BANK OF ENGLAND

Financial Stability Paper No. 38 – July 2016

Systemic risk in derivatives markets: a pilot study using CDS data

Robleh Ali and Nick Vause, Bank of England and
Filip Zikes, Federal Reserve Board of Governors

BANK OF ENGLAND

Financial Stability Paper No. 38 – July 2016

Systemic risk in derivatives markets: a pilot study using CDS data

Robleh Ali and Nick Vause, Bank of England and
Filip Zikes, Federal Reserve Board of Governors

The authors are grateful to Celso Brunetti, David Murphy, Tomohiro Ota, John Tanner and Graham Young as well as participants at Bank of England seminars and the Isaac Newton Institute workshop on Systemic Risk: Models and Mechanisms for helpful comments.

robleh.ali@bankofengland.co.uk
International Directorate, Bank of England, Threadneedle Street, London EC2R 8AH
nicholas.vause@bankofengland.co.uk
International Directorate, Bank of England, Threadneedle Street, London EC2R 8AH
filip.zikes@frb.gov
Filip Zikes, Office of Financial Stability Policy and Research, Board of Governors of the Federal Reserve System, Washington DC 20551, United States

The views expressed in this paper are those of the authors, and are not necessarily those of the Bank of England or the Federal Reserve. This paper was finalised on 8 July 2016.

ISSN 1754–4262

Contents

Systemic risk in derivatives markets: a pilot study using CDS data

Robleh Ali and Nick Vause, Bank of England and Filip Zikes, Federal Reserve Board of Governors[1]

In this paper, we draw on network analysis and a sample of derivatives data from a trade repository to demonstrate how the systemic importance of derivatives market participants may be measured. As trade repository data become more comprehensively available to authorities, the same measures could be applied more broadly. We consider the importance of market participants both to the smooth functioning of derivatives markets and in terms of their potential contribution to financial distress. In relation to market functioning, we study some measures that take into account only immediate counterparty positions and others that consider the whole counterparty network of positions. In some cases, the network of positions beyond immediate counterparties makes a significant difference to the rank ordering of the systemic importance of institutions. This means it is important for authorities responsible for financial stability to have access to data beyond the counterparty positions of institutions in their own jurisdictions. In relation to financial distress, we highlight the importance of identifying institutions which may contribute to liquidity strains, as increasing collateralisation of counterparty exposures will diminish credit risk but could at times sharply raise demand for liquid assets to post as collateral.

The views expressed in this paper are those of the authors and not necessarily those of the Bank of England, the Federal Reserve or any of their committees.

(1) Filip Zikes worked on the paper while at the Bank of England.

1 Introduction

The goal of macroprudential regulation is to reduce systemic risk. Some indicators of systemic risk such as aggregate leverage in the financial system already exist. These are based on aggregating data from individual institutions. But the financial system is a network built up from the individual transactions undertaken by its participants, and these participants bring more or less risk to the system depending on their activity and position in this network. This implies that successful macroprudential regulation has to be informed by indicators of systemic risk based on the structure of the financial system, taking into account the interconnectedness and complex relationships between financial institutions (Haldane (2009) and Yellen (2013)). This paper addresses how granular data on derivatives transactions can be used to measure which participants are systemically important to the functioning and stability of a particular derivatives market.

The systemic importance of a financial institution can be calculated in several ways from the same set of data by giving weight to different factors, thereby creating multiple 'views' of the financial system. Each view can help to inform a macroprudential regulator's judgement of which institutions are vulnerable to particular risks and where in the system additional resources may be required. Traditionally, the main focus has been on networks of direct exposures between counterparties arising from unsecured lending positions (Elsinger, Lehar and Summer (2006)). These networks help shed light on the amount of counterparty credit risk in the financial system, the systemic importance of individual institutions, and the likelihood and severity of contagion triggered by a default of a financial institution in the network.

Counterparty credit risk is, however, not the only source of contagion and systemic risk in the financial system. Ongoing regulatory reforms of over-the-counter derivatives markets aim at reducing counterparty credit risk by mandating central clearing and introducing margin requirements for non-centrally cleared transactions (FSB (2013)), but new collateral requirements could also increase liquidity risk. It is therefore equally important to develop a good understanding of the potential future exposures arising from the current network of bilateral derivatives positions, and estimate potential collateral calls associated with changes in the value of these positions. Constructing a network of potential future collateral flows may help regulators identify financial institutions vulnerable to liquidity risk and respond early to potential liquidity stresses.

Additionally, systemic vulnerabilities may also exist when some institutions are exposed to common shocks due to overlapping asset holdings, giving rise to a network of 'indirect' exposures. Although the zero-sum game nature of derivatives implies that not all market participants can be exposed symmetrically to the same shocks, there may at times exist subsets of counterparties with highly correlated books. These common exposures to market and liquidity risk thus need to be carefully assessed.

In this paper we examine all of these different views of the financial network and the associated credit, liquidity and market risks. We use granular data on the UK credit default swap (CDS) market to reconstruct monthly snapshots of market participants' derivatives holdings between 2009 and 2011. We focus on UK single-name CDS because the Bank of England's regulatory mandate allows us to obtain the full set of transactions occurring in this market. We start by studying the network defined by the gross notional amount of outstanding bilateral CDS contracts. This measure shows the largest market participants, typically the main dealers. We then look at the net notional positions (bilaterally netted positions of all CDS transactions), which measure the importance of institutions in facilitating risk transfer. These institutions bear corresponding amounts of risk, so judgements should be formed as to whether they are able to. These judgements need to be supported by other supervisory data such as capital and liquidity buffers, but we show how network analysis can add a previously unavailable layer of information. We then mark to market the value of positions using CDS quotes and examine the resulting network of exposures. If a market participant defaulted, any counterparties exposed to this institution could incur losses on these exposures. Finally, we examine potential future exposures and simulate collateral flows resulting from changes in the values of these exposures.

In summary, we find that both gross and net notional values of credit protection correlate highly with certain measures of systemic importance, notably 'eigenvector centrality'. This is helpful where authorities do not have access to data on the whole financial network but only to the bilateral positions of institutions under their jurisdiction. In other cases, notably 'betweenness centrality', however, gross and net notional values of credit protection correlate less highly with measures of systemic importance. It is therefore important to be able to distinguish the reasons why an institution can be systemically important. Broadly speaking these can be put in two categories: systemically important in relation to market functioning; and systemically important in relation to how much damage an institution can inflict on others. This means an institution could be deemed systemically important because it intermediates a large number of trades and thus helps to keep the market functioning smoothly or because the positions it maintains impose credit risk or potential liquidity strains on its counterparties.

Clearly, the data we analyse in this paper represent a relatively small subset of the total derivatives market and derivatives markets are only a subset of broader financial markets. As trade repository data become more widely and readily available, however, we could apply the methods we have developed for CDS trades to construct network views across the full range of derivatives. We therefore consider the methods we have developed for this paper to have potentially much wider application, especially as the availability of derivatives data increases with the advent of the European Market Infrastructure Regulation reporting requirement and as international authorities seek to share their respective data.

Over time, we will also be able to add information on collateral backing the derivatives exposures and other supervisory data to increase the sophistication of the systemic stress tests we run. How a shock propagates around the financial system also depends on additional exposures such as secured and unsecured interbank lending, as well as the size of banks' capital and liquidity buffers. Our long-term objective is to create an increasingly sophisticated picture of the financial system which can be used to inform the judgement of macroprudential regulators.

The rest of the paper is organised as follows. Section 2 discusses related literature. Section 3 describes the data we have used and reports some summary statistics for the CDS market during our sample period. Section 4 introduces the relevant methodology and concepts that we apply to the data. Section 5 reports our empirical results. Section 6 concludes with a summary and suggestions for future work. Technical details are contained in the appendix.

2 Related work

Our paper is broadly related to relatively recent theoretical and empirical work on financial networks and systemic risk.(1) The seminal contributions of Allen and Gale (2000) and Freixas, Parigi and Rochet (2000) show that a more interconnected financial system is more robust since it allows individual banks to better diversify risks. Others, including Vivier-Lirimonty (2006) and Blume *et al* (2011) find the opposite, that higher interconnectedness increases the likelihood of contagion. Recently, Acemoglu, Ozdaglar and Tahbaz-Salehi (2013) try to reconcile this conflicting evidence by building a framework in which financial networks may be 'robust yet fragile' depending on the scale of the shock. They find that while more-connected networks may be better able to cope with small shocks, highly connected networks may be more prone to contagion when hit by a large shock.

Moving beyond the arguably stylised theoretical models discussed above, Gai, Haldane and Kapadia (2011) and Cont, Moussa and Santos (2013), among others, study how contagion propagates in networks where financial institutions are heterogeneous in terms of their size and number of counterparties. Allen, Babus, and Carletti (2012) compare networks in which institutions have either similar or dissimilar portfolios, studying both liquidity and solvency risks in these networks.

The empirical literature closest to our work focuses on the structure and stability of networks of CDS positions. This includes Brunnermeier, Clerk and Scheicher (2013), Clerc *et al* (2013) and Peltonen, Scheicher and Vuillemey (2013), which study the network of CDS positions using transactional data for the European market. While the analysis in these papers covers a larger segment of the CDS market than our study, they focus on networks of bilateral positions expressed in terms of notional amounts, so only some forms of systemic risk are studied. Using similar data, Duffie, Scheicher and Vuillemey (2014) estimate the impact on collateral demand of new margin and clearing practices and regulations.

Besides the papers that use data on bilateral positions in derivatives, several studies employ balance sheet data together with various network reconstruction methods to estimate the bilateral positions and then perform some form of systemic risk analysis. Giansante, Markose and Shaghaghi (2012) reconstruct the US CDS market network and identify systemically important institutions; the Macroeconomic Assessment Group on Derivatives (2013) looks at the whole derivatives market using similar methodology, as does Markose (2012). Peltonen and Vuillemey (2013) develop a stress-test model for the sovereign CDS and bond market and apply it to a group of major European banks.

Our paper is also related to a number of recent papers that infer the systemic importance of financial institutions from market prices of equity, debt or derivatives, rather than from direct interbank exposures data. Bisias *et al* (2012) provide a survey of the existing systemic risk analytics and their various applications. In the context of credit derivatives, Oh and Patton (2013) propose a dynamic time-series model for large cross-sections of credit spreads and measure systemic risk as the joint probability of distress implied by the model. Billio *et al* (2014) use econometric methods to construct networks of linkages between financial institutions implied by credit spreads and study the network properties over time.

Finally, our approach to monitoring systemic financial risk is closely related to the 10-by-10-by-10 approach of Duffie (2011). In this paper, Duffie suggests studying the impact of ten stress scenarios on ten core financial intermediaries and, in each case, identifying the ten most

(1) For a more comprehensive review of the literature on networks in finance in general, see Allen and Babus (2009).

Chart 1 Summary statistics on the UK single-name CDS market

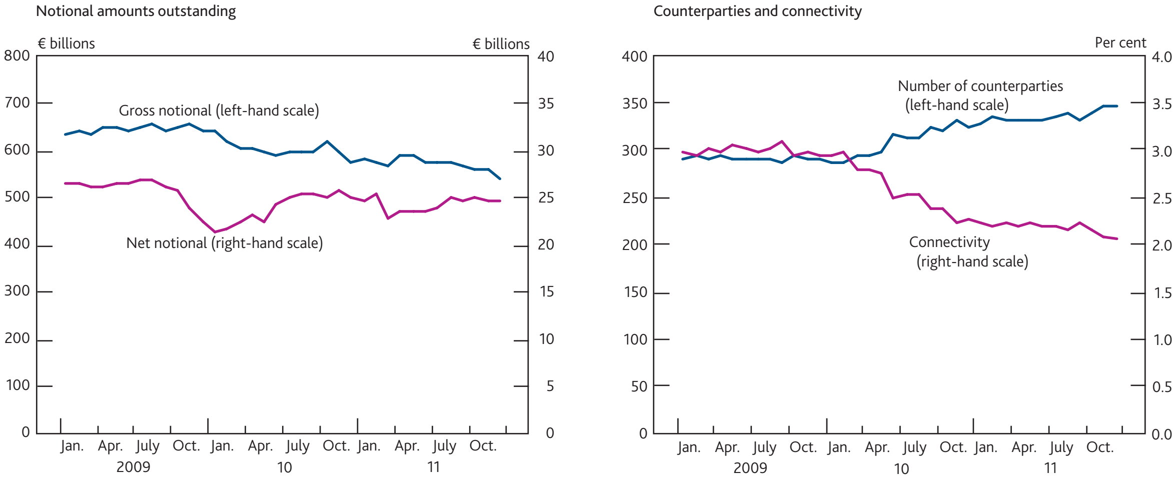

Sources: Depository Trust & Clearing Corporation and Bank calculations.

important counterparties with whom the core intermediaries make profits and losses. Although we do not consider ten stressful scenarios, we do stress test the financial network, with a particular focus on the largest 16 dealers (who we will refer to as the 'G16 dealers'), who are core intermediaries in the over-the-counter (OTC) derivatives markets.

3 Data description and preliminaries

This section describes the two sources of data used in the rest of the paper and provides a brief overview of the size and evolution of the UK CDS market between 2009 and 2011 in terms of notional amounts and market values outstanding.

3.1 Transactions data

Our main source of data is the Trade Information Warehouse (TIW) of the Depository Trust & Clearing Corporation (DTCC) from which we obtain all transaction records for selected UK single-name CDS contracts for the period 2009–11. These records include so-called price forming transactions, such as new trades, terminations and assignments, as well as non-price-forming transactions such as those resulting from trade compressions, delta-neutral auctions and novations to central counterparties. Benos, Wetherilt and Zikes (2013) provide a thorough description of these transaction reports and investigate the structure and dynamics of trading in the UK single-name CDS market in detail.

In this paper, we use the transactional reports from DTCC TIW to reconstruct monthly snapshots of outstanding CDS positions starting in January 2009 and ending in December 2011. We focus on senior CDS contracts denominated in euros as these constitute the vast majority of trading in UK single-name CDS, and restrict our sample to the largest 66 reference entities. This choice is determined by the availability of good-quality CDS quotes that we use for marking to market the CDS positions.

The left panel of **Chart 1** shows the gross and net notional amounts outstanding in our sample of UK single-name CDS. The gross notional amount decreased from €640 billion to €540 billion between 2009 and 2011, while the net notional amount outstanding decreased from €26.5 billion to €24.5 billion. The ratio of gross to net notional was quite stable during the same period, fluctuating between 3.5% and 4.5%.

In the right panel of **Chart 1** we plot the time series of the number of counterparties with outstanding CDS positions and the connectivity of the network of bilateral CDS positions. The latter is defined as the ratio of links between counterparties to the total number of possible links. This metric does not take into account the size of the positions. We find that during our sample period the number of counterparties in the network increase from around 300 to 350. At the same time the connectivity of the network dropped roughly from 3% to 2%. Thus, while the network grew in terms of the number of counterparties, it became relatively sparser.

Finally, in **Chart 2**, we plot the so-called degree distribution. This is defined as the distribution of the number of links that the nodes in the network possess. We calculate for each node the number of counterparties from which the node bought CDS protection (in-degree) and the number of counterparties to which the node sold CDS protection (out-degree). The chart shows that the network is roughly scale-free. That is, its degree distribution follows a power law: most of the parties in the network have a small number of counterparties but a few have a large number of trading relationships. Scale-free networks would generate linear plots in **Chart 2**. The few parties with many trading relationships are the major dealers who intermediate the vast majority of trading in CDS (Benos, Wetherilt and Zikes 2013).

Chart 2 Degree distribution in the UK single-name CDS market (as of June 2009)

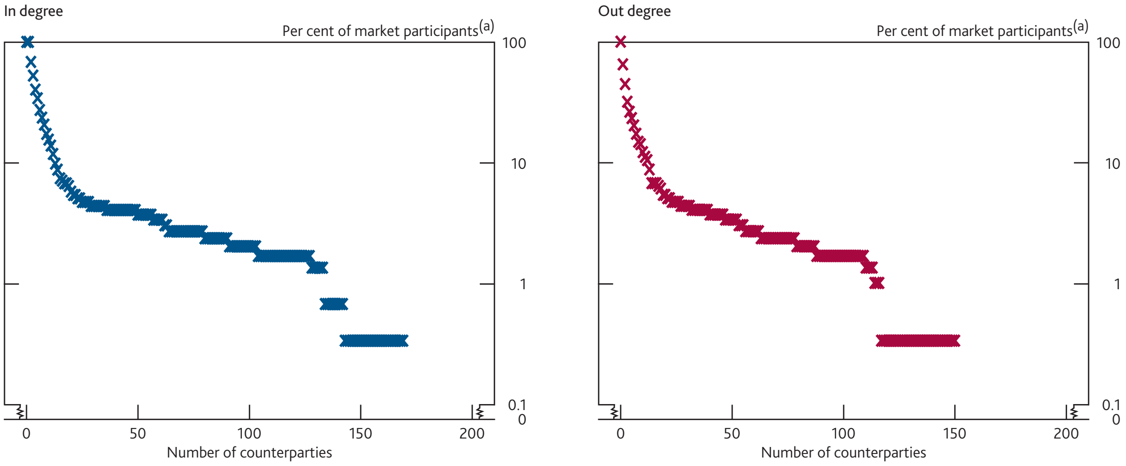

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) Logarithmic scale.

Chart 3 Market value of UK single-name CDS and its drivers

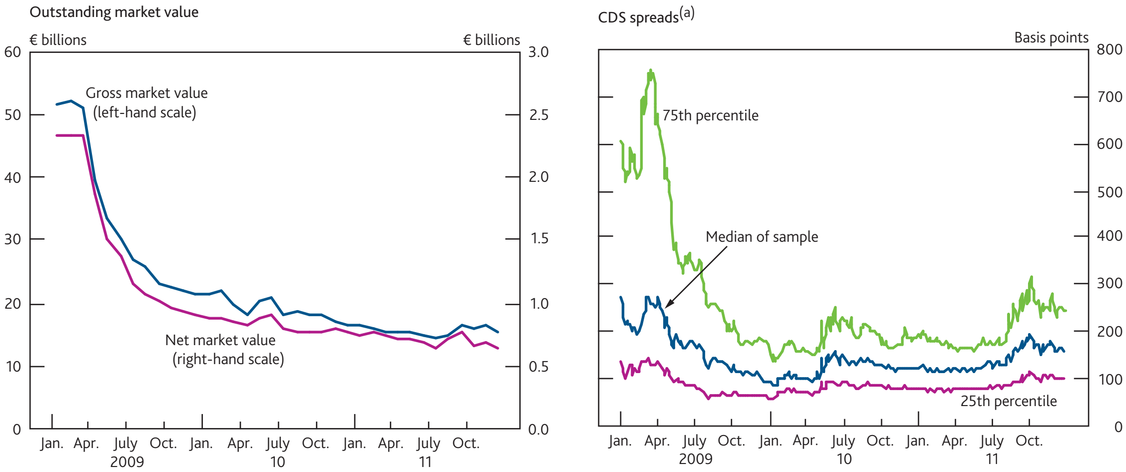

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

(a) End-of-day five-year par CDS spreads.

3.2 Quote data

We augment the CDS trade-level positions data with end-of-day par CDS spreads obtained from Markit. This allows us to mark to market the positions at any given point in time (see the appendix for details) and to simulate forward the evolution of the mark-to-market values.

Chart 3 shows the gross and net market values of outstanding CDS contracts in our sample. These values declined dramatically between 2009 and 2011. As the right panel of the chart indicates, the decline in market values was largely driven by a market-wide contraction in CDS premia during this period. As for notional amounts, the ratio of net to gross market value remained relatively stable over time, fluctuating between 4% and 5%. This reflects a large and fairly constant proportion of contracts having offsetting positive and negative values for their holders.

4 Networks and systemic risk measurement

Network theory has its roots outside finance but has developed useful measures of systemic importance which can be adapted to our purposes. Newman (2010) sets out many different methods in a textbook treatment. Not all of these methods are suitable in the context of financial networks and we therefore focus here only on those that have a straightforward and intuitive economic interpretation.

The simplest metric is based on link counting and is known as degree centrality. This gives a very basic measure of systemic importance derived from the network structure but applies no weight to the links, only registering whether a link exists between two nodes at all. Degree centrality is an important concept because it is the foundation for other methods which differ according to how they weight links. Even in a system where all links are fundamentally identical (for example on the web) it is possible to weight links according to the systemic importance of the node they connect to, meaning that a link from a more systemically important node will carry more weight.

An approach taken on the web is to assess the systemic importance of a node by reference to the importance of the nodes it is connected to. In this case, it is possible to become systemically important either by having links to a large number of nodes or by having a smaller number of links to nodes which themselves are systemically important. This is captured by *eigenvector centrality*. The most widely used application of this concept is Google's PageRank which determines the value of links according to where they came from. The PageRank of a webpage reflects the reality that one link from the BBC's website will likely generate more traffic than many links from a large number of small blogs.

Link direction is a key concept for certain networks. Taking the PageRank example, what matters is that an important website links to you, not the reverse. In derivatives networks, direction is determined by whether an entity is a buyer or a seller. Different views of the network can be derived from the same data set depending on whether sales or purchases are of interest.

The concepts of degree and eigenvector centrality are easily extended to weighted directed networks, which we focus on in this paper. In a derivatives network, a pair of nodes can have two links, one for long positions and one for short positions, with the links weighted according to the size of these positions. For each pair of nodes, we define the *net bilateral position* as the difference between the short and long positions outstanding between the two nodes. Summing up for a given node all the long positions is then the weighted-network equivalent to in-degree and is called *in-strength*, while summing up all the net short positions is equivalent to out-degree and is called *out-strength*. The difference between in-strength and out-strength is the so-called *net strength* and represents the multilaterally netted position of a given node (Brunnermeier, Clerk and Scheicher (2013)).

Turning to eigenvector centrality in directed weighted networks, note that in this case two eigenvectors exist — left and right — and hence two measure of eigenvector centrality can be constructed. In the context of CDS, for example, right eigenvector centrality could relate to credit protection purchases and left eigenvector centrality could relate to credit protection sales. *Right eigenvector centrality* identifies nodes that have a lot of important links pointing towards them (large in-strength) from nodes that themselves have a lot of important links pointing towards them and so on. In a network of exposures, for example, nodes that are central in the right eigenvector sense are those nodes that suffer large losses in a systemic stress. Correspondingly, *left eigenvector centrality* measures the extent to which nodes are pointing towards a lot of important nodes that themselves point to a lot of important nodes and so on (large out-strength). Left eigenvector centrality therefore measures the extent of damage that the failure of a node inflicts on the rest of the network.

But the net strength and eigenvector centralities of a node are not the only factors to take into consideration when considering the stability of a network. A node may have low net strength and even be relatively non-central in the eigenvector sense, but it could still have a great deal of business flow through it. This helps to connect the ultimate buyers and sellers in the market, which makes the node important to the functioning and liquidity of that market. This feature is captured by *betweenness centrality* and is calculated as the fraction of shortest paths between all pairs of nodes in the network that pass through the node of interest. Newman (2005) extends this basic betweenness measure by considering essentially all paths, though he assigns higher weight to the shorter ones. As his measure is based on random walks from one node to the other, he calls his measure random *walk betweenness*. Newman's measure is particularly appealing in the context of over-the-counter markets, where contracts pass through the inter-dealer network several times before finding its ultimate counterparties, and there is no reason to assume that these intermediation chains necessarily run along shortest paths.(1)

Chart 4 illustrates the concepts of systemic importance for a stylised network of banks. In this network, bank 10 has the highest number of links with other banks in the network, and hence the highest degree centrality. Although bank 4 has a lower number of connections than bank 10, it turns out to be connected to banks that are themselves connected to important banks in the network. This bank has the highest eigenvector centrality score. Finally, bank 8 gets the highest betweenness centrality score as it lies on many paths connecting different pairs of nodes in the network.

Given the various centrality scores of individual nodes, we can further examine the network by using the Gini coefficient approach. In this approach, the Gini coefficient is applied to the distribution of centrality scores and captures the extent to which a network is susceptible to dislocation if a small number of large nodes failed. The distribution of the centrality scores for all the nodes in the network can give a useful insight into

(1) Indeed, new contracts form between counterparties with fixed probabilities that reflect outstanding contract volumes along the alternative links from each node. One might think of these probabilities as reflecting the relative likelihoods of potential counterparties offering the best price at the moment of trading. However, we were not able to compute this centrality measure, as it applies only to networks in which every institution buys and sells at least some strictly positive amount. In the UK CDS network, some institutions had only sold protection and others had only bought it, and the random walks through the network got stuck at these nodes.

Chart 4 Illustration of centrality measures

A hypothetical network

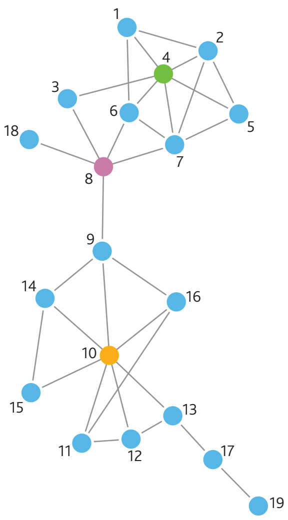

Centrality measures

Node	Degree	Eigenvector	Betweenness
1	3	0.289	0.165
2	4	0.363	0.185
3	2	0.181	0.203
4	6	0.468	0.285
5	3	0.307	0.153
6	4	0.355	0.26
7	5	0.43	0.288
8	4	0.274	0.609
9	5	0.131	0.375
10	7	0.117	0.466
11	3	0.06	0.198
12	3	0.054	0.257
13	3	0.044	0.435
14	3	0.071	0.21
15	2	0.046	0.146
16	3	0.075	0.207
17	2	0.011	0.389
18	1	0.032	0.105
19	1	0.003	0.224

the structure of the network as a whole. For example, if betweenness centrality scores are roughly equal this means that the network is not reliant on a small number of nodes for intermediation and losing a few large nodes would not cause it to dislocate. Conversely if there is an unequal distribution of betweenness centrality scores this suggests that the network would suffer greater disruption if the most systemically important nodes were to fail.

5 Results

Our results are derived from a number of different networks related to counterparty positions in UK CDS. Sections 5.1 and 5.2 respectively study networks of gross and net credit protection bought and sold, where the latter allows bought and sold positions in the same contracts with the same counterparties to offset. Sections 5.3 to 5.5 then study networks of exposures, which relate to the market values of positions between counterparties. These focus respectively on current exposures, changes in exposures and potential future changes in exposures.

5.1 Gross notional amounts

The network of gross outstanding credit protection sales and purchases reflects much about the way the CDS market functions. Hence, analysis of this network can help to identify the participants in the market who are most crucial to its smooth operation.

Our analysis begins with a map of this network for the 40 market participants holding the largest gross outstanding volumes of credit protection bought and sold (**Chart 5**). This reveals a highly interconnected core of dealers and a periphery of end users who hold positions with dealers in the core. For simplicity, the map does not include end users with smaller positions. These institutions also hold their positions with

Chart 5 Gross credit protection on UK reference entities bought and sold by the largest 40 market participants[(a)]

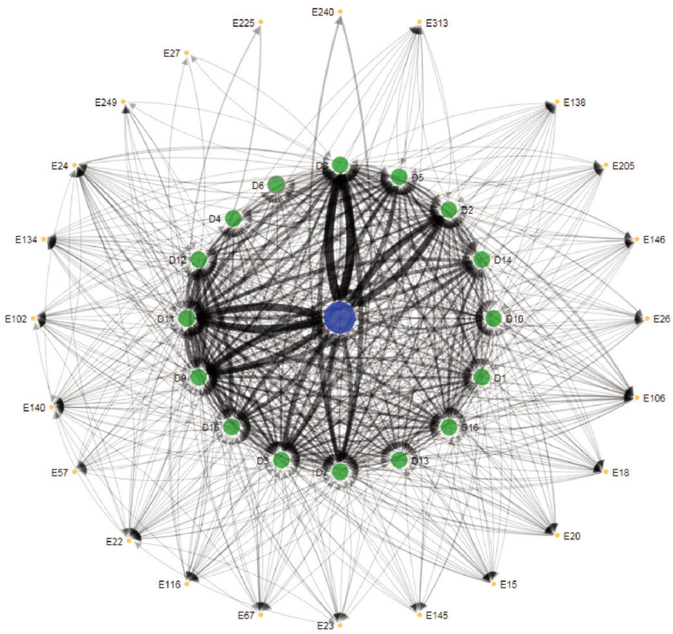

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) CCP coloured in purple, dealers in green and end users in orange. Arrows point to buyers of protection.

dealers, but typically a relatively small number of them. Note that to ensure anonymity of the counterparties in the chart, the date of the data is not reported and the nodes have been drawn with common sizes, rather than reflecting the magnitudes of outstanding positions.

This market structure is also reflected in **Chart 6**. It shows each dealer in the network holding protection-bought and protection-sold positions with virtually every other dealer, while only around one fifth of the potential linkages between dealers and end users were actually used.

Chart 6 Ratio of the number of actual links to possible links in the network of gross credit protection traded on UK reference entities[(a)]

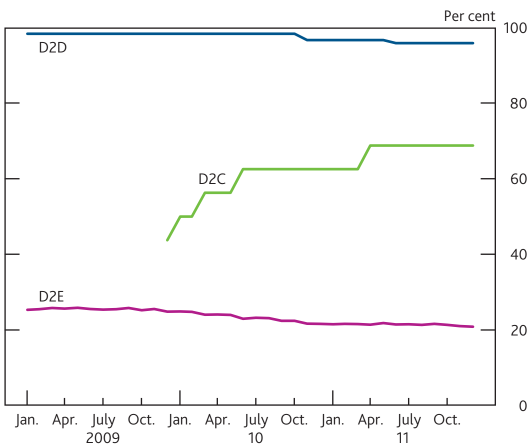

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) Broken down into dealer-to-dealer (D2D) links, dealer-to-end-user (D2E) links and dealer-to-CCP (D2C) links.

Table A Gross purchases and sales of credit protection on UK reference entities by individual market participants and their centrality in this network (as of end-2011)

Party	Gross bought (€ billion)	Party	Gross sold (€ billion)	Party	Right eigenvector centrality	Party	Left eigenvector centrality	Party	Betweenness centrality
Top five dealers									
1	62.0	1	58.0	1	0.40	1	0.38	1	0.23
2	47.9	2	53.1	2	0.35	2	0.37	2	0.20
3	47.3	3	50.2	3	0.33	3	0.35	3	0.16
4	35.8	4	37.3	5	0.26	5	0.26	6	0.12
5	35.5	5	35.6	4	0.25	4	0.24	4	0.11
Top five end users									
1	6.5	1	10.2	1	0.08	1	0.08	8	0.16
2	5.2	2	5.1	2	0.05	2	0.04	9	0.07
3	3.3	3	4.0	3	0.04	3	0.04	10	0.06
4	2.7	5	3.8	4	0.04	5	0.03	11	0.06
5	1.9	6	3.3	5	0.03	7	0.03	12	0.04
Dealers	441.7		438.4						
Per cent of total	92.2		91.5						
Non-dealers	37.5		40.8						
Per cent of total	7.8		8.5						
Total	**479.2**		**479.2**						

Sources: Depository Trust & Clearing Corporation and Bank calculations.

This core-periphery structure reflects a key aspect of CDS market functioning whereby dealers intermediate the trading requirements of end users. A dealer may, for example, sell credit protection to a client. In doing so, the dealer would become exposed to the risk that the reference entity may default. However, the dealer may subsequently offset this risk by buying protection in the interdealer market. Another dealer may be keen to sell protection if they held the opposite position, having previously bought credit protection from a client. Alternatively, several dealers may be willing to sell relatively small amounts of credit protection, dispersing the default risk across multiple dealers.

Chart 6 also shows many dealers holding positions with a central counterparty (CCP). Since the 2008 global financial crisis, when dealers incurred losses on bilateral OTC derivatives exposures, they have increasingly cleared their trades with a CCP. This involves replacing contracts between protection buyers and sellers with an equivalent contract between the buyer and a CCP and another equivalent contract between the same CCP and the seller. Regulatory plans to mandate clearing of certain types of OTC derivatives further supported this development. As a result, the proportion of potential links between dealers of UK CDS and the CCP in this market that were populated by actual sales and purchases of credit protection increased from around 40% in late 2009 to over two thirds by the end of 2011.

Given that the CDS market functions through dealers who accommodate client trades and disperse the resulting risk in the interdealer market, we next present some metrics that help to identify the institutions that are most crucial to this functioning (**Table A**).

The first and most straightforward of these indicators are the gross volumes of credit protection bought and sold by individual institutions. The table shows that dealers accounted for over 90% of the market on both of these measures, with the top five dealers accounting for almost half. This reflects the numerous trades in the interdealer market that often follow a trade with an end user as the underlying risk gets dispersed. Conversely, the largest end users had much smaller positions. But, like the dealers, these institutions (which include large non-dealer banks and hedge funds) held positions in both sold protection and bought protection. Although not reported in **Table A**, smaller end users held more unbalanced positions, having either mainly bought or mainly sold credit protection.

These patterns are reflected in left and right eigenvector centrality scores. Institutions get high left eigenvector centrality scores if they have sold a lot of credit protection to high scoring institutions, which have in turn sold a lot of credit protection to high scoring institutions, and so on. Right eigenvector centrality scores are determined in the same way except with purchases of credit protection. As for gross sales and purchases of protection, left and right eigenvector centrality scores are much larger for dealers than for end users, with similar proportional differences between these two types of market participant. This is because the outstanding gross positions of end users are almost completely held with dealers and the majority of dealers' outstanding gross

Chart 7 Rank correlations between indicators of importance of participants in the market for credit protection on UK reference entities(a)

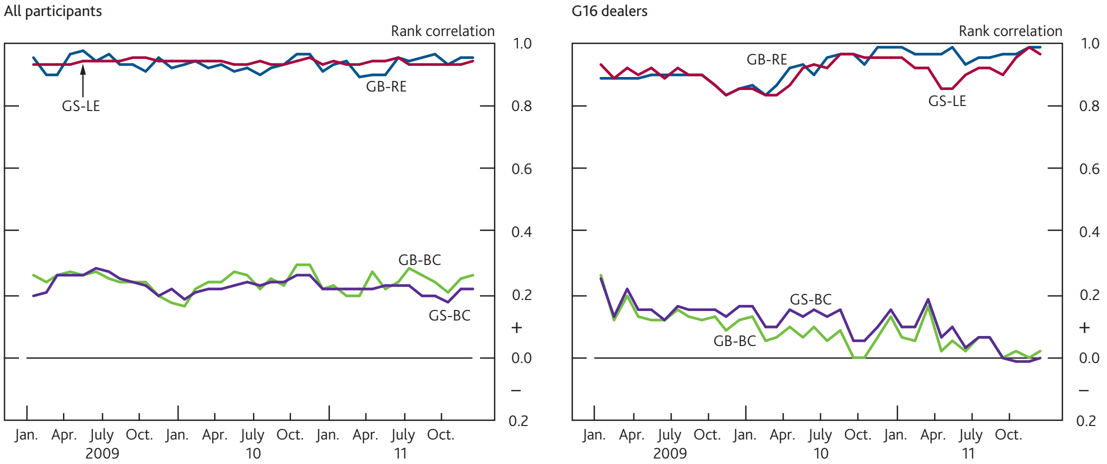

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) The indicators of importance are gross notional amounts bought (GB) and sold (GS) as well as left eigenvector (LE) and right eigenvector (RE) centrality and betweenness centrality (BC) in the network of gross outstanding UK CDS positions.

positions are held with other dealers. Hence, the network structure beyond immediate counterparties does not vary significantly across institutions. Indeed, these connections only modestly affect the rank ordering of institutions compared with those based only on gross credit protection traded with immediate counterparties (**Chart 7**).

This is helpful from one point of view. At present authorities often cannot compute left and right eigenvector centralities as these measures require data on outstanding positions across the whole network. Instead, they can typically only access data on positions involving institutions in their jurisdiction. The eigenvector centrality measures in **Table A** and **Chart 7** are an exception, constructed from the complete network of positions in CDS on UK reference entities, which UK authorities can access. However, by collecting data on gross outstanding credit protection, it seems that authorities can proxy these richer eigenvector centrality metrics.

In contrast, **Table A** and **Chart 7** show that the rank orderings of institutions based on betweenness centrality are quite different to those based on gross outstanding credit protection or eigenvector centrality. Note that there is only one set of betweenness centrality scores (no 'left' and 'right' versions) in these figures since any chain that channels protection sales in one direction must channel protection purchases in the other direction. Nevertheless, betweenness centrality adds valuable information as it identifies institutions that act like a bridge between clusters within a network. For example, if in the global CDS market a cluster of US firms had traded credit protection on US reference entities and a cluster of European firms had traded protection on European reference entities then large international dealers might bridge the two clusters, facilitating the diversification of regional credit risk. Hence, betweenness centrally would ideally be computed as a separate indicator. As this requires data on outstanding positions across the network, it is important that international authorities continue their efforts to combine data from trade repositories across the globe.

While the above helps us to identify the institutions that are most important to the smooth functioning of a particular OTC derivatives market, it remains difficult to judge how the reliance of a market on its key participants varies over time or how this reliance compares across markets. Gini coefficients can help in this regard by measuring the degree to which the importance of participants in a market — whether measured in terms of gross outstanding positions, eigenvector centrality or betweenness centrality — is concentrated. A market in which all participants were of roughly equal importance would have a Gini coefficient close to zero, whereas a market in which a small number of participants had much higher importance scores than the others would have a Gini coefficient close to one. **Chart 8** shows Gini coefficients based on the different measures of importance. Each of these coefficients increased a little over the three-year sample period, suggesting the UK CDS market became slightly more reliant on key participants during that time.

5.2 Net notional amounts

The network of net outstanding credit protection sales and purchases summarises how the CDS market redistributes credit risk. The net notional amount of credit protection bought by one institution from another is simply the difference between the gross amounts bought and sold by the first institution with the second. By taking these offsetting positions into account, net notional amounts show the volumes of risk transfer that have taken place.

Transfer of risk to parties better able to bear it is a valuable economic service. Institutions that provide this service in substantial volume might be considered systemically

Chart 8 Gini coefficients summarising the concentration of outstanding gross positions and centrality scores in the UK CDS market[(a)]

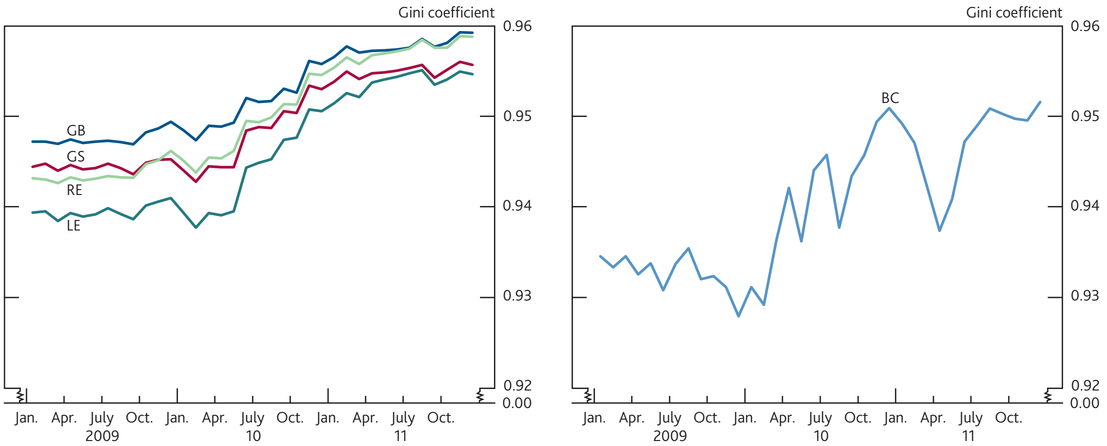

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) Gini coefficients based on gross notional amounts bought (GB) and sold (GS) as well as left eigenvector (LE) and right eigenvector (RE) centrality and betweenness centrality (BC) in the network of gross outstanding UK CDS positions.

Table B Top ten sellers of credit protection on UK financial reference entities and the top ten buyers for each seller (€ billions)[(a)]

Seller	Notional sold	Buyer type and notional bought									
1	3342	1077	427	366	275	262	186	142	135	64	45
2	2726	517	404	330	200	174	114	100	90	88	75
3	2242	391	294	281	175	145	145	102	75	58	38
4	1557	409	151	133	98	74	70	51	32	30	29
5	1251	217	131	90	87	68	60	60	55	47	44
6	1133	216	150	148	145	128	123	122	37	23	10
7	889	281	109	91	85	82	72	28	18	17	14
8	722	236	131	85	37	34	28	24	20	19	15
9	547	180	80	66	50	39	28	28	24	17	12
10	518	125	118	64	51	29	26	15	14	11	10

Sources: Depository Trust & Clearing Corporation and Bank calculations.

(a) Excluding the CCP. Buyers are identified either as a dealer (blue) or an end user (magenta).

important. At the same time, however, these same institutions could be a source of systemic risk. Were such an institution to fail, its counterparties may want to replace their positions with new counterparties. If the majority of the resulting trades were purchases of credit protection, this could drive up CDS premiums, and *vice versa* for protection sales, meaning that positions could only be replaced at a loss. These potential losses grow with the size of the positions being replaced and the price impact of trading, which is higher in less liquid markets.

By way of illustration, and taking a similar approach to that suggested by Duffie (2011), **Table B** shows the net sales of credit protection on UK financial reference entities for the ten market participants with the largest positions as of a particular date in our sample period (we withhold this date to help ensure anonymity of the counterparties). The major dealers dominate this list. For each protection-sold position, the table also shows the corresponding protection-bought positions of the ten largest counterparties (anonymised, other than their type). The larger these positions, the greater was the risk of loss should the protection seller have failed and the buyers sought to re-establish insurance against the default of UK financial institutions by purchasing protection from new counterparties. The protection buyers and the size of their positions varied significantly across the ten protection sellers, with some end users featuring alongside the major dealers. It is very helpful for authorities to be able to identify institutions vulnerable to replacement risk, so they can work to mitigate it.

Data like that in **Table B** also sheds some light on the losses that might result from reference entity defaults. Certainly, data on the top protection sellers (in the first column of the table) identifies the institutions that would incur the largest losses. But the size of these potential losses can be quite uncertain and vary substantially with the reference entity. For instance, losses were 91.5% of the notional amount following the default of Lehman Brothers, but only 8.5% for Fannie Mae. If these losses overwhelmed a protection seller, causing its failure, losses may then spread to the protection buyers (in the rows of the table) through non-payment of receivables. However, this would depend on the value of other OTC

Chart 9 Eigenvector centrality scores of G16 dealers in the network of outstanding sales and purchases of credit protection on UK reference entities

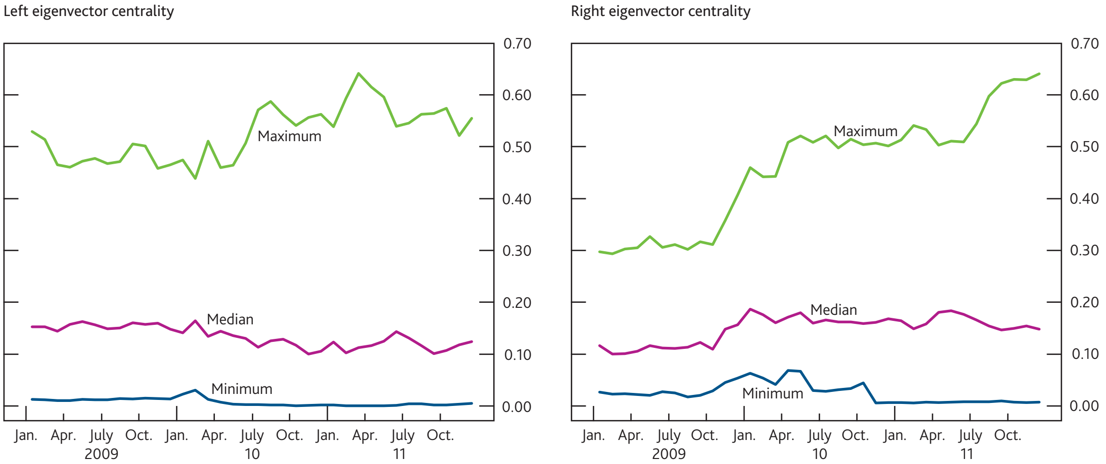

Sources: Depository Trust & Clearing Corporation and Bank calculations.

derivatives held between the two counterparties, presuming they had signed a close-out netting agreement, as this would allow contracts with positive and negative market values to be offset.

As with gross notional amounts, there is reason to look beyond immediate counterparties with net notional amounts. A protection seller may be regarded as more critical to the market if it supplies insurance to institutions that, in turn, supply insurance to other important counterparties. This is captured by right eigenvector centrality, based now on an adjacency matrix of net (rather than gross) outstanding credit protection sales and purchases. An adjacency matrix is a mathematical representation of a network. In this case, each element of that matrix represents a bilateral net position in credit protection. Conversely, institutions with high left eigenvector centrality scores buy insurance from counterparties that, in turn, buy substantial volumes of insurance from multiple counterparties. Hence, institutions may be identified as systemically important in providing or taking insurance by their right and left eigenvector centrality scores respectively. **Chart 9** shows that, even within the group of G16 dealers, some of these institutions are much more important in this sense than others. It also shows that in the three years to end-2011 the most important provider of insurance became even more important.

Net sales and purchases of credit protection correlate highly with their respective eigenvector counterparts, as was the case for gross notional amounts. When focusing only on the G16 dealers, these correlations are a little lower than the equivalents for gross notional amounts, although they are still quite high in absolute terms. This can be seen in **Table C**, which shows overlaps between the top five net buyers of credit protection and right eigenvector centrality scorers as well as between the top five net sellers of credit protection and left eigenvector centrality scorers, for both dealers and end users. Across all counterparties, the corresponding rank correlations were 0.95 and 0.86 respectively. These high rank correlations suggest that net positions with direct counterparties are good proxies for their more comprehensive eigenvector centrality counterparts. Again, this may be of some comfort to authorities who can only access data on the positions of firms in their jurisdictions.

Table C Net purchases and sales of credit protection on UK reference entities by individual market participants and their left and right eigenvector centrality

Party	Gross bought (€ billion)	Party	Right eigenvector centrality	Party	Gross sold (€ billion)	Party	Left eigenvector centrality
Top five dealers							
1	6.4	4	0.64	6	8.0	6	0.55
2	4.7	2	0.36	3	5.9	8	0.44
3	4.4	10	0.33	7	5.0	7	0.39
4	4.2	1	0.24	8	4.2	12	0.26
5	3.8	11	0.22	9	3.2	9	0.20
Top five end users							
1	1.8	1	0.17	6	0.8	6	0.19
2	0.7	4	0.06	3	0.6	3	0.07
3	0.6	2	0.06	7	0.5	9	0.07
4	0.5	10	0.05	8	0.5	7	0.06
5	0.4	11	0.04	9	0.4	12	0.05
Dealers	45.8				42.9		
Per cent of total	80.3				75.1		
Non-dealers	11.3				14.2		
Per cent of total	19.7				24.9		
Total	57.1				57.1		

Sources: Depository Trust & Clearing Corporation and Bank calculations.

5.3 Current exposures

Moving away from positions, we turn now to study networks of counterparty exposures, beginning with current exposures. The current exposure of one institution to another is the net market value of its positions with that counterparty if this is positive and zero otherwise. It is the value currently at risk should that counterparty default. We focus on net current market values, which offset contracts with positive market value against those with negative value, rather than gross market values. This is because bilateral counterparties almost always trade derivatives under close-out netting agreements, which allow them to offset positive and negative market values of different contracts in the event of default.

Table D shows the institutions which held positions in credit protection on UK reference entities across all of their counterparties that were the most in the money (ITM) or out of the money (OTM). In the absence of collateral collected against exposures, institutions with high ITM positions would have been most at risk of loss from counterparty defaults overall. However, even institutions which were OTM on a multilateral basis could still have had a significant ITM position with a particular counterparty. A map of the network of bilateral ITM and OTM positions, such as that depicted in **Chart 10**, helps to identify any such positions.

Table D Net market values of positions in credit protection on UK reference entities (€ millions)

In the money		Out of the money	
Party	Value	Party	Gross
Top five dealers			
1	300	6	473
2	106	7	230
3	96	8	181
4	85	9	77
5	79	10	49
Top five end users			
1	171	6	24
2	72	7	17
3	40	8	17
4	32	9	13
5	28	10	12
Dealers	714		1079
Per cent of total	52.1		78.6
Non-dealers	658		293
Per cent of total	47.9		21.4
Total	**1372**		**1372**

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

For example, while the market value of D11's positions across all of its counterparties was negative, **Chart 10** reveals that it had a position with D8 that was significantly in the money from its point of view. If D8 had defaulted on this exposure, D11 could have lost up to its full value. This would have depended first on whether D11 had collected any collateral from D8 against the exposure. If not, D11 would have had an unsecured claim on D8 for the net market value of the derivatives they had traded. In the worst case, this claim could have returned nothing. D11 may then have defaulted, with similar consequences for its counterparties.

Chart 10 Multilateral and bilateral net market values of positions in credit protection on UK reference entities[(a)]

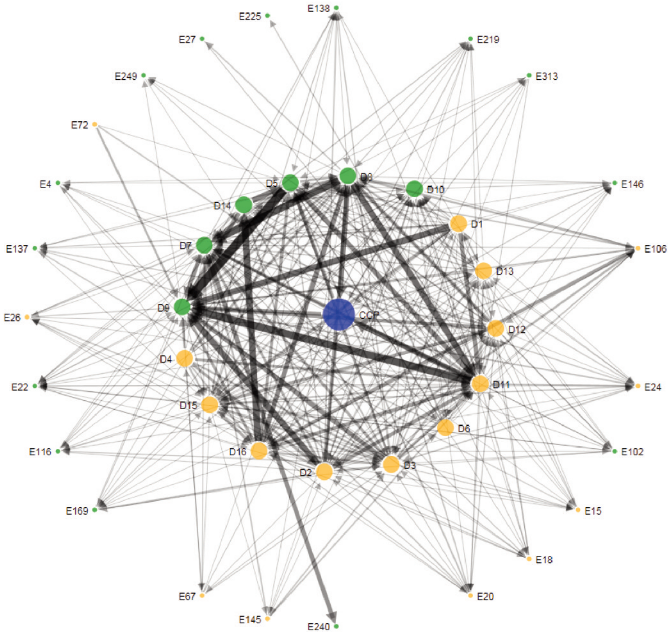

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

(a) In addition to the CCP, market participants are identified either as a dealer (D) or an end user (E). For these market participants, multilateral positions with positive net market value are coloured in green, while those with negative net market value are coloured in orange. Arrows reflect bilateral positions and point from the out-of-the-money counterparty to the in-the-money counterparty.

Credit losses could even spread without defaults. If D8 had not defaulted, but its credit quality had deteriorated, D11 should still have marked down the value it ascribed to its positions with D8 to reflect a higher probability that the receivables associated with those positions would not in fact be collected. This loss would, in turn, have weakened D11's credit quality, which may have prompted similar credit valuation adjustments (CVA) by its counterparties.

Despite the potential for credit losses to spread through a network of derivatives exposures, neither eigenvector nor betweenness centrality applied to an adjacency matrix formed from exposures would necessarily give a clear indication of the potential for individual institutions to act as a source of such contagion. The scope for losses to spread through exposures depends on the degree to which they are backed by collateral and on the volume of loss-absorbing capital held by exposed counterparty. Instead, an approach like that of Furfine (2003) which simulates the spread of losses through a network of post-collateral exposures, taking into account the capital held by institutions, is necessary.

5.4 Potential changes in current exposures

Potential changes in current exposures could represent a concern for financial stability authorities for two reasons. First, the change in exposures could reflect one or more systemically important institutions (as identified above) incurring mark-to-market (MTM) losses on their OTC derivative positions. Where these positions are not hedging

Chart 11 Probability distributions of selected G16 dealers' one-day profits and losses

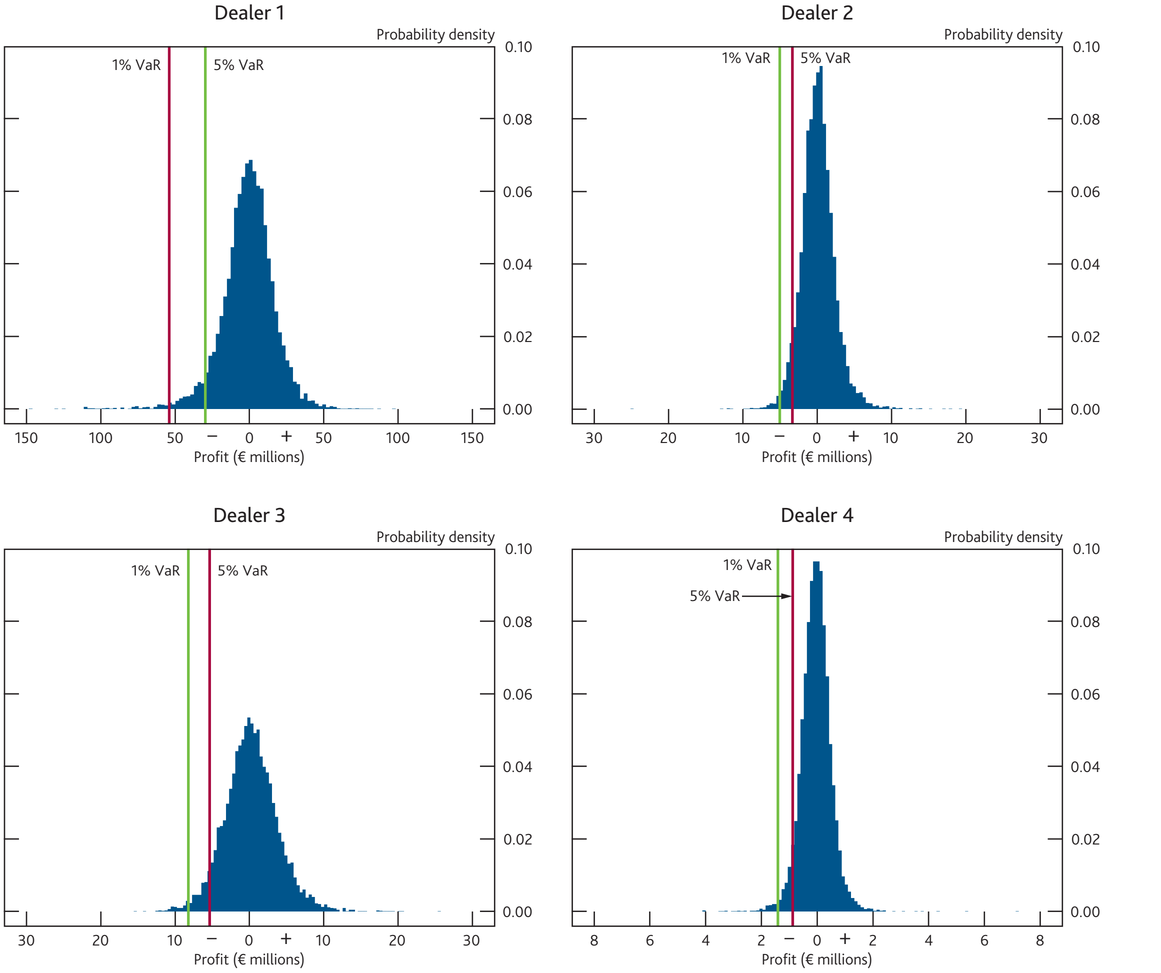

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

securities holdings or other exposures, such losses would erode the institution's capital, making default or CVA cascades more likely. Second, and regardless of whether or not OTC derivative positions were used for hedging, changes in MTM values could lead to liquidity strains where variation margins have to be posted. Variation margin is paid — usually on a daily basis — by the counterparty whose position incurs a MTM loss to the counterparty whose position makes a MTM gain. The exchange of variation margin reduces current counterparty exposures to zero. Variation margins are already exchanged on centrally cleared positions and they are due to become compulsory for most non-centrally cleared positions, with phased implementation beginning in September 2016.

We estimate potential changes in current exposures by modelling the comovement of the main drivers of CDS spreads. In particular, for each of the 66 UK CDS in our study, we model factors that drive the levels and slopes of their term structures, which relate spreads to maturities. The modelling approach, which is described in detail in the appendix, reflects the fact that large changes in current exposures are more likely when the prevailing level of volatility in CDS spreads is already high.

We use this model to simulate 10,000 potential daily changes in the market values of UK CDS. By combining these simulations with the portfolio holdings of market participants, we estimate the probability distribution of profits and losses that they face. **Chart 11** shows a set of such estimates for some of the G16 dealers as of a particular date in our sample. The 95th and 99th percentiles of these distributions represent (lower bounds of) the amounts that institutions could expect to lose one day in 20 or one day in one 100, respectively. These amounts are known as values at risk (VaR). **Table E** reports estimated VaRs for the G16 dealers around mid-2009, when the volatility of CDS spreads was still elevated following the 2008 global banking crisis, and late-2011, when volatility was closer to normal levels.

Moreover, we can use the model to compute probabilities of multiple institutions incurring large losses at the same time. For instance, **Chart 12** shows the probabilities of any two G16 dealers simultaneously incurring losses in excess of their respective 95th percentile VaRs as of a particular date in our sample. If the probability distributions of profits and losses at individual institutions were independent, the probability of simultaneous VaR breaches would be 0.25%. Some of the

Table E Estimates of G16 dealer's one-day value at risk on their UK CDS portfolios (€ millions)

	Mid-2009		Late-2011	
	95% VaR	99% VaR	95% VaR	99% VaR
D1	30	54	5	8
D2	9	15	7	12
D3	14	24	7	11
D4	3	5	2	3
D5	25	41	3	6
D6	5	8	1	2
D7	5	9	2	3
D8	12	20	2	4
D9	3	6	2	3
D10	3	5	4	7
D11	5	9	3	5
D12	7	13	1	1
D13	9	15	4	7
D14	1	1	1	1
D15	3	5	1	1
D16	36	62	5	9

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

Chart 12 Probabilities of pairs of G16 dealers simultaneously exceeding their 95th percentile value at risk

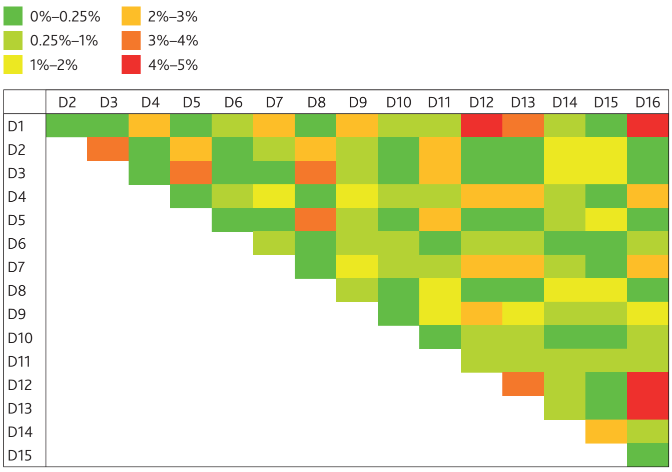

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

values in the table exceed this quite significantly, and are closer to the 5% limit at which VaR breaches at the institutions would always coincide.

While large losses at one or more systemically important institutions would be of significant concern to a financial stability authority, the increasing use of collateral against counterparty exposures at least limits the potential for contagious defaults or credit valuation adjustments. However, this instead gives rise to potential liquidity strains, which could spread between institutions. By way of illustration, **Chart 13** shows estimates of the variation margins that G16 dealers would have been expected to exchange if UK CDS spreads had suddenly increased by 100 basis points as of a particular date in our sample. In this scenario, D9 would have had to pay

Chart 13 Variation margin flows between G16 dealers following a 100 basis point increase in UK CDS spreads[(a)]

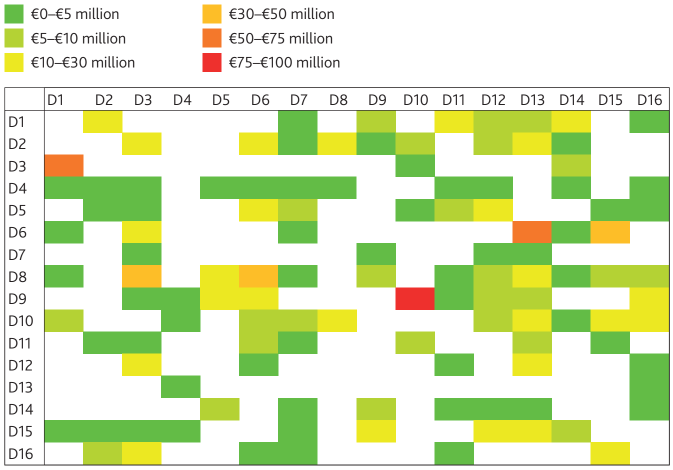

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

(a) Payers of variation margin listed in the rows; receivers in the columns.

€75–€100 million to D10. If D9 had been unable to make this payment, this could have had implications for D10's ability to make variation margin payments to the ten G16 dealers it was obliged to pay, including relatively substantial amounts to D8, D13, D15 and D16.

5.5 Potential changes in exposures/initial margins

Liquidity strains may also increase as a result of more widespread demand for initial margin. This collateral is collected when new derivatives positions are established, to protect against possible adverse movements in the value of the position should it take a number of days to close out following default of the counterparty. Additional initial margin may be called on a fixed portfolio if the potential for that portfolio to generate losses rises. Central counterparties already collect initial margin and new regulations are set to introduce the exchange of similar margins on many bilateral transactions from September 2016. Moreover, while variation margin payments transfer liquid assets from one institution to its counterparty, increases in initial margin requirements boost the demand for liquid assets from both counterparties. Hence, changes in initial margin requirements may be more likely to lead to liquidity strains across the financial system as a whole.

Initial margin requirements often increase when the prevailing level of market volatility rises. By way of illustration, **Chart 14** shows estimates of the initial margins required to cover 99th percentile ten-day losses on a fixed set of bilateral UK CDS portfolios of G16 dealers on two dates: one in May 2010, shortly after sharp credit rating downgrades of Greece heralded the onset of the euro area sovereign debt crisis; and one in March 2010, around two months earlier.[(1)] During this

(1) For CDS, margin models sometimes demand some extra margin to allow for potential reference entity defaults. This issue does not arise with other derivatives. We ignore it here.

Chart 14 Estimates of initial margins required to cover 99th percentile ten-day losses (€ millions)[(a)]

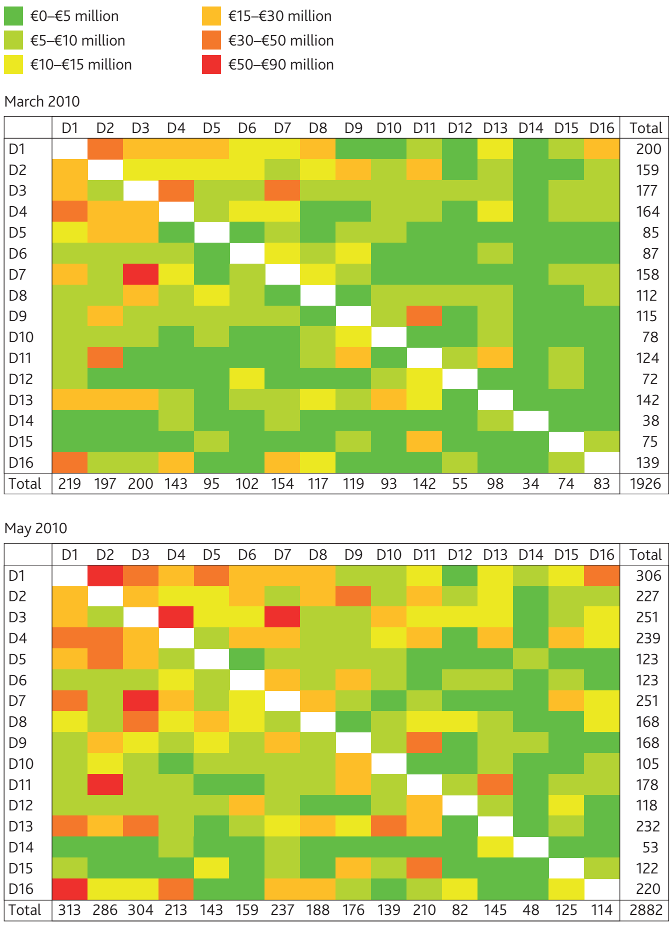

Sources: Depository Trust & Clearing Corporation, Markit and Bank calculations.

(a) Payers of initial margin listed in the rows; receivers in the columns.

short time span, initial margin requirements for G16 dealers' positions in UK CDS increased by around 50% overall. Depending on the contracts traded, changes in bilateral requirements ranged from almost nothing to more than doubling, generally increasing for both counterparties.

6 Conclusion

Authorities responsible for financial stability must decide how to deploy their resources to address the greatest risks. Such decisions are often based both on quantitative (eg data) and qualitative (eg market intelligence) information. The advent of derivatives trade reporting has significantly expanded the quantitative data available to authorities. The aim of this paper is to illustrate how a large granular data set can be processed into useful indicators which can inform regulatory judgements.

The sheer size of the datasets being generated by the derivatives reporting obligation demands an approach which recognises the need for machine processing of raw data into useful maps of the financial system. The work we have done in this paper makes a start by suggesting some measures of systemic importance that can be derived from these maps, using a large granular data set. The granularity of the data is such that we can reconstruct the positions of market participants and investigate how losses and collateral would flow through the system in response to different shocks.

It turns out that for the data we analyse here some network-based measures of systemic importance such as eigenvector centrality do not paint a significantly different picture from more basic indicators that reflect only immediate counterparty positions. However we would not discard eigenvector centrality on this basis. We have used a relatively narrow data set for this paper and it may be that eigenvector centrality adds value relative to the more basic indicators in different markets. Other measures such as betweenness centrality do give an interesting alternative picture which can help inform regulatory judgements.

This means that authorities responsible for financial stability require access to data covering the full network of counterparty positions, and not only the immediate positions of institutions in their jurisdictions. Work is in progress to grant authorities mutual access to trade repositories or to negotiate access for an international data aggregator. Additional work is also in progress to enhance the quality and ease of use of trade repository data, including developing and disseminating global standards for entity, product and trade identification.

This pilot study has shown what can be done with derivatives data in isolation. However, further systemic risk indicators could be developed by combining this data with counterparty positions in other financial instruments and with data on the balance sheets of institutions in the network. For example, to study Furfine-type simulations of credit contagion we would want to track knock-on defaults or CVA losses, evaluating how far these might spread given institutions' capital. Or, to study contagious liquidity strains, we would want to simulate collateral calls and evaluate whether these could be met given institutions' holdings and access to liquid assets.

Appendix
Valuation of CDS positions and modelling of CDS spreads

We mark to market the outstanding CDS positions following a standard methodology outlined in eg Bomfin (2005, Chapter 16). For each reference entity and date in our sample, we recover the risk-neutral survival probabilities from the term structure of end-of-day CDS par spreads obtained from Markit. We use maturities of 1, 2, 3, 5, 7 and 10 years and assume piece-wise constant hazard rates, a recovery rate of 40% and use the euro overnight indexed swap curve for discounting.

To understand the market risk associated with the CDS positions, we propose and estimate a parsimonious dynamic econometric model for the term structure of CDS spreads for the 66 reference entities included in our sample. The model allows for time-varying volatility of the individual CDS curves, as well as for time-varying correlations and tail-dependence among the curves.

Inspired by the work of Oh and Patton (2013), we employ factor copulas to model the time-series of our 66 reference entity CDS curves. The modelling exercise involves the following steps:

1. For each reference entity CDS curve, we extract the first two principle components — level and slope factors.

2. We model the time series of changes in each factor by an AR(2)-GARCH(1,1) model with skewed Student-t innovations. The 132 models are estimated separately by maximum likelihood. Standard diagnostic tests show little structure in the standardised residuals of the models.

3. We then transform the CDS factor changes into uniform variates by the conditional probability integral transform implied by the AR-GARCH model.

4. We model the 66 level/factor uniform variates by a generalised autoregressive score (GAS) Student-t copula, where the correlation matrix has a single-factor structure (six industries). The intercepts in the GAS recursion for the six factor loadings are allowed to be factor-loading specific, but the dynamic parameters are fixed across the six factor loadings. The parameter governing the number of degrees of freedom of the Student-t copula is also allowed to follow GAS dynamics, and has its own dynamic parameters. For more detail on GAS models, see Creal, Koopman and Lucas (2013).

5. We model the 66 slope/factor uniform variates using a static Student-t copula with a single-factor structure correlation matrix, similarly to the dynamic copula for the level factors. The reason for using a static copula is that the dependence between the slope factors is generally weak and exhibits little variation over time.

6. Both copulas are estimated by maximum likelihood.

We again use daily end-of-day CDS par spreads obtained from Markit for maturities 1, 2, 3, 5, 7 and 10 years. Our sample runs from October 2007 to December 2011, yielding 1,095 daily observations.

The main estimation results are summarised in **Chart A1**. In the top-left panel, we plot the median level and slope factors, which are the two factors driving the CDS curves in our model. We see that the level of CDS spreads fluctuated wildly between 2007 and 2011, peaking shortly after the default of Lehman Brothers in September 2008. The spreads also widened considerably after the demise of Bear Stearns in March 2008 and around the beginning of the eurozone sovereign crisis in May 2010. The slope factor, which measures the steepness of the term structure of CDS spreads, tended to be positive during the sample period for a typical reference entity, but flipped sign at the height of the 2008 crisis. The conditional volatility of the level and slope factors follows similar pattern, spiking in times of significant market turmoil.

In the bottom-left panel of **Chart A1**, we plot the median conditional correlation among the levels of the 66 CDS curves, together with the associated cross-sectional minima and maxima. We find that similar to volatility, the conditional correlation tends rise in times of stress. During our sample period, the market-wide increase in correlation is particularly pronounced during the eurozone sovereign crisis beginning in May 2010.

Finally, in the bottom-right panel of the chart, we plot the conditional degrees of freedom of the copula driving the level factors. This parameter determines, together with the correlation coefficients, the degree to which extreme events occur jointly for a number of reference entities rather than individually — the so-called tail dependence. Lower values of the degree-of-freedom parameter imply higher likelihood of joint occurrence of extreme events. We find that the degree-of-freedom parameter fluctuates wildly during our sample period, taking value between five (high tail dependence) and 40 (low tail dependence). Interestingly, the times when the parameter is particularly low are not always characterised by market stress.

Chart A1 Summary statistics of the factor copula fitted to UK CDS premia

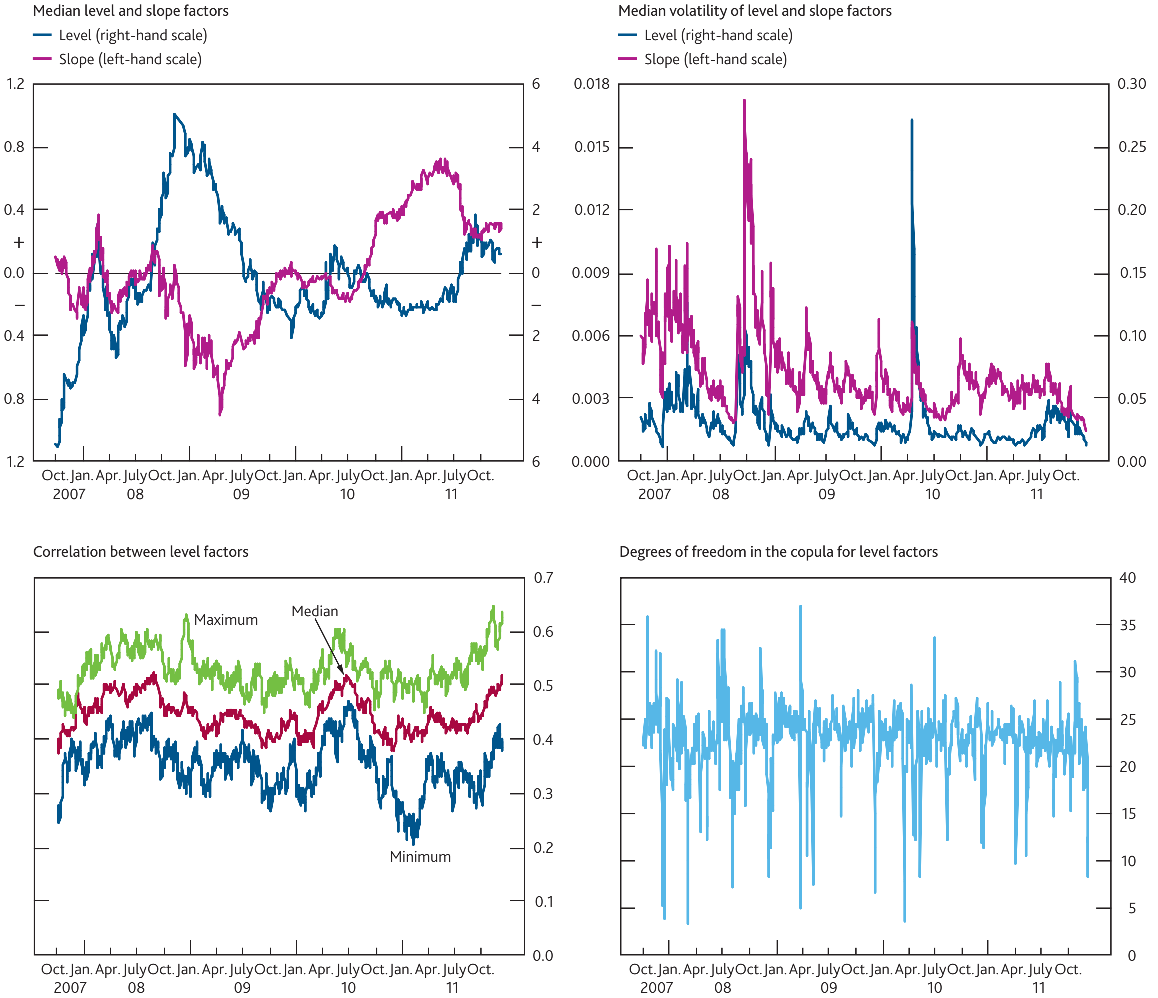

Sources: Markit and Bank calculations.

References

Acemoglu, D, Ozdaglar, A and Tahbaz-Salehi, A (2013), 'Systemic risk and stability in financial networks', *MIT Working Paper*.

Allen, F and Babus, A (2009), 'Networks in finance', in Kleindorfer, P R and Yoram, W (eds), *The network challenge: strategy, profit, and risk in an interlinked world*, Wharton School Publishing, pages 367–82.

Allen, F and Gale, D (2000), 'Financial contagion', *Journal of Political Economy*, Vol. 108, No. 1, pages 1–33.

Allen, F, Babus, A and Carletti, E (2012), 'Asset commonality, debt maturity and systemic risk', *Journal of Financial Economics*, Vol. 104, No. 3, pages 519–34.

Benos, E, Wetherilt, A and Zikes, F (2013), 'The structure and dynamics of the UK credit default swap market', *Bank of England Financial Stability Paper No. 25*; available at www.bankofengland.co.uk/financialstability/Documents/fpc/fspapers/fs_paper25.pdf.

Billio, M, Getmansky, M, Gray, D, Lo, A W, Merton, R C and Pelizzon, L (2014), 'Sovereign, bank and insurance credit spreads: connectedness and system networks ', *SYRTO Working Paper No. 8*; available at http://syrtoproject.eu/wp-content/uploads/2014/05/8_UNIVE26.pdf.

Bisias, D, Flood, M, Lo, A W and Valavanis, S (2012), 'A survey of systemic risk analytics', *Office of Financial Research Working Paper No. 0001*, US Department of the Treasury.

Blume, L, Easley, D, Kleinberg, J, Kleinberg, R and Tardos, E (2011), 'Network formation in the presence of contagious risk'. Proceedings of the 12th ACM Conference on Electronic Commerce.

Bomfin, A (2005), *Understanding Credit Derivatives and Related Instruments*, Academic Press.

Brunnermeier, M, Clerk, R and Scheicher, M (2013), 'Assessing contagion risk in the CDS market', *Banque de France Financial Stability Review*, No. 17.

Clerc, L, Gabrieli, S, Kern, S and El Omari, Y (2013), 'Assessing contagion risk through the network structure of CDS exposures on European reference entities', *mimeo*.

Cont, R, Moussa, A and Santos, E B (2013), 'Network structure and systemic risk in banking systems', in Fouque, J-P and Langsam, J A (eds), *Handbook on systemic risk*, Cambridge University Press; available at http://ebooks.cambridge.org/chapter.jsf?bid=CBO9781139151184&cid=CBO9781139151184A106.

Creal, D, Koopman, S J and Lucas, A (2013), 'Generalized autoregressive score models with applications', *Journal of Applied Econometrics*, Vol. 28, No. 5, pages 777–95.

Duffie, D (2011), 'Systemic risk exposures: a 10-by-10-by-10 approach', *mimeo*, Stanford University.

Duffie, D, Scheicher, M and Vuillemey, G (2014), 'Central clearing and collateral demand', *Journal of Financial Economics*, Vol. 116, No. 2, pages 237–56.

Elsinger, H, Lehar, A and Summer, M (2006), 'Risk assessment for banking systems', *Management Science*, Vol. 52, No. 9, pages 1,301–14.

Freixas, X, Parigi, B M and Rochet, J-C (2000), 'Systemic risk, interbank relations and liquidity provision by the central bank', *Journal of Money, Credit and Banking,* Vol. 32, No. 3, pages 611–38.

FSB (2013), 'OTC derivatives market reforms: fifth progress report on implementation', April.

Furfine, C (2003), 'Interbank exposures: quantifying the risk of contagion', *Journal of Money, Credit and Banking*, Vol. 35, No. 1, pages 111–28.

Gai, P, Haldane, A G and Kapadia, S (2011), 'Complexity, concentration and contagion', *Journal of Monetary Economics*, Vol. 58, No. 5, pages 453–70; available at www.sciencedirect.com/science/article/pii/S0304393211000481.

Giansante, S, Markose, S and Shaghaghi, A R (2012), 'Too interconnected to fail financial network of US CDS market: topological fragility and systemic risk', *Journal of Economic Behaviour & Organization*, Vol. 83, No. 3, pages 627–46.

Haldane, A G (2009), 'Rethinking the financial network'; available at www.bankofengland.co.uk/archive/Documents/historicpubs/speeches/2009/speech386.pdf.

Macroeconomic Assessment Group on Derivatives (2013), *Macroeconomic impact assessment of OTC derivative reforms*; available at www.bis.org/publ/othp20.pdf.

Markose, S (2012), 'Systemic risk from global financial derivatives: a network analysis of contagion and its mitigation with super-spreader tax', *IMF Working Paper No. 12/282*.

Newman, M E J (2005), 'A measure of betweenness centrality based on random walks', *Social Networks*, Vol. 27, No. 1, pages 39–54.

Newman, M E J (2010), *Networks: an introduction*, Oxford University Press.

Oh, D H and Patton, A J (2013), 'Time-varying systemic risk: evidence from a copula model of CDS spreads', *Working Paper*, Duke University; available at http://public.econ.duke.edu/~ap172/.

Peltonen, T A, Scheicher, M, Vuillemey, G (2013), 'The network structure of the CDS market and its determinants', *European Central Bank Working Paper No. 1583*.

Peltonen, T and Vuillemey, G (2013), 'Disentangling the bond-CDS nexus: a stress test model of the CDS market', *European Central Bank Working Paper No. 1599*.

Vivier-Lirimonty, S (2006), 'Contagion in interbank debt networks', *Working Paper*; available at www.researchgate.net/publication/242244605_Contagion_in_Interbank_Debt_Networks.

Yellen, J L (2013), 'Interconnectedness and systemic risk: lessons from the financial crisis and policy implications', remarks at the American Economic Association/American Finance Association joint luncheon, San Diego, California.