and other damage to property.

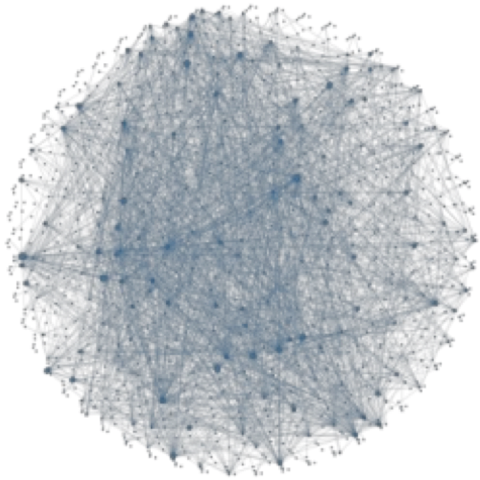

(d) General liability.

Figure 12: Visualisation of major lines of insurance business and potential retrocession spirals (network cycles). Node size proportional to PageRank centrality, and edge width to contract size. PageRank centrality here incorporates both the direction of links and their weight, and was originally developed for website ranking in search engines. The principle of the algorithm is that the vulnerability of a node increases with the number of connections to other vulnerable counterparties.

E Computer Software

The network analysis in this article was prepared using open source software for R, using the following packages:

- `data.table`: Matt Dowle and Arun Srinivasan (2019). data.table: Extension of 'data.frame'. R package version 1.12.8.; `https://CRAN.R-project.org/package=data.table`.
- `DT`: Yihui Xie, Joe Cheng and Xianying Tan (2018). DT: A Wrapper of the JavaScript Library 'DataTables'. R package version 0.5.; `https://CRAN.R-project.org/package=DT`.
- `igraph`: Csardi G, Nepusz T: The igraph software package for complex network research, InterJournal, Complex Systems 1695. 2006. `http://igraph.org`.
- `lpSolve`: Michel Berkelaar and others (2015). lpSolve: Interface to 'Lp_solve' v. 5.5 to Solve Linear/Integer Programs. R package version 5.6.13.; `https://CRAN.R-project.org/package=lpSolve`.
- `plotly`: Carson Sievert (2018) plotly for R.; `https://plotly-book.cpsievert.me`.
- `RODBC`: Brian Ripley and Michael Lapsley (2017). RODBC: ODBC Database Access. R package version 1.3-15.; `https://CRAN.R-project.org/package=RODBC`.
- `scales`: Hadley Wickham (2018). scales: Scale Functions for Visualization. R package version 1.0.0.; `https://CRAN.R-project.org/package=scales`.
- `shiny`: Winston Chang, Joe Cheng, JJ Allaire, Yihui Xie and Jonathan McPherson (2018). shiny: Web Application; Framework for R. R package version 1.2.0. `https://CRAN.R-project.org/package=shiny`.
- `shinyBS`: Eric Bailey (2015). shinyBS: Twitter Bootstrap Components for Shiny. R package version 0.61.; `https://CRAN.R-project.org/package=shinyBS`.
- `shinycssloaders`: Andras Sali (2017). shinycssloaders: Add CSS Loading Animations to 'shiny' Outputs. R package version 0.2.0.; `https://CRAN.R-project.org/package=shinycssloaders`.

- `shinyjs`: Dean Attali (2018). shinyjs: Easily Improve the User Experience of Your Shiny Apps in Seconds. R package; version 1.0. `https://CRAN.R-project.org/package=shinyjs`.

- `shinythemes`: Winston Chang (2018). shinythemes: Themes for Shiny. R package version 1.1.2.; `https://CRAN.R-project.org/package=shinythemes`.

- `shinyWidgets`: Victor Perrier, Fanny Meyer and David Granjon (2018). shinyWidgets: Custom Inputs Widgets for Shiny. R package version 0.4.4.; `https://CRAN.R-project.org/package=shinyWidgets`.

- `stringr`: Hadley Wickham (2019). stringr: Simple, Consistent Wrappers for Common String Operations. R package version 1.4.0.; `https://CRAN.R-project.org/package=stringr`.

- `V8`: Jeroen Ooms (2017). V8: Embedded JavaScript Engine for R. R package version 1.5.; `https://CRAN.R-project.org/package=V8`.

- `visNetwork`: Almende B.V., Benoit Thieurmel and Titouan Robert (2018). visNetwork: Network Visualization using 'vis.js' Library. R package version 2.0.5.; `https://CRAN.R-project.org/package=visNetwork`.

See `https://github.com/bank-of-england/NetworkApp` for further details.