BANK OF ENGLAND

Staff Working Paper No. 726

Multiplex network analysis of the UK OTC derivatives market

Marco Bardoscia, Ginestra Bianconi and Gerardo Ferrara

November 2018

This is an updated version of the Staff Working Paper originally published on 18 May 2018

BANK OF ENGLAND

Staff Working Paper No. 726
Multiplex network analysis of the UK OTC derivatives market

Marco Bardoscia,[1] Ginestra Bianconi[2] and Gerardo Ferrara[3]

Abstract

In this paper, we analyze the network of exposures constructed by using the UK trade repository data for three different categories of contracts: interest rate, credit, and foreign exchange derivatives. We study how liquidity shocks related to variation margins propagate across the network and translate into payment deficiencies across different derivative markets. A key finding of the paper is that, in extreme theoretical scenarios where liquidity buffers are small, a handful of institutions may experience significant spillover effects due to the directionality of their portfolios. Additionally, we show that a variant of a recently introduced centrality measure — Functional Multiplex PageRank — can be used as a proxy for the vulnerability of financial institutions, outperforming in this respect the commonly used eigenvector centrality.

Key words: Liquidity shock, multiplex networks, systemic risk, financial networks, central counterparty (CCP).

JEL classification: D85, G01, G17, L14.

(1) Bank of England. Email: marco.bardoscia@bankofengland.co.uk
(2) School of Mathematical Sciences, Queen Mary University of London. Email: g.bianconi@qmul.ac.uk
(3) Bank of England. Email: gerardo.ferrara@bankofengland.co.uk

Andrew Haldane, Mark Manning, Sean McGrath, Angus Moir, Pedro Gurrola-Perez, Radoslav Raykov, Andrea Serafino, Yedidiah Solowiejczyk, John Tanner, Michalis Vasios, Nicholas Vause, Paul Nahai-Williamson, Michael Yoganayagam, and an anonymous referee provided valuable comments. We are particularly grateful for the assistance given by George Barton and Katia Pascarella in collecting the data.

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

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Since the global financial crisis in 2007, the G20 has overseen an ambitious program of regulatory reform in financial markets. One goal of the reform program is to make derivative markets safer by reducing systemic risk and improving counterparty risk management. For this reason, many standardized over-the-counter (OTC) derivative contracts must now be cleared through central counterparties (CCPs). For example, in the EU, specific classes of interest rate derivatives[1] must be cleared through CCPs. More recently, as CCPs have increased notably in their size and many institutions have become more exposed to CCPs, authorities are examining the role CCPs may play as a source of stress in the financial system (Alfranseder et al., 2018; Bank for International Settlements, 2018; Markose et al., 2017).

In this paper we analyze the nature of the interconnectedness in the UK derivative markets to understand existing interdependencies across different derivative markets and to give a high-level overview of potential channels for the transmission of liquidity shocks in the system. In order to do this we follow two complementary approaches. First, we use a state-of-the-art centrality measure, based on the work of Iacovacci et al. (2016), to assess the vulnerability of financial institutions. Second, we draw on the recent contributions on contagion mechanisms in OTC derivative markets (Heath et al., 2016; Paddrik et al., 2016; Paddrik and Young, 2017) to gain a better understanding of the transmission of liquidity stress. A recently growing body of research on financial networks shows that multiplex network[2] analysis reveals important built-in correlations that can be used to extract information otherwise unobtainable when aggregating across the layers (Berndsen et al., 2016; Bianconi, 2013, 2018; Boccaletti et al., 2014; Iacovacci et al., 2016; Menichetti et al., 2014; Molina-Borboa et al., 2015). Furthermore, Poledna et al. (2015) and Bargigli et al. (2015) show that considering single layers in isolation can lead to large misrepresentations of systemic risk.

This paper extends the existing academic literature on the OTC derivative markets in three main directions. First, we use UK trade state reports from DTCC and Unavista (as of 30th June 2016) that include data on both centrally and non-centrally cleared trades of CCP clearing members. This allows us to build a network of exposures between those institutions across three OTC markets: interest rate, credit, and foreign exchange derivatives. To the best of our knowledge, Abad et al. (2016) is the only other study on the same set of asset classes, but it only offers comprehensive survey of each derivative market separately. In contrast, we analyze the structural properties of the three derivative markets simultaneously in a multiplex form.

Second, we compare the extensions of the eigenvector centrality and of the PageRank centrality to weighted multiplex networks. To do this, we develop an extension of the centrality measure

[1] These comprise a subset of OTC interest rate derivative contracts denominated in the G4 currencies (EUR, GBP, USD, and JPY) including fixed-to-float interest rate swaps (IRS), basis swaps, forward rate agreements, and overnight index swaps. There are also OTC interest rate derivative contracts denominated in some non-G4 currencies (SEK, PLN, and NOK), which include fixed-to-float interest rate swaps (IRS) and forward rate agreements (FRA).

[2] A multiplex network is simply a collection of networks, called *layers*, each corresponding to a different market.

introduced in Iacovacci et al. (2016) to weighted networks. We observe that the computed eigenvector centrality is large for one CCP and very small for all the other institutions, largely because the eigenvector centrality is biased towards highly interconnected institutions, such as CCPs. In contrast, the picture emerging from the PageRank centrality is much more nuanced, with CCPs not ranking among the very top institutions.

Third, we extend the analysis first developed by Paddrik et al. (2016) from a financial network with only one layer for the Credit Default Swaps (CDSs) to a multiplex network made of three different markets – interest rate, credit, and foreign exchange derivatives, thereby capturing the interactions between them. We compute the impact of variation margin (VM) shocks in terms of potential deficiencies in expected payments between market participants. We do this by using an approach similar to Heath et al. (2016) in which shocks are randomly drawn from a normal distribution and by analyzing how much financial stress these shocks create as they propagate across the network. Paddrik et al. (2016) make different assumptions about the contagion mechanism that focus on solvency risks. For example, they allow institutions to meet their obligations by also using the initial margins they have collected, which should normally be usable only after the default of the institutions that posted them. On the contrary, we focus solely on the short-term dynamics of liquidity shocks across different OTC derivative markets in the form of reduced payments, neglecting potential solvency problems which might subsequently ensue.

Our results have important implications. We observe that our version of the PageRank centrality ranking, rather than the eigenvector centrality ranking, better highlights the most important institutions involved in the propagation of VM shocks in the system. Finally, our contagion model also allows us to estimate the extent to which each individual node contributes to liquidity stress amplification. It shows that, in extreme theoretical scenarios where liquidity buffers are small, a handful of institutions may experience significant spillover effects due to the directionality of their portfolios.

The paper is organized as follows: Section 2 describes the data and some descriptive statistics. Section 3 presents the structural properties of the multiplex network and compares the results coming from different centrality measures. Section 4 presents an analysis of the propagation of VM shocks across the multiplex network, and Section 5 concludes the paper.

2 Data

In response to the global financial crisis, G20 leaders agreed to make it mandatory for counterparties to derivatives transactions to report details of such contracts to trade repositories (TRs) in order to increase transparency. This initiative, known as the reporting obligation, allowed the access, for the first time, to more granular transactional data. The reporting obligation is being implemented across different jurisdictions: for example, in the United States, trade reporting is

part of the Dodd-Frank Act and was implemented by the Commodity Futures Trading Commission (CFTC) in December 2012. In the European Union, the reporting obligation has been implemented through the article 4 of the European Market Infrastructure Regulation (EMIR).

Under EMIR, all OTC and exchange-traded derivatives transactions undertaken by EU counterparties since August 2012 (or open at that point) have had to be reported by the following business day to a TR. The EMIR reporting obligation covers CCPs; financial counterparties, such as banks, insurance firms or pension schemes; and non-financial counterparties that are EU legal entities. However, since we have only access to the Bank of England's data, we can read trade reports as per the conditions stated in EMIR under Article 2 of Commission Delegated Regulation (EU) No 151/2013. In summary, this means that we can access reports on *a*) trades cleared by a CCP supervised by the Bank of England, *b*) trades where one of the counterparties is a UK entity *c*) trades where the derivative contract is referencing an entity located in the UK (e.g., CDSs on UK banks) or derivatives on UK sovereign debt, *d*) trades where one of the counterparties is supervised by the Prudential Regulatory Authority at the Bank of England, and *e*) aggregated position data for all derivative contracts referencing sterling.[3]

In this paper we focus on the three asset classes that correspond to the three largest markets for OTC derivatives: (i) interest rates derivatives (IR); (ii) credit derivatives (CD); and (iii) foreign exchange derivatives (FX). Using the TRs DTCC and UnaVista, for each asset class we build the network of exposures between institutions at the 30th June 2016. In this context, the exposure in a given asset class between two institutions is simply the aggregate net mark-to-market value of the outstanding contracts between them that fall within the specified asset class. Given that exposures are directional (the fact that i is exposed to j does not imply that j is exposed to i) the networks are directed and weighted, the weights being the exposures. We consider all the transaction of clearing members that are centrally or non-centrally cleared. The three networks corresponding to the different asset classes can be seen as three *layers* of a single multilayer, or more precisely multiplex network.[4]

The multiplex network of UK OTC derivatives comprises $M = 3$ layers and a total of $N = 2174$ nodes. Formally we denote with $\mathbf{A}^{\alpha}$ the directed and weighted adjacency matrix of the layer α whose elements A^{α}_{ij} are larger than zero if i has a net exposure (or net mark-to-market value of their position) to j of monetary values A^{α}_{ij}. Here $\alpha = \text{IR, CD, FX}$, while $i, j = 1, \ldots, N$. For simplicity we convert all exposures in US dollars.

In Figure 1 we illustrate the network topology associated with each of the three markets under consideration. Here we show the undirected and un-weighted multiplex network, which in any layer contains information only on which pairs of institutions have an outstanding contract between them that falls within that asset class. For each layer α, this information is encoded in the symmetric

[3]For this reason, a large number of EU and non-EU legal entities are also covered in our dataset.

[4]A multilayer network is more general than a multiplex network because it allows for edges that originate in one layer and terminate in another layer, which in our case do not have any economic meaning.

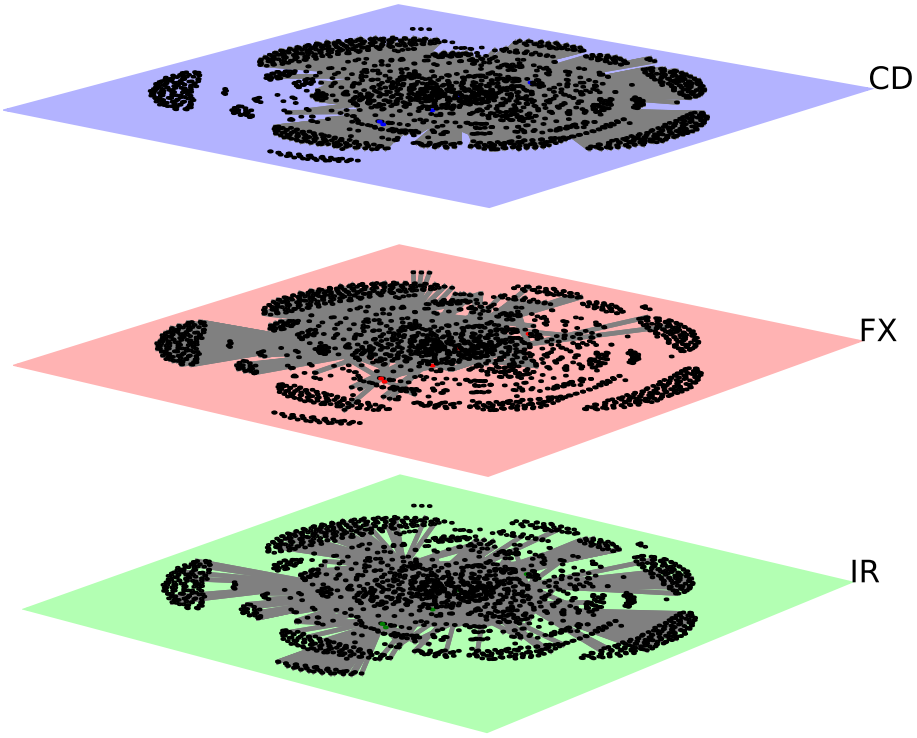

Figure 1: The network topology is shown for the three layers: credit (top-layer), FX (mid-layer), and interest rate (bottom-layer). CCP nodes are highlighted in a different colour for each layer. Loosely speaking, CCPs belong to the core of the three layers.

adjacency matrix $\mathbf{B}^{\alpha}$, in which $B^{\alpha}_{ij} = 1$ if i and j have at least one open position and $B^{\alpha}_{ij} = 0$ otherwise.

Table 1 shows the fraction of transactions centrally and non-centrally cleared in each market. IR transactions dominate the cleared market, followed by CD transactions. This is largely expected, as many of them are centrally cleared mandatorily. Consistently with the fact that central clearing is not mandatory for FX transactions, we find that very few of them are centrally cleared.

In order to carry out the analysis in Section 4, we use the liquidity buffer of CCPs in our sample. The liquidity buffer is intended as the sum of their cash and cash equivalents, which can be taken from their mid-year financial reports, which is consistent with the date of our exposures.

	Centrally cleared	Non-centrally cleared
IR	68.69%	31.31%
CD	8.47%	91.53%
FX	0.92%	99.08%

Table 1: Fraction of centrally cleared derivatives transactions as of 30th June 2016.

3 Structural properties of the UK derivative markets

Here we characterize the structure of this multiplex network using the recently introduced tools of multilayer network theory (Bianconi, 2013, 2018; Boccaletti et al., 2014; Menichetti et al., 2014). The first step in exploring interconnectedness of a financial system is to study properties of the unweighted network, such as the distribution of the activity of the nodes, the density of the different layers, the correlations among the degree of the same node in different layers, and the significance of specific patterns of interconnections across the layers compared to a random hypothesis.

Second, we focus on properties of the weighted network in which each layer is considered isolated from the others, such as the correlations between the number of interconnections of an institution and its total exposures.

Third, we focus on centrality measures specifically designed for multiplex networks. More specifically, we compare the extensions to weighted multiplex networks of the eigenvector centrality and of the PageRank centrality, arguing in favor of the latter. The goal of this comparison is to evaluate the validity of using different centrality measures by conducting an empirical study. While different centrality measures have previously been used in network analysis research to identify key nodes in a network, the validity of the centrality measures themselves is debatable. In other words, the research question that remains unanswered is: how exactly do these measures correlate with the real world? After all, the financial network is not based on just the ties between institutions. Additionally, market infrastructure rules can change the importance of any given node. For example, it is important to consider whether institutions will always remain crucial in a contagion model just because they occupy a more central location in a given financial network. Obviously, in such situations, false positives can be extremely detrimental because they could result in a waste of resources to investigate the problem or could imply unfounded market concerns. On the contrary, false negatives can also cause serious problems because relevant authorities might not be able to detect possible vulnerabilities in the financial system in a timely manner. Thus, there certainly is a need to test the efficacy and validity of individual centrality measures to correctly identify influential institutions in multiplex networks. We tackle this important issue in Section 4 by comparing the ranking of institutions given by the centrality measures to the ranking implied by our contagion model.

	Active nodes	Density	Average degree
IR	473	0.85%	7.99
CD	1469	0.17%	4.91
FX	553	0.33%	3.58

Table 2: Number of nodes active in each of the three layers (IR, CD, FX), alongside the density and average degree for each layer.

Institutions active in:		
One layer	Two layers	Three layers
89.3%	6.6%	4.1%

Table 3: Fraction of financial institutions active in one, two, or three layers.

3.1 Unweighted multiplex network properties

In order to characterize the dataset we start from the properties of the undirected and unweighted multiplex network shown in Figure 1.

We first investigate the activity of the nodes in the multiplex network. A financial institution is active in one layer if it has least one contract in that layer. In network terms, institution i is active in layer α if it is not an isolated node in that layer since its degree in layer α is larger than zero. The degree κ_i^α of node i in layer α is simply defined as the number of the edges to which node i participates, shown by $\kappa_i^\alpha = \sum_{j=1}^{N} B_{ij}^\alpha$.

Each layer is characterized by the number of financial institutions active in it, reported in Table 2. The CD layer is the layer that includes most active nodes while the IR and FX layers have the same order of magnitude of active nodes. Furthermore, in Table 2, we also report the density of each layer (i.e., the fraction of existing exposures over all the possible exposures) and the average degree of institutions (i.e., the average number of exposures that each institution has). The IR layer is by far the densest, followed by FX and the CD layers. Additionally we can define the activity a_i of a financial institution i as the number of layers in which the financial institution is connected (i.e., it has at least a contract). It follows that an activity $a_i = 1$ indicates that the financial institution is only active in one layer, $a_i = 2$ indicates that the financial institution is active in two layers, and $a_i = 3$ indicates that it is active in each of the three layers. In Table 3 we report the fraction of financial institutions $\delta(a)$ with activity $a_i = a \in \{1, 2, 3\}$. From this table we can see that most of the financial institutions in the dataset are active in just one layer and that the percentage of institutions that are connected in more than one layer is 10.7%.

We then characterized the degree correlations among different layers. The degree correlations reveal whether highly connected institutions in one market (layer) are also typically highly connected institutions in another market or if the reverse is true. These correlations among the degree

	IR	CD	FX
IR	1.00	-0.24(4)	0.14(3)
CD	-0.24(4)	1.00	-0.36(5)
FX	0.14(3)	-0.36(5)	1.00

Table 4: Kendall τ correlation coefficient between the degrees in each pair of layers. The error evaluated by randomizing the node labels in the different layer is always in the last digit and for each pair of layer is always lower than 0.08.

of the same node in different layers can be captured by the Kendall τ correlation coefficient K. In Table 4 we report the results and assess the error by comparing these correlations with the results obtained using the same procedure when we randomize independently the node labels in the different layers. We observe that although the correlations are low they are always significant. The degree in the layer CD appears to be negatively correlated with the degree in the layers FX and IR, while the degree in the layers IR and FX are positively correlated.

Finally, we characterize the connections between institutions in terms of multilinks. The multilink between nodes i and j is defined as $\vec{m}_{ij} = (B_{ij}^{\mathrm{IR}}, B_{ij}^{\mathrm{CD}}, B_{ij}^{\mathrm{FX}})$. For example, if nodes i and j have an open position both in the IR market and in the FX market, but not in the CD market, then $\vec{m}_{ij} = (1, 0, 1)$. In Figure 2 we show an illustrative example of a small multilayer network that includes all possible kinds of multilinks. We denote with $\mathcal{N}(\vec{m})$ the number of multilinks of the kind $\vec{m}$ present in the network. For example, $\mathcal{N}((1, 0, 1))$ is the number of multilinks of the kind $(1, 0, 1)$. In Bianconi (2013) several classes of null models for multilayer networks are introduced and $\overline{\mathcal{N}}(\vec{m})$, the corresponding average number of multilinks of kind $\vec{m}$ (for a network of a given size), are computed. We indicate with $\overline{\mathcal{N}}(\vec{m})$ the average and with $\sigma(\mathcal{N}(\vec{m}))$ the standard deviation of the number of multilinks $\vec{m}$ across the ensemble of all networks consistent with the features that one has chosen to use. Therefore, $\overline{\mathcal{N}}(\vec{m})$ provides a natural benchmark for $\mathcal{N}(\vec{m})$, the number of multilinks of kind $\vec{m}$ actually observed. By defining the zeta-score

$$\theta(\vec{m}) = \frac{\mathcal{N}(\vec{m}) - \overline{\mathcal{N}}(\vec{m})}{\sigma(\mathcal{N}(\vec{m}))}\,, \tag{1}$$

we can quantify the relative abundance or scarcity of the multilink of kind $\vec{m}$. For example, if $\theta(\vec{m}) = 2$, the number of multilinks of kind $\vec{m}$ is two standard deviations away from the number that would be observed in a typical large network having the chosen features. Here we fix the number of nodes and the degree of each node in each layer. In Figure 3, we show the indicator $\theta(\vec{m})$ for all kinds of multilinks $\vec{m}$. We can see that the multilinks $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$ are under-represented, while all the multilinks for which more than one edge per layer exists are over-represented. In particular, the multilink corresponding to an edge present both in the IR and in the FX market is 30% more likely to occur than in the null model. This means a shock in the

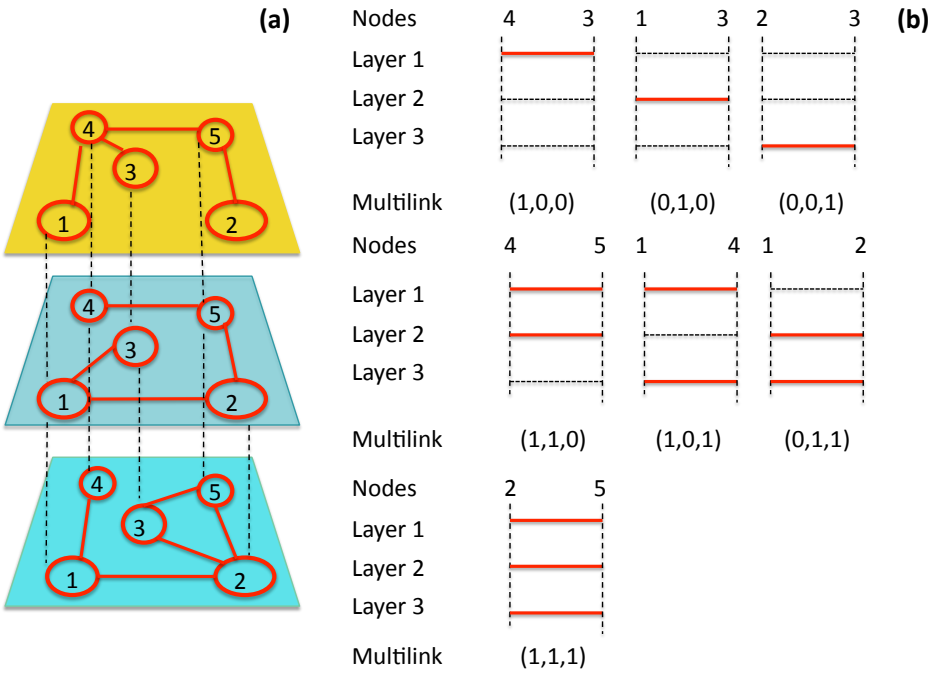

Figure 2: Illustrative example of a multilayer network of five nodes ($N = 5$) and three layers ($M = 3$) that includes all possible kinds of multilinks.

IR market will likely impact institutions operating in the FX market as well.

3.2 Weighted multiplex network properties

An important feature of several weighted networks is the existence of correlations between strengths and the degrees. The strength of a node is simply the sum of weights of all incoming and outgoing edges. Such correlations indicate that strengths are not independent on the degree of institution. In our specific case we want to test whether more interconnected institutions typically also have bigger exposures. To this end, for each layer α and for each node i, we measure the correlations existing between the degree κ_i^α and the total *un-netted* exposure s_i^α given by

$$s_i^\alpha = \sum_{j=1}^{N} A_{ij}^\alpha + \sum_{j=1}^{N} A_{ji}^\alpha \,, \tag{2}$$

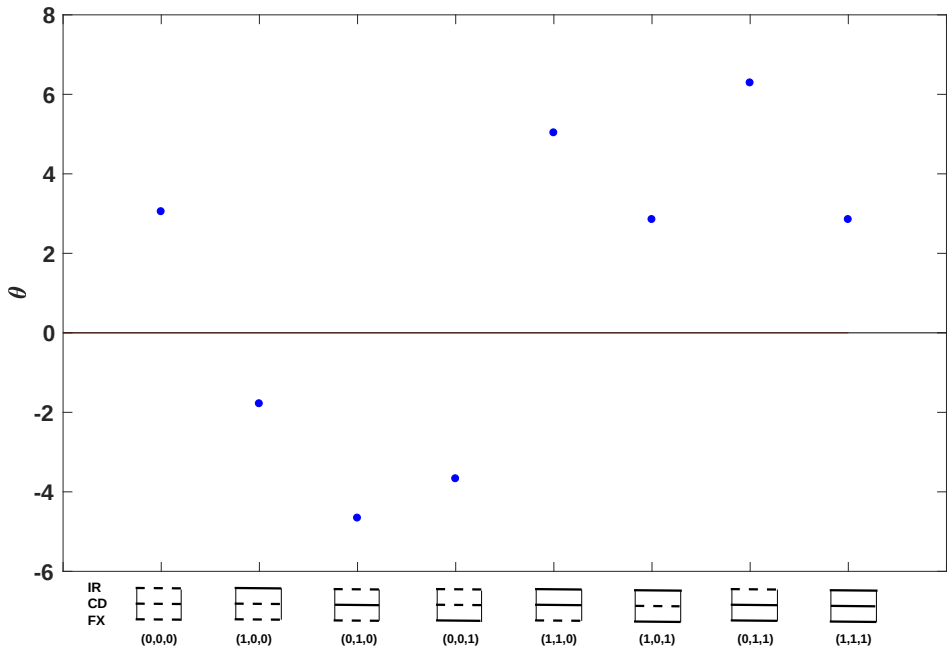

Figure 3: Relative abundance (positive values) or scarcity (negative values) of all kinds of multi-links with respect to a typical large network with the same degree sequence.

where the first term represent the total exposure of i towards other institutions, while the second term represent the total exposure of other institutions towards i.

In Figure 4 we plot the average exposure $\bar{s}_k^\alpha$ of contracts in layer α for financial institutions having a degree in layer α given by $\kappa^\alpha = k$. The data seem to suggest a super-linear increase of $\bar{s}_k^\alpha$. More precisely, it points to the functional dependence:

$$\bar{s}_k^\alpha \simeq k^{\beta_1} , \tag{3}$$

where $\beta_1 > 1$. In order to check this hypothesis we run a linear regression using the associated linear model:

$$\ln \bar{s}_k^\alpha = \beta_0 + \beta_1 \ln k + \epsilon. \tag{4}$$

The values of β_1, its standard deviation σ_{β_1} and the corresponding values of R^2 are reported in Table 5. As expected, the estimated value of β_1 is greater than one for all layers. This means that the most connected institutions are also the institutions with larger exposures. However, this result is more significant for the IR and CD layers than for the FX layer.

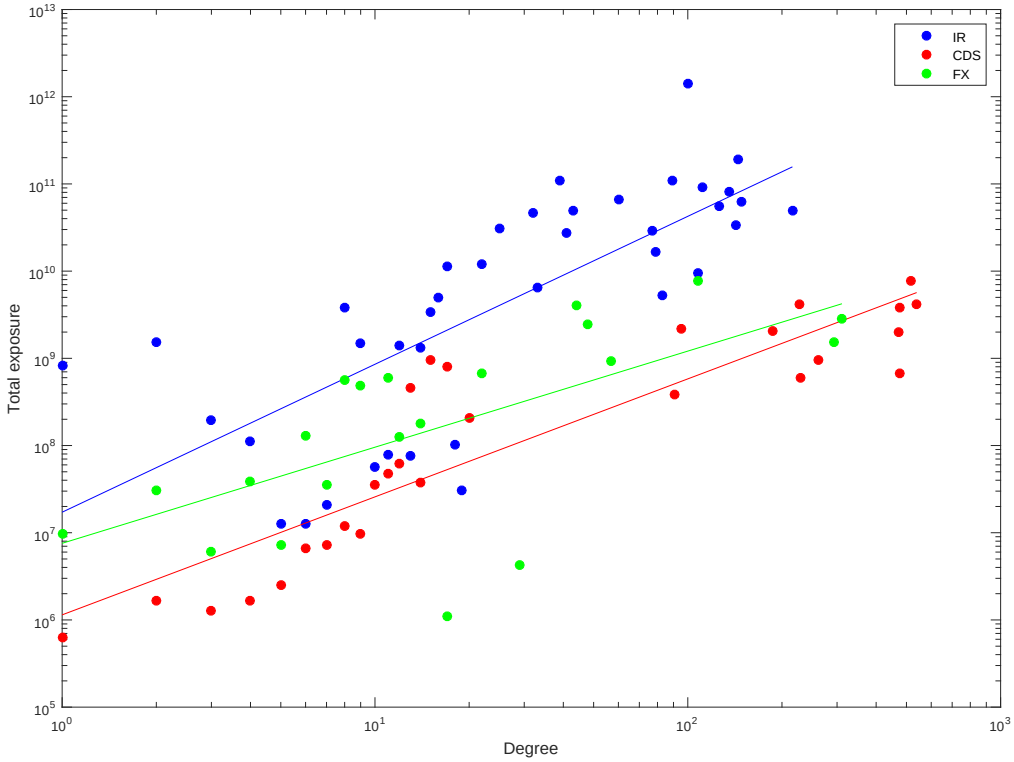

Figure 4: The total exposure $\bar{s}_k^\alpha$ of institutions with degree $\kappa^\alpha = k$ is shown for the three layers (IR, CD, FX) together with the fit provided by the statistical model of Eq. (4).

3.3 Functional Multiplex PageRank and vulnerability of financial institutions

One standard way to gauge the relative importance of financial institutions is to compute their centrality. In the literature of financial networks the centrality of a given node is normally considered a proxy for the vulnerability of that node to shocks (Alter et al., 2015; Heath et al., 2016; Markose et al., 2017; Siebenbrunner, 2017). While there is certainly no shortage of centrality measures, it is only recently that suitable centrality measures for multiplex networks have been proposed. In this section we introduce two centrality measures for single-layered networks – the eigenvector centrality and the PageRank centrality – and we re-interpret them from the perspective of the propagation of shocks in a financial network. We will then generalize both centralities to a multiplex setting.

In graph theory, the eigenvector centrality is a popular measure of the importance of a node (or institution) in a network. Relative scores are assigned to all institutions in the network based on the assumption that connections to high-scoring institutions contribute more to the score of a

	IR	CD	FX
β_1	1.70	1.35	1.10
σ_{β_1}	0.25	0.13	0.29
R^2	0.55	0.812	0.43

Table 5: The regression coefficients β_1 of the statistical model Eq. (4) are shown for the three layers (IR, CD, FX) together with its standard deviation σ_{β_1} and the corresponding values of R^2.

given institution than equal connections to low-scoring institutions. Therefore, a high score means that an institution is connected to many institutions who themselves have high scores. However, by construction, CCPs will get a very high score since they are always connected to institutions with high scores. The PageRank centrality, like the eigenvector centrality, is based on the idea that the vulnerability of a financial institution increases if it has a net exposure to already vulnerable institutions. Crucially, when applied to our networks, the PageRank centrality overcomes one of the biggest disadvantages of the eigenvector centrality. In fact, the eigenvector centrality tends to be highly degenerate as it is always vanishing for each node without exposures to other institutions. In this paper, we apply an extension of the Functional Multiplex PageRank (FMP) centrality recently introduced in Iacovacci et al. (2016) to a weighted multiplex network. The FMP centrality is a natural extension of the PageRank centrality, widely used for single-layer networks. It provides one of the most successful ranking algorithms and it has been used in a broad range of fields, ranging from on-line information to off-line and on-line social networks.

We denote the centrality of node i with $X_i \geq 0$ and the corresponding vector of centralities with $\mathbf{X}$. In all the cases that we consider, the centrality of each node (i.e., institution) will depend on the centrality of its neighbors (i.e., of its counterparties). However, the centrality of those neighbors will depend on the centrality of their own neighbors, and so on. Ultimately, the centrality of each node will be a function of the centrality of all the other nodes, given by

$$\mathbf{X} = f(\mathbf{X})\,, \tag{5}$$

where the function f depends on the specific centrality measure. The centrality of all nodes is computed by iterating (5) until a suitably accurate approximation of its fixed point is found. In practice, one chooses an initial point $\mathbf{X}^{(0)}$ and progressively computes the subsequent iterations $\mathbf{X}^{(1)} = f(\mathbf{X}^{(0)})$, $\mathbf{X}^{(2)} = f(\mathbf{X}^{(1)})$, and so on.[5] Let us now interpret $\mathbf{X}^{(0)}$ as the vector of initial exogenous shocks. $\mathbf{X}^{(1)}$ will incorporate the effect of the first round of shock propagation and, in general, $\mathbf{X}^{(s)}$ the effect of the first s rounds of shock propagation. The centrality measure, given by the solution of (5), will incorporate the effects of infinite rounds of propagation and therefore

[5] In most cases the solution of $\mathbf{X} = f(\mathbf{X})$ such that $X_i \geq 0$, $\forall i$ is unique and therefore independent of the initial point $\mathbf{X}^{(0)}$.

represent the steady state reached by the system. Typically, centrality measures are normalized to one, i.e. $\sum_i X_i = 1$ and we can imagine that f is suitably defined so that such normalization is kept at all rounds. As a consequence, the overall size of the shock will be the same regardless of the centrality measure that we choose. However, different centrality measure will locally propagate shocks to their neighboring nodes in different ways.

The eigenvector centrality is defined by the normalized eigenvector corresponding to the largest eigenvalue of the (transposed) weighted adjacency matrix:

$$\lambda_1 X_i^{\mathrm{EC}} = \sum_j A_{ji} X_j^{\mathrm{EC}} \,. \tag{6}$$

In practice, $\mathbf{X}^{\mathrm{EC}}$ can be computed by the power iteration algorithm, which consists in iterating

$$X_i^{\mathrm{EC},(s+1)} = \sum_j A_{ji} X_j^{\mathrm{EC},(s)} \tag{7}$$

and by normalizing the centralities at each step. Let us illustrate one step of the algorithm with a small network consisting of four institutions and the following non-zero exposures: $A_{31} = 1$, $A_{41} = 2$, $A_{32} = 1$. Starting from initial shocks $X_1^{(0)} = X_2^{(0)} = X_3^{(0)} = X_4^{(0)} = 1/4$, after the first round the shock is transmitted from institution 1 to institutions 3 and 4 and from institution 2 to institution 3, so that $X_3^{(1)} = X_4^{(1)} = 1/2$ and $X_1^{(1)} = X_2^{(1)} = 0$. It is worth noting that the eigenvector centrality of all institutions that are not exposed to other institutions is equal to zero.[6] Hence, if we used the eigenvector centrality as a proxy for the vulnerability of institutions, we would find that institutions that are not exposed to other institutions are not vulnerable to shocks. In other words, the vulnerability of an institutions depends only on the fact that it might be exposed to other institutions, and it does not depend on any external factors. Moreover, because their centrality is zero, institutions that are not exposed to other institutions do not even transmit any shock to their own counterparties.

The weighted PageRank centrality is defined by

$$X_i^{\mathrm{PR}} = \mu \sum_j \frac{A_{ji}}{g_j} X_j^{\mathrm{PR}} + \nu \tag{8a}$$

$$g_j = \begin{cases} \sum_k A_{jk} & \text{if } \sum_k A_{jk} > 0 \\ 1 & \text{otherwise} \end{cases} \tag{8b}$$

where μ is a parameter (usually called the damping factor) ranging between zero and one, while ν is

[6] If node i is not exposed to any other institution, then $A_{ji} = 0$, $\forall j$. As a consequence, from the definition of eigenvector centrality it readily follows that $X_i = 0$.

set such that $\sum_i X_i = 1$.[7] The definition of g_j ensures the first term in (8a) is always well-defined. Specifically, it is equal to $A_{ji}/\sum_k A_{jk}$ if $A_{ji} \neq 0$, and it is equal to zero otherwise. Analogously to (7) we can write:

$$X_i^{\mathrm{PR},(s+1)} = \mu \sum_j \frac{A_{ji}}{g_j} X_j^{\mathrm{PR},(s)} + \nu \,. \tag{9}$$

Let us initially discuss the case in which $\mu = 1$. Expanding on the example above, we can see that in this case $X_3^{(1)} = 2/3$, $X_4^{(1)} = 1/3$, and $X_1^{(1)} = X_2^{(1)} = 0$. The key point is that here shocks are locally transmitted to counterparties pro rata. This means that, before centralities are normalized, each institution redistributes its shock to its neighbors so that the sum of shocks received by them is equal to the shock transmitted. Such a transmission mechanism resembles more closely the Eisenberg and Noe model (Eisenberg and Noe, 2001), a variation of which will be introduced in Section 4. Moreover, differently from the eigenvector centrality, all institutions (including those that are not exposed to other institutions) now receive a contribution to their centrality equal to the normalization constant ν. Indeed, in the extreme case in which $\mu = 0$, the centrality is equal to $\nu = 1/N$ for all institutions and does not depend on the network of exposures at all. In the intermediate case in which $0 < \mu < 1$, the parameter μ determines how important the contribution of the network of exposures is in the computation of centralities.[8]

In a multiplex network the centrality of a node will be unavoidably dependent on the relevance attributed to edges within different layers. In practice, this corresponds to the fact that the exposure of an institution in one layer can contribute more or less to its vulnerability depending on the type of external shock that is affecting the market. From the perspective of defining the centrality, this translates into introducing a set of layer-specific weights, the so-called *influences*, that determine the relative importance of the each layer. In a nutshell, edges in a layer with a larger influence will contribute more to the centrality than financial institutions connected in layers with a smaller influence. As a consequence, the centrality will ultimately depend on the chosen set of influences. Rather than trying to estimate precise values for the influences, here we take an agnostic approach by focusing on the full distribution of the centrality across the whole range of values for influences. More specifically, we will compute both the typical value, which corresponds to the average over all the possible values of the influences, and the worst case value, which correspond to the maximum across all the possible values of the influences.

Formally we denote the set of influences of the three layers with $\mathbf{z} = (z^{\mathrm{IR}}, z^{\mathrm{CD}}, z^{\mathrm{FX}})$ and the centrality of node i, which will depend on the influences, with $X_i(\mathbf{z})$. The Functional Multiplex

[7] To do this, we assume that $\nu = 1/N \sum_j \left(1 - \mu \mathbb{1}_{\sum_k A_{jk}}\right) X_j$, where the indicator function $\mathbb{1}_x$ is equal to one when $x > 0$ and is equal to zero otherwise.

[8] For the damping factor μ, we use the same values as Langville and Meyer (2004) taking $\mu = 0.85$.

Eigenvector Centrality (FMEC) is defined by

$$\lambda_1 X_i^{\mathrm{FMEC}}(\mathbf{z}) = \sum_j \sum_\alpha A_{ji}^\alpha z^\alpha X_j^{\mathrm{FMEC}}(\mathbf{z}) \,, \tag{10}$$

where λ_1 is the largest eigenvector of the matrix C, with $C_{ji} = \sum_\alpha A_{ji}^\alpha z^\alpha$. Eq. (10) is the straightforward extension of (6) in which each layer is weighted with the corresponding influence. Similarly, the Functional Multiplex PageRank Centrality (FMP) is defined by[9]

$$X_i^{\mathrm{FMP}}(\mathbf{z}) = \mu \sum_{j,\alpha} \frac{A_{ji}^\alpha}{g_j} z^\alpha X_j^{\mathrm{FMP}}(\mathbf{z}) + \nu \tag{11a}$$

$$g_j = \begin{cases} \sum_{k,\alpha} A_{jk}^\alpha z^\alpha & \text{if } \sum_{k,\alpha} A_{jk}^\alpha z^\alpha > 0 \\ 1 & \text{otherwise} \end{cases} , \tag{11b}$$

where again ν is the normalization constant.[10] Without loss of generality we can consider the case in which $\sum_\alpha z^\alpha = 1$. Indeed, if $\sum_\alpha z^\alpha \neq 1$ we can perform the replacements $z^\alpha \to z^\alpha / \sum_\alpha z^\alpha$, which leaves (11) unchanged. Hence, it is possible to represent the entire function $X_i(\mathbf{z})$ as a contour-line plot over a triangle.

For the reasons above, we now focus our analysis on the FMP centrality and, whenever this does not generate confusion, we drop the superscript in the notation for the centralities. We perform a quantitative comparison between the FMP centrality and the FMEC centrality in Section 3.3.1.

In order to rank financial institutions according to their centrality, first we introduce the Absolute Functional Multiplex PageRank (Abs FMP) $\hat{X}_i^{\mathrm{FMP}}$ of node i given by

$$\hat{X}_i^{\mathrm{FMP}} = \left\langle X_i^{\mathrm{FMP}}(\mathbf{z}) \right\rangle_{\mathbf{z}} \,, \tag{12}$$

where the average is computed across all the values of $\mathbf{z}$. The resulting ranking of the nodes can be taken as proxy for which financial institution is more vulnerable by averaging over different scenarios modelled by the different values of the parameter $\mathbf{z}$. Second, we introduce the Maximum Functional Multiplex PageRank (Max FMP) $X_i^{\star\mathrm{FMP}}$ given by

$$X_i^{\star\mathrm{FMP}} = \max_{\mathbf{z}} X_i^{\mathrm{FMP}}(\mathbf{z}) \,. \tag{13}$$

which can be taken as a proxy for the vulnerability of financial institutions in the worst case

[9] We note that this definition of the FMP differs from the original one proposed in Iacovacci et al. (2016). Here a different influence is associated to each layer, while in the original version of the algorithm a different influence is associated to each different type of multilink. This modification allows us to extend the use of the FMP originally proposed for multiplex networks of $M = 2$ to the multiplex network of the UK derivative market formed by $M = 3$ layers. This definition can also be used with weighted adjacency matrices.

[10] In this case $\nu = 1/N \sum_j \left(1 - \mu \mathbb{1}_{\sum_{k,\alpha} A_{jk}^\alpha z^\alpha}\right) X_j(\mathbf{z})$.

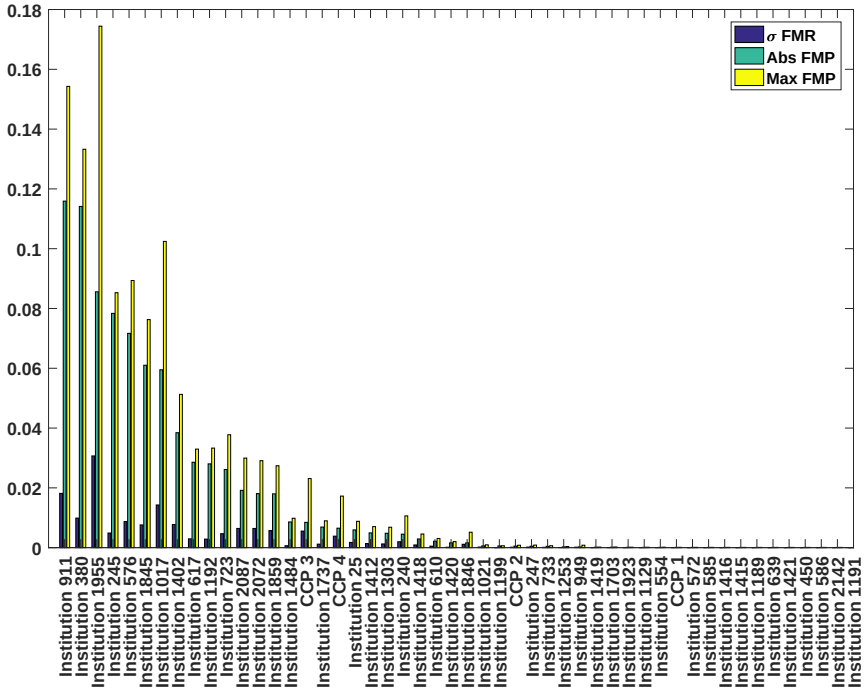

Figure 5: The ranking of the top 50 most vulnerable financial institutions according to the Abs FMP is compared with to maximum and the standard deviation.

scenario. Finally every institution i can also be characterized by the standard deviation of the Functional Multiplex PageRank from its mean (σ FMP) σ_i^{FMP}

$$\sigma_i^{\mathrm{FMP}} = \left\langle \left[X_i^{\mathrm{FMP}}(\mathbf{z}) - \hat{X}_i^{\mathrm{FMP}} \right]^2 \right\rangle_{\mathbf{z}} . \tag{14}$$

A high value of σ_i^{FMP} indicates that the financial institution i can be much more vulnerable in some scenarios than in others, whereas a low value of σ_i^{FMP} indicates that the financial institution i is more or less equally vulnerable with respect to different possible scenarios.

In Figure 5 we show the institutions with the largest 50 values of Abs FMP, and for the same set of institutions we report the Max FMP, and the σ FMP as well. We observe that the top 15 most vulnerable financial institutions according to the Abs FMP do not include CCPs. Additionally, for these institutions, σ FMP is typically much smaller than Abs FMP and Max FMP. Although the CCPs are not listed among the most vulnerable financial institutions, it should be noted that they are characterized by a σ FMP very close to the Abs FMP, which indicates that their vulnerability varies significantly across the full range of influences.

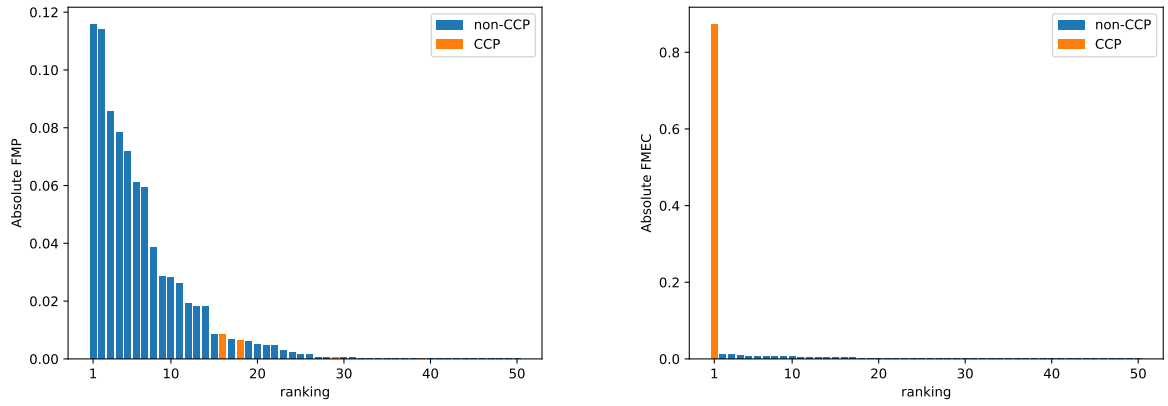

Figure 6: The ranking of the top 50 financial institutions according to the Abs FMP (left-chart) is compared with the ranking of the top 50 financial institutions according to the Abs FMEC (right-chart).

3.3.1 PageRank versus eigenvector centrality

In this section we perform a quantitative comparison between the FMP and the FMEC centrality. To this end, in analogy with the Abs FMP we define the Absolute Functional Multiplex Eigenvector Centrality (Abs FMEC) as

$$\hat{X}_i^{\mathrm{FMEC}} = \left\langle X_i^{\mathrm{FMEC}}(\mathbf{z}) \right\rangle_{\mathbf{z}} . \tag{15}$$

Using the FMP to study a directed multiplex network has great advantages with respect to the FMEC. In fact, as with the eigenvector centrality for single layers, the FMEC also gives almost no value to the nodes (that are not CCPs) in the out-component of multiplex networks, often giving highly degenerate results. One possible explanation for these results is that, by construction, high eigenvector scores are always assigned to CCPs since they are connected to many institutions which have high scores themselves. Here we demonstrate the advantage of considering the FMP over the FMEC in the case of our dataset. In Figure 6 we compare the ranking provided by the FMP to the ranking provided by the FMEC. It is evident that by associating a vanishing centrality with each node in the out-component the FMEC produces a very degenerate ranking in which many institutions have very small and comparable Abs FMEC while a single CCP is associated with the largest vulnerability. In the next section, we will also compare the results from the contagion model with the ranking provided by the centrality measures discussed above.

4 Analysis of the variation margin shocks propagation

In the previous section we argued that the FMP centrality is a better proxy for the vulnerability of a financial institution than the eigenvector centrality. Given that centrality measures do not take into consideration the various details of the contagion mechanism, they can be used when shocks are not known a priori or when the details on how shocks propagate between institutions are not known. This includes the case in which the contagion mechanism is known, but some of the parameters required are particularly difficult to estimate. However, when the contagion mechanism is known, it is possible to compare the results of the two approaches. This makes it possible to check whether our centrality measures are a good indicator of the vulnerability of financial institutions.

In this section we take a complementary approach by studying how a VM shock would propagate across a network of financial institutions. Here we assume an exogenous VM shock (i.e., the shock is known a priori), and we make specific assumptions about how the shock propagates across the network. To do this, we simulate the key design features of a liquidity contagion model and trace directly the channels by which stress could be transmitted through the system in extreme market conditions. Finally, we compare the ranking derived from the contagion model with the ranking implied by the centrality measures described in the previous section.

4.1 Liquidity contagion model

We assume that VM shocks are triggered by movements in the mark-to-market value of derivative contracts. When the value of the contract held by an institution undergoes a particularly large negative variation, its counterparty will typically ask the institution to post some VM. In Paddrik et al. (2016) VMs are computed based on the impact the CCAR stress test has on credit ratings, which in turn influence the net present value of the open derivative positions. Given the greater complexity of our network, which includes different classes of contracts, here we take a simpler approach whereby shocks are estimated by means of an impact valuation model (Bardoscia et al., 2015, 2017; Cerezetti et al., 2017; Heath et al., 2016). This allows us to disregard the complexity of repricing all derivative contracts and, at the same time, it permits us to study the response of the system to a wide class of shocks. We compute VM shocks by using the approach followed in Heath et al. (2016), which simulates extreme changes to OTC derivative exposures. There, VM shocks are drawn from a normal distribution with a zero mean and a variance proportional to both the size of the exposures and to the average volatility of a given market. Here we simply extend that approach to the three markets (IR, FX, CD) in our analysis. More precisely, by denoting with σ_α^2 the average volatility of the returns of contracts in the layer α and by denoting with $\bar{p}_{ij}^{\alpha}$ the variation margin that institution i is required to post to institution j for contracts in the layer α,

we have that $\bar{p}_{ij}^{\alpha} \sim A_{ij}^{\alpha}\mathcal{N}(0, \sigma_{\alpha}^2)$, where again A_{ij} represents the exposure from i to j in layer α.[11]

The starting point for the analysis is that each node needs to deliver and/or receive VM payments from its counterparties. If node i fails to post the required VM to its counterparty j, this might have knock-on effects, such as node j not being able to post the required VM to its own counterparties. This is likely to depend on whether j has a sufficient buffer of liquid assets and on whether its other counterparties are able to make the due payments. Ultimately, the payments that each node is able to make will depend on the payments that all the other nodes are able to make.

The formal set-up of our experiment is an extension to a multiplex network of Paddrik et al. (2016), which is in turn based on the Eisenberg and Noe algorithm (Eisenberg and Noe, 2001). However, as we explain below, an important difference in our algorithm, compared with the one used by Paddrik et al. (2016), is the treatment of initial margins. As mentioned previously, $\bar{p}_{ij}^{\alpha}$ denotes the payment obligation of node i towards node j in the layer α. Similarly, we denote with p_{ij}^{α} the actual payment made by node i to node j for contracts in the layer α. Total payment obligations and total realized payments will be, respectively, $\bar{p}_i = \sum_{j,\alpha} \bar{p}_{ij}^{\alpha}$ and $p_i = \sum_{j,\alpha} p_{ij}^{\alpha}$. When $\bar{p}_i > 0$, we denote the relative amount owed by node i to node j in the layer α with $L_{ij}^{\alpha} = \bar{p}_{ij}^{\alpha}/\bar{p}_i$, otherwise we set $L_{ij}^{\alpha} = 0$. If $p_{ij}^{\alpha} < \bar{p}_{ij}^{\alpha}$, the shortfall faced by j from obligations with i in the layer α will be equal to $\bar{p}_{ij}^{\alpha} - p_{ij}^{\alpha}$. In Paddrik et al. (2016) the initial margin collected by node j from node i for contracts in layer α can be used by node j to offset the failure of node i to post the full VM requested for contracts in the layer α. However, here we will neglect initial margins. In fact, initial margins normally can be used only when the counterparties that posted it default. Neglecting initial margins corresponds to assuming that the failure of an institution to make a payment does not automatically trigger its default. In other words, institutions are willing to grant forbearance to their counterparties, hoping that they are able to secure the liquid resources they need in the near future. From this perspective, the difference between expected and realized payments should be interpreted as a measure of the short term payment shortfalls.

We define the stress faced by node i as the difference between all the payments it is required to make (the payment obligations) and all the payment inflows from its counterparties (the realized payments):

$$s_i = \sum_{j,\alpha} \bar{p}_{ij}^{\alpha} - \sum_{j,\alpha} p_{ji}^{\alpha} \,. \tag{16}$$

Ultimately, the payments that node i will be able to make will depend on whether it has any buffer of liquid assets that is able to further offset the stress it faces. However, for most institutions it is not easy to reliably estimate the extent of such buffers. We follow Paddrik et al. (2016) in assuming that

[11] We use the price volatility estimates deriving from the observations in our dataset. However, the covariance between price changes across products is assumed to be zero. In particular, we use the price dispersion measure of Jankowitsch et al. (2011), which they define as the average of the relative differences between individual execution prices and the average execution price.

stress is transmitted linearly between counterparties. This amounts to assuming that, whenever the payments made are smaller than the payments due, the difference is proportional to the stress faced by the node who is making the payment:

$$\bar{p}_i - p_i = \min(\tau_i s_i, \bar{p}_i)\,, \tag{17}$$

where the coefficient τ_i represents the transmission factor. The presence of the minimum in the equation above ensures that shortfalls in payments are not larger than the sum of all the payments due. A further assumption is that all creditors have the same seniority, which implies that, whenever obligations are not met in full, payments are made pro rata to creditors, shown by $p^\alpha_{ij} = L^\alpha_{ij} p_i$. From Eq. (17), we can derive $\bar{p}^\alpha_{ij} - p^\alpha_{ij} = \min(\tau_i L^\alpha_{ij} s_i, \bar{p}^\alpha_{ij})$ or equivalently:

$$p^\alpha_{ij} = \left[\bar{p}^\alpha_{ij} - \tau_i L^\alpha_{ij} s_i\right]_+ . \tag{18}$$

Because the stress s_i is the sum of payments unbalanced for institution i across all layers, an important consequence of (18) is that imbalances in payments in one layer will spill over to other layers as well. This is different from running the algorithm independently on all layers and then summing the final payments across layers. The reason is that, in that case, imbalances in payments occurring with different signs in different layers will not offset. However, it is possible to show that (18) leads to the same aggregate payments that we would get if we aggregate all the VM payments across all layers from the start. Given that we will present results also disaggregated by layer, we prefer to maintain the notation in which payments depend explicitly on the layers. From Eq. (17) we can see that, when $\tau_i = 1$ the difference between payments expected and payments made by node i is equal to the stress of node i. This corresponds to a situation in which node i does not have any additional liquidity buffer, and it therefore fully passes all the stress to its counterparties. Conversely, when $\tau_i = 0$ node i is always able to make the payments due, regardless of its stress. This corresponds to a situation in which node i can access an infinite liquidity buffer. By choosing τ_i within the range $[0, 1]$ one can explore all the intermediate scenarios that interpolate between those two extreme situations.

However, the liquidity buffers of some nodes can be known. For example, it is relatively easy to quantify liquidity buffers for CCPs. To keep the notation light, we will imagine that we know liquidity buffers only for node 0, but the formulation can also be extended to the cases in which liquidity buffers of more than one node are known.[12] In this case we do not need to make additional assumptions about how such nodes transmit stress to their counterparties. In fact, they will simply transmit all the stress that exceeds their liquidity buffers. This amounts to setting $\tau_0 = 1$ and to

[12] This decision is consistent with Paddrik et al. (2016), in which there is only one CCP whose index is 0.

redefining the stress for node 0 as:

$$s_0 = \left[\sum_{j,\alpha} \bar{p}^{\alpha}_{0j} - \sum_{j,\alpha} p^{\alpha}_{j0} - \gamma_0 \right]_+ , \tag{19}$$

where γ_0 is the size of the liquidity buffer. Realized payments are also computed, in this case via Eq. (20):

$$p^{\alpha}_{0j} = \left[\bar{p}^{\alpha}_{0j} - L^{\alpha}_{0j} s_0 \right]_+ . \tag{20}$$

By putting Eq. (16), Eq. (18), Eq. (19), and Eq. (20) together we get:

$$p^{\alpha}_{ij} = \left[\bar{p}^{\alpha}_{ij} - \tau_i L^{\alpha}_{ij} \left[\sum_{j,\alpha} \bar{p}^{\alpha}_{ij} - \sum_{j,\alpha} p^{\alpha}_{ji} \right] \right]_+ , \qquad \forall i \neq 0 \tag{21a}$$

$$p^{\alpha}_{0j} = \left[\bar{p}^{\alpha}_{0j} - L^{\alpha}_{0j} \left[\sum_{j,\alpha} \bar{p}^{\alpha}_{0j} - \sum_{j,\alpha} p^{\alpha}_{j0} - \gamma_0 \right]_+ \right]_+ . \tag{21b}$$

The Knaster-Tarski theorem ensures that the system in Eq. (21) has at least one fixed point. Moreover, there exists one greatest fixed point at which all institutions simultaneously maximize the payments made. It is easy to prove that the greatest fixed point can be computed by iterating Eq. (21) and by using the payment obligations $\bar{p}^{\alpha}_{ij}$ as the starting point. As a consequence, computing the greatest fixed point amounts to simulating the following process. Initially, all institutions assume that they will receive the payments due from their counterparties in full. If one institution faces a liquidity shortfall it will transmit the shortfall to its counterparties. If the deficiency in payments suffered by the counterparties is large enough, some of them will face a shortfall themselves. Those counterparties will in turn propagate the shortfall to their own counterparties. Deficiencies continue to propagate across the network until an equilibrium (the greatest fixed point) is reached. The payments corresponding to the greatest fixed point can therefore be interpreted as the realized payments between nodes.

A measure of the liquidity stress experienced by the system as a whole is the aggregate deficiency in VM payments D. By defining $d^{\alpha}_{ij} = \bar{p}^{\alpha}_{ij} - p^{\alpha}_{ij}$ as the deficiency in payments[13] from node i to node j on the layer α, we simply have:

$$D = \sum_{i,j,\alpha} d^{\alpha}_{ij} . \tag{22}$$

In all our simulations we make the additional assumption that transmission factors are homogeneous across the system, given by $\tau_i = \tau$, $\forall i$, and we display the average over all the realizations of the shock. In Figure 7 we plot the aggregate deficiencies D as a function of τ. For $\tau = 0$

[13] The deficiency from node i to node j is equal to the shortfall faced by j from obligations with i.

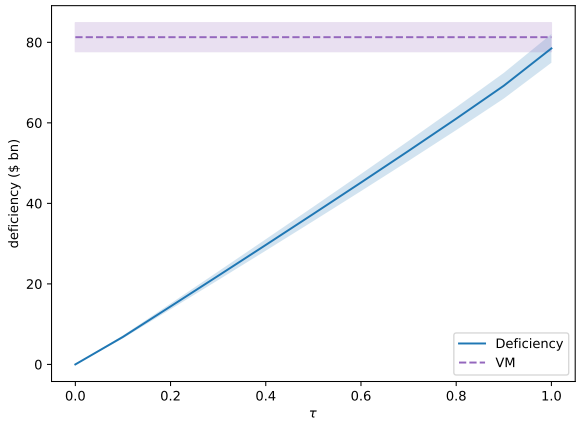

Figure 7: Aggregate VM payments that institutions are requested to make (purple dashed line) and aggregate deficiencies in those VM payments (blue solid line) as a function of the transmission factor τ, both expressed in US billion dollars. Lines are averaged over 100 realizations of the shock, while error regions span three standard deviations from the mean in both directions.

every node is able to sustain any shock and deficiencies are equal to zero. Conversely, for $\tau = 1$ deficiencies approach the total VM. In this case, each node is fully passing the stress not covered by their available resources to its counterparties. In fact, if a CCP faces a liquidity shortfall it will use all the cash at its disposal, regardless of its origin (e.g., margins, default fund contributions, the CCP's own capital, etc.). Therefore, the algorithm simply redistributes VM payments not covered by liquidity buffers. Between $\tau = 0$ and $\tau = 1$, we observe that deficiencies increase approximately linearly.

In Figure 8 we break down deficiencies by layers. We can see that the IR layer is by far the largest contributor to aggregate deficiencies. The reason is that the interest rate segment continues to account for the vast majority of outstanding OTC derivatives. At the end of June 2016, the notional amount of outstanding OTC interest rate derivative contracts totalled around 80% of the aggregate UK OTC derivatives market. Deficiencies in the IR layer increase linearly with τ, both in aggregate and for the CCPs. In CD and FX layers aggregate deficiencies increase linearly for most of the range of τ, but deficiencies of CCPs are a lot smaller. This is consistent with the smaller fraction of contracts that are centrally cleared in those layers.

In order to assess how the size of the shock can impact deficiencies we multiply the square root of the variance of the distribution from which shocks are drawn by a factor β, meaning variation margins in the layer α have a value equal to $\beta\sigma_\alpha$. From Figure 9 we see that for $\tau \geq 0.5$ deficiencies increase almost linearly with β, while for $\tau = 0.25$ deficiencies increase more rapidly when $\beta > 1$. As expected, deficiencies are very close to the initial VM shock for all values of β, when $\tau = 1$.

Next, we estimate how much each institution contributes to the aggregate deficiency. The contribution of institution i is computed in the following way. First, we run a baseline scenario in

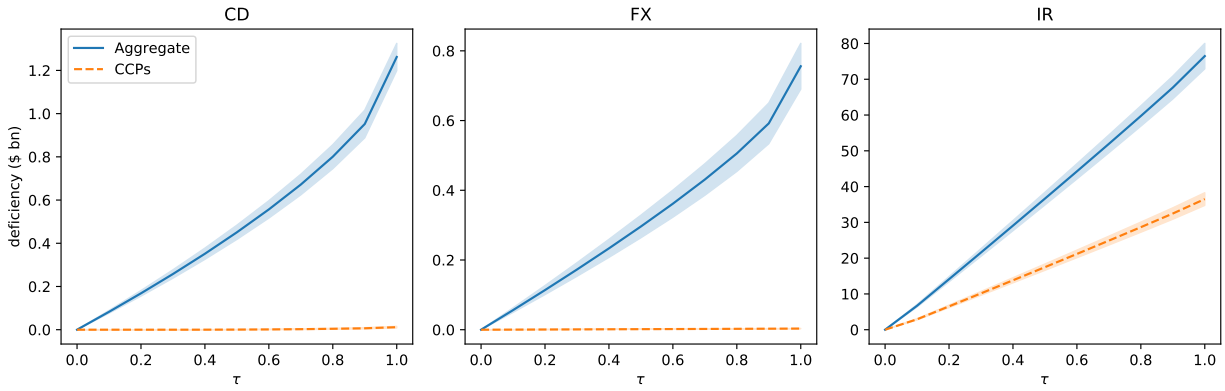

Figure 8: Deficiencies in VM payments broken down by layer as a function of the transmission factor τ. Blue solid lines refer to US billion dollar aggregate deficiencies, orange dashed lines refer to US billion dollar deficiencies of CCPs only. Lines are averaged over 100 realizations of the shock, while error regions span three standard deviations from the mean in both directions.

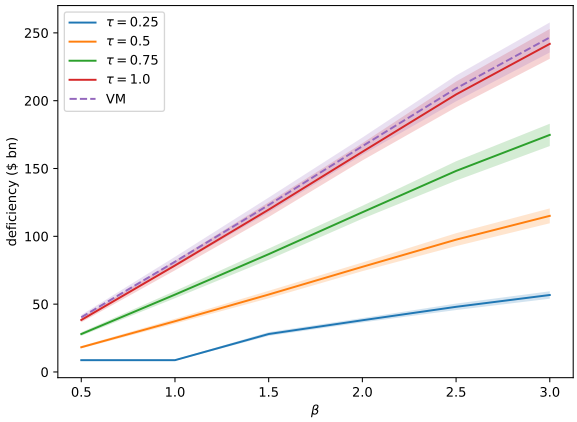

Figure 9: Aggregate deficiencies in VM payments expressed in US billion dollars as a function of the shock factor β, for different values of τ. Lines are averaged over 100 realizations of the shock, while error regions span three standard deviations from the mean in both directions.

which the transmission factor τ is equal for all institutions. Second, we set $\tau_i = 0$, while keeping $\tau_j = \tau$, for $j \neq i$. This corresponds to a fictitious scenario in which an infinite liquidity buffer has been assigned to node i, so that it will always be able to absorb any shortfall. As a consequence, it will not pass any stress to its counterparties. The difference between deficiencies computed in the two scenarios is taken as a measure of the contribution of institution i to the aggregate deficiency. In the left panel of Figure 10 we plot the 250 largest contributions to deficiencies in descending order. The distribution of contributions is quite broad, with about 0.5% of the institutions having contributions larger than $1 billion for $\tau = 0.5$. About 1.2% of the institutions

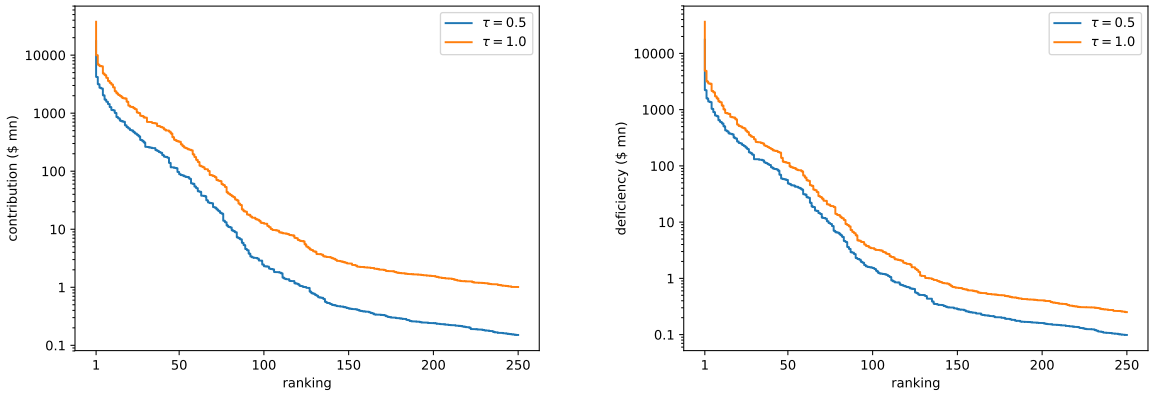

Figure 10: Top 250 individual institutions contributing to deficiencies (left panel) and experiencing deficiencies (right panel) in VM payments expressed in US million dollars. Rankings are based on deficiencies averaged over 100 (left panel) and 100 (right panel) realizations of random shocks, respectively.

have contributions larger than \$1 billion for $\tau = 1$. It should be noted that, when simulating shocks, the contributions to deficiencies need to be averaged over multiple realizations of random shocks, making the computational cost very large. For this reason, we used only 100 realizations of possible random shocks.

The contribution of an institution to the aggregate deficiency is a measure of how much that institution affects the spreading of shocks. However, in principle, it does not give us any information about how vulnerable that institution is to shocks. For this reason, we also compute the aggregate deficiencies for all the institutions in the network. In the right panel of Figure 10 we plot the largest deficiencies in descending order. For $\tau = 0.5$ about 0.3% of the institutions have deficiencies larger than \$1 billion, while for $\tau = 1$ about 0.6% have deficiencies larger than \$1 billion.

Finally, we compare the rankings obtained with the centrality measures Abs FMT and Abs FMEC (displayed in Figure 6) with the rankings of deficiencies in payments obtained with the contagion algorithm, for several values of τ. In order to avoid the noise from institutions with small deficiencies, we limit our analysis to the first 50 positions of the rankings. In Table 6 we report the corresponding Kendall rank correlation coefficients. We can see that, consistently across all values of τ and for different values of β ranging between 0.5 and 3, the correlation between the Abs FMT centrality and the deficiencies in payments is larger than the correlation between the Abs FMEC centrality and the deficiencies in payments. These results suggest that the Abs FMP centrality can be seen as a proxy for the vulnerability of institutions, whenever a more sophisticated contagion or stress testing mechanism is not available.

	$\sigma = 0.5$		$\sigma = 1.0$		$\sigma = 1.5$		$\sigma = 2.0$		$\sigma = 2.5$		$\sigma = 3.0$	
τ	FMT	FMEC	FMT	FMEC	FMT	FMEC	FMT	FMEC	FMT	FMEC	FMT	FMEC
0.1	0.453	0.352	0.510	0.373	0.420	0.358	0.491	0.285	0.492	0.342	0.468	0.393
0.2	0.453	0.352	0.495	0.324	0.422	0.357	0.428	0.280	0.494	0.344	0.463	0.357
0.3	0.453	0.355	0.497	0.325	0.425	0.360	0.428	0.280	0.494	0.344	0.468	0.362
0.4	0.455	0.357	0.497	0.328	0.428	0.363	0.430	0.282	0.495	0.345	0.471	0.365
0.5	0.456	0.358	0.503	0.335	0.430	0.365	0.433	0.282	0.495	0.345	0.473	0.366
0.6	0.458	0.360	0.503	0.335	0.435	0.370	0.438	0.286	0.499	0.348	0.459	0.321
0.7	0.458	0.360	0.503	0.338	0.438	0.373	0.443	0.291	0.499	0.348	0.461	0.326
0.8	0.440	0.340	0.508	0.343	0.422	0.326	0.448	0.293	0.484	0.301	0.464	0.329
0.9	0.445	0.345	0.508	0.346	0.425	0.329	0.453	0.298	0.484	0.304	0.471	0.335
1.0	0.458	0.365	0.515	0.320	0.432	0.345	0.461	0.309	0.504	0.324	0.480	0.345

Table 6: Kendall rank correlation coefficient between either of the two centralities ABS FMT and ABS FMEC and the deficiencies of individual institutions for different values of β. Rankings are based on deficiencies averaged over 100 realizations of random shocks.

5 Conclusions

In this paper we carry out a multiplex network analysis of the UK OTC market. Our dataset includes both centrally and non-centrally cleared contracts between CCP clearing members. We focus on assets that constitute the vast majority of traded contracts: interest rate, credit, and foreign exchange derivatives.

A key point of our analysis is that we fully account for the complexity implied by different asset classes, rather than aggregating across them. We do so by building a separate network for each derivative market and by joining them together in a single multiplex network. In particular, we are able to assess the systemic vulnerability of individual institutions by developing and computing centrality measures specifically designed for multiplex networks. Metrics such as the Abs FMP centrality measure described in Section 3 provide a good first approximation of the vulnerability of institutions, helping to identify those among them that may warrant closer attention. This could be a useful tool for regulators.

We also simulate how a liquidity shock arising from requests to post VM propagates across the system. If an institution does not have adequate liquidity buffers, it will pass on liquidity stress to its counterparties as it will not be able to make the expected payments. By extending the algorithm recently introduced in Paddrik et al. (2016) to multiplex networks, we compute both aggregated and layer-specific deficiencies in payments under multiple shock scenarios and quantify the systemic impact of individual institutions on such deficiencies. We find that, for τ between 0.5 and 1, between 0.3% and 0.6% of institutions experience materially large deficiencies, while between 0.5% and 1.2% of institutions give material contributions to the aggregate deficiency. We show that the rankings of vulnerability based on the Abs FMP centrality correlate reasonably well with rankings based on deficiencies computed via the contagion algorithm. This suggests that the Abs FMP centrality

could be used as a proxy for calculating the potential vulnerability of institutions.

This is particularly useful because running a realistic liquidity contagion algorithm can sometimes be problematic. For instance, retrieving the liquidity buffers for all the institutions under analysis is impractical when there are many institutions involved in the analysis. Furthermore, when simulating shocks, results should be averaged over multiple realizations, making the computational cost very large. On the contrary, centrality measures only depend on exposures and can be computed independently of shocks, thus not incurring large computational costs.

A possible future extension of our work is to use a multi-layer contagion mechanism like the one developed here to study solvency problems in which the initial margin requirements are specific to the different layers of the network representing the different OTC markets under analysis. Our work could also be extended to tie VM shocks to specific macroeconomic scenarios, this would allow us to carry out a full-fledged liquidity stress test for the UK OTC derivative markets.

References

Abad, J., Aldasoro, I., Aymanns, C., D'Errico, M., Rousová, L. F., Hoffmann, P., Langfield, S., Neychev, M., and Roukny, T. (2016). Shedding light on dark markets: First insights from the new eu-wide otc derivatives dataset. *European Systemic Risk Board Occasional Paper Series*, (11). Available from: https://www.esrb.europa.eu/pub/pdf/occasional/20160922_occasional_paper_11.en.pdf.

Alfranseder, E., Fiedor, P., Lapschies, S., Orszaghova, L., and Sobolewski, P. (2018). Indicators for the monitoring of central counterparties in the EU. Available from: https://www.esrb.europa.eu/pub/pdf/occasional/esrb.op14.en.pdf.

Alter, A., Craig, B., and Raupach, P. (2015). Centrality-based capital allocations and bailout. *International Journal of Central Banking*, 11(3):329–377.

Bank for International Settlements (2018). Analysis of Central Clearing Interdependencies. Available from: https://www.bis.org/cpmi/publ/d181.htm.

Bardoscia, M., Battiston, S., Caccioli, F., and Caldarelli, G. (2015). Debtrank: A microscopic foundation for shock propagation. *PloS One*, 10(6):e0130406.

Bardoscia, M., Battiston, S., Caccioli, F., and Caldarelli, G. (2017). Pathways towards instability in financial networks. *Nature Communications*, 8:14416.

Bargigli, L., di Iasio, G., Infante, L., and Pierobon, F. (2015). The multiplex structure of interbank networks. *Quantitative Finance*, 15:673–691.

Berndsen, R. J., León, C., and Renneboog, L. (2016). Financial stability in networks of financial institutions and market infrastructures. *Journal of Financial Stability*, 35:120–135.

Bianconi, G. (2013). Statistical mechanics of multiplex networks: Entropy and overlap. *Physical Review E*, 87(6):062806.

Bianconi, G. (2018). *Multilayer Networks: Structure and Function*. Oxford University Press.

Boccaletti, S., Bianconi, G., Criado, R., Del Genio, C. I., Gómez-Gardenes, J., Romance, M., Sendina-Nadal, I., Wang, Z., and Zanin, M. (2014). The structure and dynamics of multilayer networks. *Physics Reports*, 544(1):1–122.

Cerezetti, F., Sumawong, A., Shreyas, U. P., and Karimalis, E. (2017). Market liquidity, closeout procedures and initial margin for CCPs. *Bank of England Working Papers No. 643*. Available from: https://ssrn.com/abstract=2912261.

Eisenberg, L. and Noe, T. H. (2001). Systemic risk in financial systems. *Management Science*, 47(2):236–249.

Heath, A., Kelly, G., Manning, M., Markose, S., and Shaghaghi, A. R. (2016). CCPs and network stability in OTC derivatives markets. *Journal of Financial Stability*, 27:217–233.

Iacovacci, J., Rahmede, C., Arenas, A., and Bianconi, G. (2016). Functional multiplex pagerank. *Europhysics Letters*, 116(2):28004.

Jankowitsch, R., Nashikkar, A., and Subrahmanyam, M. G. (2011). Price dispersion in OTC markets: A new measure of liquidity. *Journal of Banking & Finance*, 35(2):343–357.

Langville, A. and Meyer, C. (2004). Deeper inside pagerank. *Internet Mathematics*, 1(3):335–380.

Markose, S., Giansante, S., and Shaghaghi, A. R. (2017). A systemic risk assessment of otc derivatives reforms and skin-in-the-game for CCPs. *Banque de France Financial Stability Review*, 21:111–126.

Menichetti, G., Remondini, D., Panzarasa, P., Mondragón, R. J., and Bianconi, G. (2014). Weighted multiplex networks. *PloS One*, 9(6):e97857.

Molina-Borboa, J.-L., Martinez-Jaramillo, S., van der Leij, M., and López-Gallo, F. (2015). A multiplex network analysis of the mexican banking system: link persistence, overlap and waiting times. *Journal of Network Theory in Finance*, 1(1):99–138.

Paddrik, M., Rajan, S., and Young, H. P. (2016). Contagion in the CDS market. *Office of Financial Research Working Paper*, 16-12. Available from: https://www.financialresearch.gov/working-papers/files/OFRwp-2016-12_Contagion-in-the-CDS-Market.pdf.

Paddrik, M. and Young, P. (2017). How safe are central counterparties in derivatives markets? Available from: https://ssrn.com/abstract=3067589.

Poledna, S., Molina-Borboa, J. L., Martínez-Jaramillo, S., Van Der Leij, M., and Thurner, S. (2015). The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability*, 20:70–81.

Siebenbrunner, C. (2017). Clearing algorithms and network centrality. Available from: https://ssrn.com/abstract=2959680.