

On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms

Francis X. Diebold
University of Pennsylvania
and NBER

Kamil Yilmaz
Koç University

June 23, 2011

*** Preliminary Draft: Not For Circulation Without Authors' Permission ***

Abstract: We propose several connectedness measures built from pieces of variance decompositions, and we argue that they provide natural and insightful measures of connectedness among financial asset returns and volatilities. We also show that variance decompositions define weighted, directed networks, so that our connectedness measures are intimately-related to measures of network topology. Building on these insights, we track both average and daily time-varying connectedness of major U.S. financial institutions' stock return volatilities in recent years, including during the financial crisis of 2007-2008.

Key Words: Risk measurement, risk management, portfolio allocation, market risk, credit risk, systemic risk, asset markets, degree distribution

JEL codes: C3, G2

Contact Author: F.X. Diebold, fdiebold@sas.upenn.edu

Acknowledgments: For helpful comments we thank seminar/conference participants at UCSD, the Federal Reserve Board, the Federal Reserve Bank of Philadelphia, the University of Pennsylvania, Rice University, the Federal Reserve Bank of Kansas City, the International Monetary Fund, the Commodity Futures Trading Commission, the Pew Charitable Trusts, and the Federal Reserve Bank of Cleveland. Special thanks go to Celso Brunetti. The usual disclaimer applies.

“When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science.”

[Kelvin (1891)]

“None of us anticipated the magnitude of the ripple effects.”

[Merrill Lynch President Gregory Fleming
on the financial crisis of 2007-2008,
as reported in Lowenstein (2010)]

1 Introduction

Connectedness would appear central to modern risk measurement and management, and indeed it is. It features prominently in key aspects of market risk (return connectedness and portfolio concentration), credit risk (default connectedness), counter-party and grid-lock risk (bilateral and multilateral contractual connectedness), and not least, systemic risk (system-wide connectedness). It is also central to understanding underlying fundamental macroeconomic risks, in particular business cycle risk (intra- and inter-country real activity connectedness).

Perhaps surprisingly, then, connectedness remains a rather elusive concept, in many respects incompletely defined and poorly measured. Correlation-based measures remain widespread, for example, yet they measure only pairwise association and are uncomfortably (for financial markets) wed to linear, Gaussian thinking. Different authors chip away at this situation in different ways. The equi-correlation approach of Engle and Kelly (2009), for example, uses average correlations across all pairs. The CoVaR approach of Adrian and Brunnermeier (2008) and the marginal expected shortfall approach of Acharya et al. (2010) track association between individual-firm and overall-market movements and also rely less on linear, Gaussian methods. Although these and various other measures are certainly of interest, they measure different things, and a unified framework remains lacking.

To address this situation, in this paper we develop and apply a unified framework for conceptualizing and empirically measuring connectedness at a variety of levels, from pairwise

through system-wide, via variance decompositions from approximating models. In section 2 we introduce the conceptual framework and measures, in population. In section 3 we relate our measures to some of those in the burgeoning network literatures; the relationships turn out to be direct and important. In section 4 we treat the sample estimation of connectedness, allowing explicitly for time-variation. Finally, in the penultimate section 5, we use our framework to study connectedness at all levels among a large set return volatilities of U.S. financial institutions during the last decade, including during the financial crisis of 2007-2008. We conclude in section 6.

2 Population Connectedness

Our approach to connectedness is based on assessing shares of forecast error variation in various locations (firms, markets, countries, etc.) due to shocks arising *elsewhere*. This is intimately related to the familiar econometric notion of a variance decomposition, in which the the forecast error variance of variable i is decomposed into parts attributed to the various variables in the system. We denote by d_{ij}^H the ij -th H -step variance decomposition component; that is, the fraction of variable i 's H -step forecast error variance due to shocks in variable j . All of our connectedness measures – from simple pairwise to system-wide – are based in the “cross” variance decompositions, d_{ij}^H , $i, j = 1, \dots, N$, $i \neq j$. At the risk of belaboring the obvious, the key is $i \neq j$.

2.1 The Population Data-Generating Process

Consider an N -dimensional data-generating process (DGP) with orthogonal shocks: $x_t = \Theta(L)u_t$, $\Theta(L) = \Theta_0 + \Theta_1 L + \Theta_2 L^2 + \dots$, $E(u_t u_t') = I$. Note that Θ_0 need not be diagonal. All aspects of connectedness are contained in this very general representation. In particular, contemporaneous aspects of connectedness are summarized in Θ_0 , and dynamic aspects in $\{\Theta_1, \Theta_2, \dots\}$. Nevertheless, attempting to understand connectedness by staring at (literally) hundreds of elements of $\{\Theta_0, \Theta_1, \Theta_2, \dots\}$ is typically fruitless. One needs a transformation of $\{\Theta_0, \Theta_1, \Theta_2, \dots\}$ that better reveals connectedness. Variance decompositions achieve this.

2.2 The Population Connectedness Table

The simple Table 1, which we call a *connectedness table*, proves central for understanding the various connectedness measures and their relationships. Its main upper-left $N \times N$ block

	x_1	x_2	\dots	x_N	From Others
x_1	d_{11}^H	d_{12}^H	\dots	d_{1N}^H	$\sum_{j=1}^N d_{1j}^H, j \neq 1$
x_2	d_{21}^H	d_{22}^H	\dots	d_{2N}^H	$\sum_{j=1}^N d_{2j}^H, j \neq 2$
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
x_N	d_{N1}^H	d_{N2}^H	\dots	d_{NN}^H	$\sum_{j=1}^N d_{Nj}^H, j \neq N$
To Others	$\sum_{\substack{i=1 \\ i \neq 1}}^N d_{i1}^H$	$\sum_{\substack{i=1 \\ i \neq 2}}^N d_{i2}^H$	\dots	$\sum_{\substack{i=1 \\ i \neq N}}^N d_{iN}^H$	$\frac{1}{N} \sum_{\substack{i,j=1 \\ i \neq j}}^N d_{ij}^H$

Table 1: Connectedness Table Schematic. See Text for details.

contains the variance decompositions. For future reference we call that upper-left block a “variance decomposition matrix,” and we denote it by $D^H = [d_{ij}^H]$. The connectedness table simply augments D^H with a rightmost column containing row sums, a bottom row containing column sums, and a bottom-right element containing the grand average, in all cases for $i \neq j$.

The off-diagonal entries of D^H are the parts of the N forecast-error variance decompositions of relevance from a connectedness perspective; in particular, they measure *pairwise directional connectedness*. In general the pairwise directional connectedness from j to i is

$$C_{i \leftarrow j}^H = d_{ij}^H.$$

Note that in general $C_{i \leftarrow j}^H \neq C_{j \leftarrow i}^H$. Hence there are $N^2 - N$ separate pairwise directional connectedness measures. They are analogous to bilateral imports and exports for each of a set of N countries.

Sometimes we are interested in *net* pairwise directional connectedness, in a fashion analogous to a bilateral trade balance. We have

$$C_{ij}^H = C_{j \leftarrow i}^H - C_{i \leftarrow j}^H.$$

There are $\frac{N^2 - N}{2}$ net pairwise directional connectedness measures.

We call the off-diagonal row and column sums, labeled “from” and “to” in the connectedness table, the *total directional connectedness* measures. The total directional connectedness from others to i is

$$C_{i \leftarrow \bullet}^H = \sum_{\substack{j=1 \\ j \neq i}}^N d_{ij}^H,$$

and the total directional connectedness to others from j is

$$C_{\bullet \leftarrow j}^H = \sum_{\substack{i=1 \\ i \neq j}}^N d_{ij}^H.$$

Hence there are $2N$ total directional connectedness measures, N “to others,” or “transmitted,” and N “from others,” or “received,” analogous to total exports and total imports for each of a set of N countries.

Just as with pairwise directional connectedness, sometimes we are interested in net total effects. Net total directional connectedness is

$$C_i^H = C_{\bullet \leftarrow i}^H - C_{i \leftarrow \bullet}^H.$$

There are N net total directional connectedness measures, analogous to the total trade balances of each of a set of N countries.

Finally, the grand total of the off-diagonal entries in D^H (equivalently, the sum of the “from” column or “to” row) measures *total connectedness*. We have¹

$$C^H = \frac{1}{N} \sum_{\substack{i,j=1 \\ i \neq j}}^N d_{ij}^H.$$

There is just one total connectedness measure, as total connectedness distills a system into a single number analogous to total world exports or total world imports (the two are of course identical).

The connectedness table makes clear how one can begin with the most disaggregated (e.g., microeconomic, firm-level pairwise directional) connectedness measures and aggregate them in various ways to obtain macroeconomic economy-wide total directional and total connectedness. Different agents may be relatively more interested in one or another of the measures. For example, firm i may be maximally interested in how various others connect to it ($C_{i \leftarrow j}^H$, for various j), or how *all* others connect to it, $C_{i \leftarrow \bullet}^H$. In contrast, regulators might be more concerned with identifying systemically important firms j , in the sense of large total directional connectedness to others from j , $C_{\bullet \leftarrow j}^H$, and they might also be more concerned with monitoring total (system-wide) connectedness C^H .

¹Note that we construct total connectedness by taking off-diagonal D^H variation relative to total D^H variation (N), so that C^H is expressed as a decimal share, as with “from” total directional connectedness.

2.3 Correlated Shocks

In the orthogonal structural system discussed thus far, the variance decompositions are easily calculated, because orthogonality ensures that the variance of a weighted sum is simply an appropriately-weighted sum of variances. But reduced-form shocks are rarely orthogonal. To identify uncorrelated structural shocks from correlated reduced-form shocks, one must, inescapably, make assumptions. Sometimes this is more-or-less transparent, as in the Cholesky-factor identifications that trace to Sims (1980), or in any of the scores of subsequent “structural VAR” identifications. Sometimes it is less clear, but equally true, as in the generalized variance decomposition (GVD) framework of Koop et al. (1996) and Pesaran and Shin (1998), or in deeply-structural dynamic stochastic general equilibrium approaches like DiNicolo and Lucchetta (2010), which embody large sets of maintained assumptions.

Any set of identifying assumptions may fail. Results based on traditional Cholesky-factor identification, for example, may be sensitive to ordering, as Cholesky-factor identification amounts to assumption of a particular recursive ordering. No set of identifying assumptions is necessarily better or worse than any other, particularly as many models are exactly- as opposed to over-identified, so that the identifying restrictions can not be tested. Reasonable people may disagree as to their preferred assumptions, and they often do. We have nothing to add to such debates; one must simply take a stand and march onward conditional upon (and cognizant of) the maintained assumptions.

Our own preferences run toward Cholesky and related identifications. We often find that total connectedness is robust to Cholesky ordering; that is, the range of total connectedness estimates across orderings is often quite small. Directional connectedness, however, is sometimes more sensitive to Cholesky ordering, which enhances the appeal of GVDs. Like Cholesky-based variance decompositions, GVDs are based on a largely-atheoretical identification scheme, but they are independent of ordering.²

We refer the interested reader to Pesaran and Shin (1998), which builds on Koop et al. (1996), for motivation and background regarding the GVD identification scheme and instead proceed directly to a precise statement of the resultant GVD. The H -step generalized variance decomposition matrix $D^{gH} = [d_{ij}^{gH}]$ has entries

$$d_{ij}^{gH} = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} (e_i' \Theta_h \Sigma e_j)^2}{\sum_{h=0}^{H-1} (e_i' \Theta_h \Sigma \Theta_h' e_i)},$$

²GVDs of course make other assumptions, most notably normality, and hence may be more useful for assessing connectedness of (log) volatilities, which are well-approximated as Gaussian, than for returns, which are not.

where e_j is a selection vector with j^{th} element unity and zeros elsewhere, Θ_h is the coefficient matrix multiplying the h -lagged shock vector in the infinite moving-average representation of the non-orthogonalized VAR, Σ is the covariance matrix of the shock vector in the non-orthogonalized VAR, and σ_{jj} is the j^{th} diagonal element of Σ .³ Because shocks are not necessarily orthogonal in the *GVD* environment, sums of forecast error variance contributions are not necessarily unity (that is, row sums of D^g are not necessarily unity).⁴ Hence we base our generalized connectedness indexes not on D^g , but rather on $\tilde{D}^g = [\tilde{d}_{ij}^g]$, where $\tilde{d}_{ij}^g = \frac{d_{ij}^g}{\sum_{j=1}^N d_{ij}^g}$. Obviously, by construction $\sum_{j=1}^N \tilde{d}_{ij}^g = 1$ and $\sum_{i,j=1}^N \tilde{d}_{ij}^g = N$. Armed with \tilde{D}^g , we can immediately calculate generalized connectedness measures \tilde{C} , $\tilde{C}_{\bullet \leftarrow j}$, $\tilde{C}_{i \leftarrow \bullet}$, \tilde{C}_i , $\tilde{C}_{i \leftarrow j}$, $\tilde{C}_{j \leftarrow i}$, and \tilde{C}_{ij} .

3 Network Topology and Connectedness

Networks are everywhere in modern life, from power grids to Facebook. Not surprisingly, research on networks has grown explosively in recent years.⁵ Once one starts thinking about networks, one is naturally led to think about network connectedness. Deep questions abound. Just what *is* network connectedness? Is it a pairwise or system-wide concept, or both, or neither? How might it be related to the notion of connectedness that we have thus far emphasized, based on variance decompositions? Given an acceptable definition of network connectedness in theory, can we *estimate* it in practice, from real data?

Interestingly, it turns out that our connectedness measures, early variants of which were proposed in Diebold and Yilmaz (2009) independently of the network literature, are closely related to several measures commonly used to describe network topology. Indeed as we shall soon see in detail, variance decompositions effectively *are* networks.

A network is composed of N nodes and L links between nodes. In the classical development, a network is simply an $N \times N$ adjacency matrix A of zeros and ones,

$$A = [A_{ij}], \tag{1}$$

where $A_{ij} = 1$ if nodes i and j are linked, and $A_{ij} = 0$ otherwise. Note that A is symmetric, because if i and j are connected, then so too must be j and i . It is no accident that we

³Note the typo in the original paper of Pesaran and Shin (1998), p. 20. They write σ_{ii}^{-1} but should have written σ_{jj}^{-1} .

⁴We now drop the “ H ” superscripts, because from this point onward they are not needed for clarity.

⁵Newman (2010) and Jackson (2008) provide good general and economic introductions, respectively.

elevate the status of equation (1), separating it from the text and giving it its own number. Mathematically (i.e., algebraically), the adjacency matrix A is the network, and all network properties are embedded in A .

Assessing network connectedness and other properties is challenging, however, and there is no single, all-encompassing measure. The only thing that’s all-encompassing is the matrix A itself, yet staring for example at an A matrix of dimension 500×500 (let alone $500,000 \times 500,000$) is not likely to be revealing. Hence we analyze A in various complementary ways. The most useful for our purposes is based on the idea of node degree, to which we now turn.⁶

3.1 Node Degree and the Degree Distribution

A node’s *degree* is its number of links to other nodes. Immediately the degree of node i is

$$d_i = \sum_{j=1}^N A_{ij} = \sum_{j=1}^N A_{ji}.$$

Degree is a single-node property, but we can of course examine the pattern of degrees across nodes. The *degree distribution* is the probability distribution of degrees across nodes. It is a univariate distribution with support $0, \dots, (N - 1)$, and it summarizes certain aspects of network connectedness.⁷ Important aspects include its location, scale, skewness and tail thickness, which we discuss in greater detail below.⁸

3.1.1 Location, Scale and Symmetry

Global connectedness considerations suggest examining the location of the degree distribution. Hence an obvious measure of network connectedness is the mean of the degree distribution. A high mean degree indicates high overall connectedness. If the degree distribution is asymmetric, perhaps with a long right tail due to a few outlying high-degree nodes, then the median may prove a more robust and reliable location measure. Related, we may want to examine dispersion of degrees around the mean (that is, examine degree scale in addition to location), as assessed for example by standard deviation or interquartile range.

⁶ Bech and Atalay (2010) and Adamic et al. (2010) provide good reviews and financial applications of network-theoretic connectedness measures. Our interest, however, is more directly and exclusively focused on node degrees and their distribution.

⁷The support of $0, \dots, (N - 1)$ stems from our adoption of the standard convention that $A_{ii} = 0, \forall i$.

⁸One could also examine aspects of the degree distribution beyond its moments, such as “inequality” as measured by Gini coefficients or Herfindahl indexes, which effectively measure deviation of a distribution from uniformity.

3.1.2 Tail Thickness

The shape of the degree distribution governs many network properties. We have already mentioned the first three moments, linked to notions of location, scale and symmetry. Another important aspect of degree distribution shape is tail thickness, which is related to the fourth moment, or kurtosis.

The binomial $B((N-1), \theta)$ is one of two important benchmark degree distributions. It is motivated by mathematical simplicity rather than empirical relevance; in particular, it flows immediately from the simplest imaginable probabilistic model of link formation with fixed Bernoulli probability θ . Erdős and Rényi (1959) constructed their “random networks” in this way a half-century ago, obtaining:

$$f(d) = \binom{N-1}{d} \theta^d (1-\theta)^{N-1-d},$$

where f is the degree distribution, d is degree, N is the number of network nodes, and θ is the link-formation probability. For $N \rightarrow \infty$ and $\theta \rightarrow 0$ with $N\theta = \lambda$, the usual Poisson approximation holds,

$$f(d) = \frac{\lambda^d e^{-\lambda}}{d!}.$$

Similarly, for $N \rightarrow \infty$ and θ fixed the usual Gaussian approximation holds,

$$f(d) = \frac{1}{\sqrt{N\theta(1-\theta)2\pi}} e^{-\frac{1}{2}\left(\frac{d-N\theta}{\sqrt{N\theta(1-\theta)}}\right)^2}.$$

Whether viewed as binomial, Poisson or Gaussian, the key feature of the degree distribution in the Erdős-Rényi random networks environment is its “thin tails,” which decay exponentially quickly. This makes huge-degree nodes unlikely.

Binomial thin tails contrasts with the second important benchmark degree distribution, with tails governed by a power law:

$$f(d) = cd^{-\lambda},$$

or equivalently, in the distribution’s tails,

$$\ln(f(d)) = \ln c - \lambda \ln(d).$$

In contrast to the quickly- (exponentially-) decaying tails of Gaussian distributions, power-

law distributions have slowly- (hyperbolically-) decaying tails, and they are scale-free in that $\frac{f(d_1)}{f(d_2)} = \frac{f(cd_1)}{f(cd_2)}$, $\forall c$. The fat tails associated with power law degree distributions encourage the appearance of huge-degree nodes, sometimes called hubs. The power law was noticed in empirical degree distributions and later shown to arise theoretically when nodes form new links sequentially with probability proportional to their degree, as in the “scale-free networks” of Barabási and Albert (1999).

The tension between Gaussian (thin-tailed) and power-law (fat-tailed) distributions is a key theme in modern asymptotic statistical theory. Under very general conditions, the large-sample distributions of appropriately-standardized sums are members of the stable family, but the stable family spans both Gaussian and power-law distributions.⁹ As we have noted, the Gaussian and power-law benchmarks arise theoretically under certain conditions, and one or the other sometimes appears to accord with observation, but the degree distribution for any particular network is ultimately an empirical matter and may conform neither to a Binomial distribution nor to a power law.

3.2 k-Step Degree Distributions and the Distance Distribution

The just-described adjacency matrix and degree distribution are what we will call “1-step,” as the links are direct. We write the 1-step adjacency matrix A with greater precision as $A^{(1)}$. However, even if i is not directly linked to j , i may be linked to k , and k to j , so that i and j are linked at a distance of two steps rather than one. Hence we also have a 2-step adjacency matrix $A^{(2)}$.¹⁰

Higher-order degree distributions in general, and higher-order mean degrees in particular, capture subtle but potentially very important deep network connectedness properties. For example, 2-degrees are closely-related to so-called “eigenvector centralities,” which are like 1-degrees but with links *weighted* by numbers of *their* links rather than simply added.¹¹ This captures the idea that a node’s 1-degree may be low, but if its few 1-links are in fact to very-highly-connected nodes, then it is effectively very highly connected.

It is of interest to examine the full set of k -step adjacency matrices $A^{(k)}$, the corresponding sequence of k -step degree distributions $f^{(k)}(d^{(k)})$, and the k -step mean degrees $E(d^{(k)})$, which must be non-decreasing in k , $k = 1, \dots, (N - 1)$. Of central interest are the speed and

⁹See, for example, Embrechts et al. (1997) and the references therein for rigorous scientific development, and Taleb (2007) for popular exposition.

¹⁰Note that if two nodes are 1-step linked, they are automatically 2-step linked.

¹¹See, for example, Newman (2010).

pattern with which the k -step degree distribution approaches placing unit probability mass on $d^{(k)} = (N - 1)$ as $k \rightarrow (N - 1)$.

A *path* is a sequence of links connecting one node to another. The *distance* s_{ij} between two nodes i and j is the smallest number of links that must be traversed to go from i to j .¹² Distance is a two-node property, in contrast to degree, which is a single-node property. A network is *connected* if $s_{ij} \leq N - 1, \forall i, j$. $s_{max} = \max_{i,j} s_{ij}$ is called the network *diameter*. We can assemble all distances s_{ij} into an $N \times N$ distance matrix S .

In fashion parallel to the degree distribution, we can learn from the *distance distribution*, the univariate distribution of the s_{ij} , defined on $0, \dots, N - 1$. There is only one distance distribution, as the distance matrix blends information from the various k -step adjacency matrices, and similarly, the distance distribution blends information from the k -step degree distributions. Mean distance is of particular interest in characterizing connectedness.¹³

A beautiful large- N approximation relates network diameter, size and mean 1-degree in Erdős-Rényi random networks:¹⁴

$$s_{max} \approx \frac{\ln N}{\ln E(d^{(1)})}.$$

In particular, network diameter grows only with the *logarithm* of network size. This $\ln N$ diameter approximation provides a mathematical characterization of the “six-degree” or “small-world” phenomenon, namely that diameters tend to be small even for huge networks, provided that $E(d^{(1)}) > 1$. For example, for $N = 300,000,000$ (roughly the U.S. population) and mean 1-degree $E(d^{(1)}) = 20$, we have network diameter $s_{max} \approx 6$.

Erdős-Rényi random networks are often poor descriptions of real-world networks, however, due for example to empirical power-law degree distributions or node clustering.¹⁵ Interestingly, however, Watts and Strogatz (1998) have shown that the $\ln N$ diameter approximation nevertheless holds in networks with small clusters of linked neighbors with just a few “long-range” links.

¹²Sometimes distance is called path length.

¹³Sometimes mean distance is called characteristic path length.

¹⁴See, for example, Newman (2010), p. 420.

¹⁵Clustering refers to the fact that in real networks two “people” with a common “friend” are more likely to be friends than two randomly-selected people. Such clustering triangles sometimes occur disproportionately often in real networks but not in Erdős-Rényi random networks, in which any given link simply exists with probability θ .

3.3 Variance Decompositions as Weighted, Directed Networks

Having introduced some elementary aspects of networks and their description, we now arrive at the fundamental insight on which this paper builds: *variance decompositions are networks*. More completely, the variance decomposition matrix D , which defines our connectedness table and all associated connectedness measures, is a network adjacency matrix A . Hence variance decompositions define networks, and network connectedness measures may be used in conjunction with variance decompositions of a set of variables x to understand connectedness among components.

The networks defined by variance decompositions, however, are rather more sophisticated than the classical network structures sketched thus far. First, the adjacency matrix $A^{(1)}$ (variance decomposition matrix D) is not filled simply with 0-1 entries; rather, the entries are *weights*, with some potentially strong and others potentially weak. Second, the links are *directed*; that is, the strength of the ij link is not necessarily the same as the strength of the ji link, so the adjacency matrix is generally not diagonal. Third, there are constraints on the row sums of $A^{(1)}$. In particular, each row must sum to 1 because the entries are variance shares. Hence we write the diagonal elements as $A_{ii}^{(1)} = 1 - \sum_{\substack{j=1 \\ j \neq i}}^N A_{ij}^{(1)}$. Note in particular that the diagonal elements of $A^{(1)}$ are no longer 0.

Weighted, directed versions of the earlier-introduced network connectedness statistics are readily defined, including degrees and degree distributions, and distances and the distance distribution. For example, node degrees are now obtained not by summing zeros and ones, but rather by summing pairwise weights, and there are now “to-degrees” and “from-degrees,” corresponding to row sums and column sums.¹⁶ The 1-step from-degree of node i is $d_i^{(1)from} = \sum_{\substack{j=1 \\ j \neq i}}^N A_{ij}^{(1)}$. The 1-step from-degree distribution is the probability distribution of from degrees across nodes. It is a univariate distribution with support on the unit interval. Similarly, the 1-step to-degree of node j is $d_j^{(1)to} = \sum_{\substack{i=1 \\ i \neq j}}^N A_{ij}^{(1)}$. The 1-step to-degree distribution is the probability distribution of to degrees across nodes. It is a univariate distribution with support on $[0, N]$.

Two key insights emerge readily. First, our earlier-defined total directional connectedness measures are precisely 1-step from-degrees and to-degrees. Second, our total connectedness measure is simply the mean 1-degree (to or from – it’s the same either way).

¹⁶Our to-degrees and from-degrees are often called “out-degrees” and “in-degrees” in the network literature.

4 Sample (Time-Varying) Connectedness

Clearly C depends on the set of variables x whose connectedness is to be examined, the predictive horizon H for variance decompositions, and the dynamics $A(L)$, so we write $C(x, H, A(L))$.¹⁷ In reality A is unknown and must be approximated (e.g., using a finite-ordered vector autoregression). Recognizing the centrality of the approximating model adopted, we write $C(x, H, A(L), M(\theta))$, where θ is a finite-dimensional parameter.

In addition, we want to allow for time-varying connectedness, which allows us to move from the static, unconditional, perspective implicitly adopted thus far, to a dynamic, conditional perspective. Time-varying $A(L)$, and hence time-varying connectedness, may arise for a variety of reasons. $A(L)$ may evolve slowly with evolving tastes, technologies and institutions, or it may vary with the business cycle, or it may shift abruptly with financial market environment (e.g., crisis, non-crisis). Whether and how much $A(L)$ varies is ultimately an empirical matter and will surely differ across applications, but in any event it would be foolish simply to assume it constant. Hence we allow the connection table and all of its elements to vary over time, and we write $C_t(x, H, A_t(L), M(\theta_t))$.

Finally, everything we have written thus far is in *population*, whereas in reality we must use an approximating model *estimated* using data $1 : T$, so we write $\widehat{C}_t(x, H, A_t(L), \widehat{M}_{1:T}(\hat{\theta}_t))$. To economize on notation we henceforth drop $A(L)$, because it is determined by nature rather than a choice made by the econometrician, relying on the reader to remember its relevance and writing $\widehat{C}_t(x, H, \widehat{M}_{1:T}(\hat{\theta}_t))$.

In what follows we successively discuss aspects of x , H and $\widehat{M}_{1:T}(\hat{\theta}_t)$.

4.1 The Reference Universe, x

Connectedness measurements are defined only with respect to a reference universe, namely the set of x 's defining the object of interest to be studied. Choice of x has important implications for the appropriate approximating model; for example, x may (or may not) be strongly serially correlated, conditionally heteroskedastic, or highly disaggregated. Connectedness measurements generally will not, and should not, be robust to choice of reference universe.

Three sub-issues arise, which we call the “ x object,” the “ x choice,” and the “ x frequency.” By x object we refer to the type of x variable studied, typically either returns or return volatilities. By x choice we mean precisely which (and hence how many) x variables are

¹⁷The same holds, of course, for the various directional connectedness measures, so we use $C(x, H, A(L))$ as a stand-in for all our connectedness measures.

chosen for study. By x frequency we refer to the observational frequency of the x variables (daily, monthly, ...). In this paper the x object is realized equity return volatility, the x choice is approximately fifteen major U.S. financial institutions, and the x frequency is daily. We will provide details subsequently.

4.2 The Predictive Horizon, H

Certain considerations in certain contexts may help guide selection of connectedness *horizon*, H . For example, in risk management contexts, one might focus on H values consistent with risk measurement and management considerations. $H = 10$, for example, would cohere with the 10-day value at risk (Var) required under the Basel accord. Similarly, in portfolio management contexts one might link H to the rebalancing period.

The connectedness horizon is important particularly because it is related to issues of dynamic connectedness (in the fashion of contagion) as opposed to purely contemporaneous connectedness. To take a simple pairwise example, shocks to j may impact the forecast error variance of i only with a lag, so that $C_{i \leftarrow j}$ may be small for small H but nevertheless large for larger H .¹⁸ Intuitively, as the horizon lengthens there may be more chance for connectedness to appear. Thus, in a sense, varying H lets us break connectedness into “long-run,” “short-run,” etc. More precisely, as H lengthens we obtain a corresponding sequence of conditional prediction error variance decompositions for which the conditioning information is becoming progressively less valuable. In the limit as $H \rightarrow \infty$, we obtain an *unconditional* variance decomposition.

In this paper we anchor on a horizon of $H = 12$ days, and we also examine a range of nearby H values. In a sense this provides a robustness check, but as we argued above, there is no reason why connectedness should be “robust” to H . Hence we view examination of a menu of H values simply as an interesting part of a phenomenological investigation.

4.3 The Approximating Model, $M(\theta_t)$

For return volatilities, conditional mean dynamics in $M(\theta)$ will be important and will surely need to be modeled. The obvious workhorse approximating model is a vector autoregression, $VAR(p)$. A logarithmic VAR is often appropriate for volatilities, as in Andersen et al. (2003).

¹⁸Such dynamic phenomena, and the rich patterns that are possible, are closely related to aspects of multi-step Granger causality, as treated for example in Dufour and Renault (1998), Dufour and Taamouti (2010), and the references therein.

Indeed, throughout this paper we use a third-order vector autoregressive approximating model for log volatilities.

Connectedness is just a transformation of system coefficients. Hence if the coefficients are time-varying, so too will be connectedness. Tracking (“nowcasting”) real-time connectedness movement is of central interest. Connectedness may be a highly nonlinear phenomenon, and time-varying parameters are an important way to allow for nonlinearity. Indeed as “White’s theorem” makes clear (see Granger (2008)), linear models with time-varying parameters are actually very general approximations to arbitrary nonlinear models. Ultimately we want to try to escape the shackles of linear, Gaussian, correlation-based analysis, as in much recent work including Härdle et al. (2011).

In this paper we capture parameter variation by using a rolling estimation window; we write $\hat{C}_t(x, H, \widehat{M}_{t-w:t}(\hat{\theta}))$, where w denotes the width of the uniform one-sided window. That is, we sweep through the sample, on each day using only the most recent w business days to estimate the approximating vector autoregression and calculate connectedness measures.

The rolling-window approach has the advantages of tremendous simplicity and coherence with a wide variety of possible DGPs involving time-varying parameters. However it does of course require choice of the tuning parameter w , in a manner precisely analogous to bandwidth choice in density estimation. In this paper we anchor on $w = 100$ days, but we also explore robustness to alternative choices.

5 U.S. Financial Institution Connectedness

Thus far we have introduced tools for connectedness measurement and related them to tools for describing the structure of weighted directed networks. We now put those tools to work, using them to monitor and characterize the evolution of connectedness among major U.S. financial institutions before and during the 2007-2008 financial crisis. Understanding such financial connectedness is of interest not only in terms of understanding financial crises, but also in terms of understanding the business cycle, as the financial system’s health has important implications for the pace of real activity.

We proceed in four steps. First, in section 5.1, we describe the data that we use to measure financial institution connectedness. Next, in section 5.2, we perform a full-sample (static) analysis, in which we effectively characterize average, or unconditional, connectedness. This is not only of intrinsic interest, but it also sets the stage for section 5.3, where we perform a rolling-sample (dynamic) analysis of conditional connectedness. Our ultimate

interest lies there; we monitor high-frequency (daily) connectedness as conditions evolve, sometimes gradually and sometimes abruptly. Finally, in section 5.4, we “zoom in” on financial institution connectedness during the global financial crisis of 2007-2008.

5.1 Data

Financial institutions are connected directly through counter-party linkages associated with positions in various assets, through contractual obligations associated with services provided to clients and other institutions, and through deals recorded in their balance sheets. High-frequency analysis of financial institution connectedness therefore might seem to require high-frequency balance sheet and related information, which is generally unavailable.

Fortunately, however, we have available stock market returns and return volatilities, which reflect forward-looking assessments of many thousands of smart, strategic and often privately-informed agents as regards precisely the relevant sorts of connections. We use that data to measure connectedness and its evolution. It is important to note that we remain agnostic as to how connectedness arises; rather, we take it as given and seek to measure it correctly for a wide range of possible underlying causal structures.¹⁹

In this paper we study *volatility* connectedness, for at least two reasons. First, if volatility tracks investor fear (e.g., the VIX is often touted as an “investor fear gauge”), then volatility connectedness is the “fear connectedness” expressed by market participants as they trade. We are interested in the level, variation, paths, patterns and clustering in precisely that fear connectedness. Second, volatility connectedness is of special interest because we are particularly interested in crises, and volatility is particularly crisis-sensitive.

Volatilities are latent and hence must be estimated. In this paper we use realized volatilities, which have received significant attention in recent years.²⁰ For a given firm on a given day, we construct daily realized return volatility using high-frequency intra-day data from the Trade and Quote (TAQ) database. In particular, we calculate daily realized volatility as the sum of squared log price changes over the 78 5-minute intervals during trading hours, from 09:00-12:00 and 13:00-16:30.

We treat realized volatility as the object of direct interest, as in Andersen et al. (2003).²¹

¹⁹Obviously there are tradeoffs, but we prefer an approach that potentially achieves much under minimal assumptions, in contrast to a more deeply structural approach that in principle could achieve even more, but only under heroic assumptions, and that may not be robust to violations of those assumptions.

²⁰For surveys see Andersen et al. (2006), Andersen et al. (2010) and Andersen et al. (2011).

²¹This contrasts with an alternative approach that views realized volatility not as the direct object of interest, but rather as an estimate of underlying quadratic variation. In that case one might want to

Institution	Ticker	Business	Market Cap.	
			12/29/06	12/31/09
J.P. Morgan	JPM	C-Bank	169	171
Wells Fargo	WFC	C-Bank	121	137
Bank of America	BAC	C-Bank	241	131
Citigroup	C	C-Bank	274	76
US Bancorp	USB	C-Bank	64	43
Bank of NY Mellon	BK	C-Bank	30	34
PNC Bank	PNC	C-Bank	22	24
American Express	AXP	Credit Cards	74	49
Goldman Sachs	GS	I-Bank	86	86
Morgan Stanley	MS	I-Bank	85	40
Fannie Mae	FNM	Mortgages	59	1.3
Freddie Mac	FRE	Mortgages	47	0.9
AIG	AIG	Insurance	187	4

Table 2: U.S. Financial Institution Detail. C-Bank denotes a commercial bank, and I-Bank denotes an investment bank. Market capitalizations are in billions of U.S. dollars. Fannie Mae and Freddie Mac were placed in government conservatorship on September 7, 2008, and AIG began government ownership on September 17, 2008.

This is appropriate because for the large, heavily-traded firms that we examine, five-minute sampling is frequent enough largely to eliminate measurement error, yet infrequent enough such that microstructure noise (e.g., due to bid-ask bounce) is not a concern. In addition, and importantly, realized volatility actually *is* an object of direct interest, traded in the volatility swap markets, in contrast to underlying quadratic variation or any other object that realized volatility may or may not be construed as estimating.

Volatilities tend to be strongly serially correlated – much more so than returns, particularly when observed at relatively high frequency. We will capture that serial correlation using vector-autoregressive approximating models, as described earlier. Volatilities also tend to be distributed asymmetrically, with a right skew, and approximate normality is often obtained by taking natural logarithms. Hence we work throughout with log volatilities.

5.2 Static (Full-Sample, Unconditional) Analysis

Here we study stock return volatilities for thirteen major U.S. financial institutions that survived the crisis of 2007-2008. In Table 2 we list the firms, tickers, market capitalization before and after the crisis, and critical episodes/dates during the crisis. Our sample includes seven commercial banks, two investment banks, one credit card company, two mortgage finance companies and one insurance company. Stocks of all firms except Fannie Mae and Freddie Mac were included in the S&P500 prior to the sub-prime crisis of 2007.

Our sample begins in May 1999 and ends in April 2010. Starting in 1999 allows us to include among our firms Goldman Sachs, Morgan Stanley and U.S. Bancorp, all of which went public in the late 1990s. Our sample also spans several important financial market episodes in addition to the crisis of 2007-2008.²² These include the dot-com bubble collapse of 2000, the Enron scandal of October 2001, and the WorldCom/MCI scandal and bankruptcy of July 2002. Hence we can not only assess connectedness of our firms during the crisis of 2007-2008, but also compare and contrast connectedness during other episodes.

Perhaps our decision to include AIG deserves some discussion. We include AIG because it was a major supplier of “financial insurance” in the 2000s, selling credit default swaps (CDSs) through its AIG Financial Products arm in London. Although CDSs provided lucrative business for AIG early on, contributing close to seventeen percent of revenue in 2005, they singlehandedly brought down AIG as the financial crisis of 2007-2008 spread across financial markets.

The full-sample connectedness table appears as Table 3. Let us begin with the pairwise directional connectedness measures, $\tilde{C}_{i \leftarrow j}^H$, which are the off-diagonal elements of the 13x13 matrix. A quick inspection of Table 3 shows that the highest pairwise connectedness measure observed from is between from Freddie Mac to Fannie Mae ($\tilde{C}_{FNM \leftarrow FRE}^H = 22\%$). In return, the pairwise connectedness from Fannie Mae to Freddie Mac ($\tilde{C}_{FRE \leftarrow FNM}^H = 17.6\%$) is ranked second. The two mortgage finance companies have been viewed very much as twins by the markets and it is quite normal that their pairwise connectedness measures are quite high. When we net the two gross measures out, the resulting net pairwise directional connectedness from Freddie Mac to Fannie Mae is 4.6%, that is, $\tilde{C}_{FRE, FNM}^H = 4.6\%$.

The next largest pairwise directional connectedness takes place from Morgan Stanley to Goldman Sachs ($\tilde{C}_{GS \leftarrow MS}^H = 13.3\%$), two top investment banks that were able to survive the

acknowledge estimation error explicitly, as in Hansen and Lunde (2010). Doing so will generally increase the persistence of the estimated volatility dynamics.

²²The 2007-2008 crisis may itself be split into the sub-prime / liquidity crisis of 2007 and the financial crisis of 2008.

	AXP	BAC	BK	C	GS	JPM	MS	PNC	USB	WFC	AIG	FNM	FRE	FROM
AXP	20.0	8.5	7.1	10.3	5.8	9.8	8.8	5.1	8.0	7.8	3.2	2.6	3.0	80.0
BAC	8.3	19.1	6.0	10.6	5.8	8.0	7.4	6.1	7.1	9.2	4.2	3.5	4.6	80.9
BK	8.4	8.3	18.8	8.4	6.2	9.3	8.5	5.7	8.4	8.3	4.2	2.4	3.0	81.2
C	9.5	9.6	5.4	20.4	4.9	8.7	7.8	5.2	7.0	8.0	5.4	3.5	4.7	79.6
GS	8.2	8.6	6.8	7.6	22.1	8.8	13.3	4.0	6.0	7.6	2.4	1.9	2.6	77.9
JPM	10.2	8.6	7.1	10.6	6.2	18.8	9.5	5.2	7.8	7.3	3.6	2.5	2.6	81.2
MS	9.2	8.3	7.1	8.9	9.8	9.7	20.5	4.2	5.5	7.1	3.4	2.8	3.6	79.5
PNC	7.7	8.8	7.4	8.5	4.6	7.6	6.6	18.1	7.6	8.8	5.2	4.2	4.9	81.9
USB	9.3	9.9	7.6	9.9	5.7	8.7	6.4	5.4	20.1	8.5	4.3	1.6	2.7	79.9
WFC	8.3	10.2	6.5	9.8	6.2	7.6	7.1	5.9	7.3	18.0	3.8	3.8	5.3	82.0
AIG	5.3	7.3	4.9	8.8	2.6	5.2	4.9	6.2	6.0	5.6	27.5	6.6	9.0	72.5
FNM	4.2	5.4	2.5	6.0	2.3	3.5	3.8	5.5	1.9	6.8	6.5	29.6	22.0	70.4
FRE	4.3	6.3	2.9	6.5	2.6	3.3	4.1	5.2	2.9	7.3	7.4	17.6	29.6	70.4
TO	92.9	99.7	71.3	106.1	62.7	90.2	88.2	63.7	75.5	92.2	53.8	53.1	68.1	78.3
NET	13.0	18.8	-9.9	26.5	-15.2	8.9	8.7	-18.2	-4.4	10.2	-18.7	-17.4	-2.3	

Table 3: Full-Sample Connectedness Table. The sample is May 4, 1999 through April 30, 2010. The ij -th entry of the upper-left 13x13 firm submatrix gives the ij -th pairwise directional connectedness; i.e., the percent of 12-day-ahead forecast error variance of firm i due to shocks from firm j . The rightmost (“FROM”) column gives total directional connectedness (from); i.e., row sums (“from all others to i ”). The bottom (“TO”) row gives total directional connectedness (to); i.e., column sums (“to all others from j ”). The bottom-most (“NET”) row gives the difference in total directional connectedness (to-from). The bottom-right element (in boldface) is total connectedness (mean “from” connectedness, or equivalently, mean “to” connectedness).

2007-08 financial crisis.²³ While the connectedness from Goldman Sachs to Morgan Stanley is also high ($\tilde{C}_{MS \leftarrow GS}^H = 9.8\%$), in net terms the directional connectedness takes place from Morgan Stanley to Goldman Sachs stock ($\tilde{C}_{GS,MS}^H = 3.5\%$).

The highest values of pairwise directional connectedness measures among the commercial bank stocks are observed to take place from Citigroup, on the one hand, and Bank of America and J.P. Morgan, on the other ($\tilde{C}_{BAC \leftarrow C}^H = \tilde{C}_{JPM \leftarrow C}^H = 13.3\%$). A high value of pairwise connectedness from Citigroup to either of Bank of America and/or J. P. Morgan shows that being the worst hit institution among the top five commercial banks, Citigroup's stock spread its troubles to the stocks of other top commercial banks.

As we have seen above, Fannie Mae and Freddie Mac are tightly connected to each other. Their pairwise connectedness with AIG also indicated that they are well connected with AIG as well. Pairwise directional connectedness of the stocks of these three institutions with the stocks of each of the remaining financial institutions tend to be much lower than connectedness of other bank stocks in our sample. We need to remind that these three institutions had lots of difficulties during the 2007-08 financial crisis and could have gone bankrupt had the U.S. government not intervened in financial markets in September 2008.

The row sum of the pairwise connectedness measures results in the total directional connectedness from others to each of the thirteen stocks (see Section 2). In other words, the "FROM" column measures the share of volatility shocks received from other financial firm stocks in the total variance of the forecast error for each stock. By definition, it is equal to 100 % minus the own share of the total forecast error variance. As the own-effects (diagonal elements of the matrix) range between 18 and 30 %, the total directional connectedness in the "FROM" column range between 70 and 82 %.

Similarly, the column sum of all pairwise connectedness measures results in the corresponding stock's total directional connectedness to others. As each stock's contribution to others' forecast error variances is not constrained to add up to 100 %, entries in the "TO" row can exceed 100 %. While the financial stocks are largely similar in terms of receiving volatility shocks from others, they are highly differentiated as transmitters of volatility shocks to others. The stark difference between the distributions of the two connectedness measures is clearly observed in their respective empirical survivor functions presented in Figure 1. Compared to the very steep survivor function defined over a narrow range for the connectedness from others, the survivor function for the connectedness to others is quite

²³As other three investment banks ceased to exist in 2008, they are not included in the full sample connectedness table.

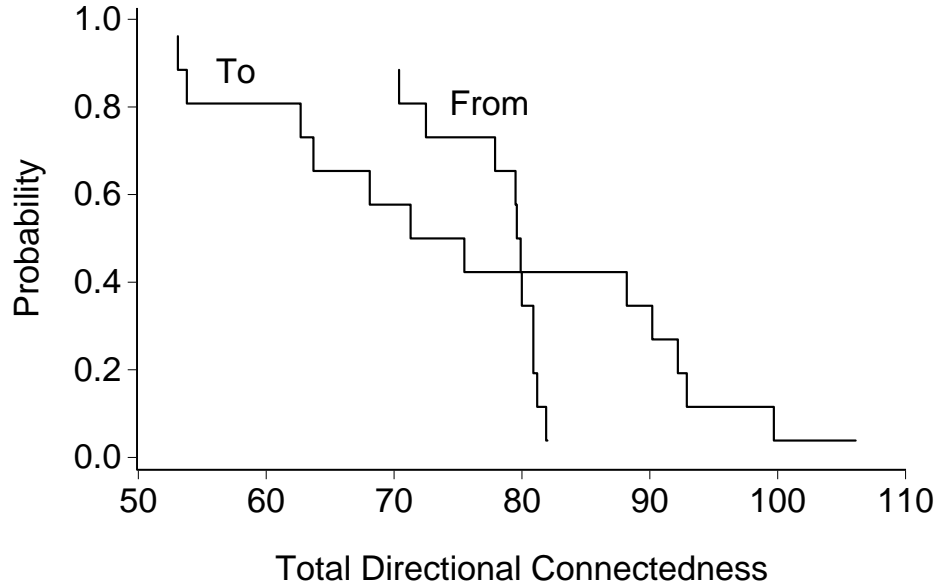


Figure 1: Full-Sample Total Directional Connectedness: Empirical Survivor Functions. We plot the empirical survivor functions for total directional connectedness “to” others and “from” others. The predictive horizon for the underlying variance decomposition is 12 days.

flat and defined over a wider range. Starting at a minimum of 70 % for Fannie and Freddie and increasing only up to a maximum of 82 % for Wells Fargo and PNC Bank, the total directional connectedness from others is distributed rather tightly. The total directional connectedness to others, on the other hand, vary from a low of 53 % for Fannie Mae, to all the way up to 106 % for the Citigroup: A range of 53 points for the connectedness to others compared to a range of just 12 points for the connectedness from others.

The largest commercial banks (as of 2010) were the ones that have the highest values of connectedness (all exceeding 90 %) to others. Being the most vulnerable among them, Citibank generated a total directional connectedness measure of 106 % to others. Besides the four top commercial banks, American Express Bank also generated significant (93 %) volatility connectedness to others.

The difference between the total directional connectedness to others and the total directional connectedness from others gives the net total directional connectedness to others. In terms of the net total directional connectedness Citigroup (26.5 %) leads the way, followed by Bank of America (18.8 %), American Express Bank (13 %), and J. P. Morgan (8.9 %). AIG (-19 %), PNC Bank (-18 %), Fannie Mae (-17 %), Goldman Sachs (-15 %) and Bank

of New York Mellon (BK, -10 %) are the financial institutions with negative values of net total directional connectedness to others.

Finally, with a value of 78.3 % the measure of total connectedness among the thirteen financial stocks is higher than the total connectedness measures we obtained in other settings, such as the connectedness among different asset classes, or among international stock markets. Given the large number of stocks included in the sample, there is a high degree of connectedness for the full sample. As we will see below there is always high degree of connectedness even during tranquil times. In the case of a set of financial stocks there is another reason why we would expect the connectedness index to be high compared to the connectedness index for a set of major national stock markets around the world and for a set of asset classes in the U.S. As the institutions included in our analysis are all operating in the finance industry, both industry-wide and macroeconomic shocks affect each one of these stocks one way or the other. As some of these institutions are more vulnerable to external or industry-wide shocks than others, they are likely to be transmitting more of these shocks to other financial institutions, generating a higher degree of connectedness to others. Obviously, to the extent that they have important implications for the rest of the industry, idiosyncratic volatility shocks are also transmitted to other stocks. For that reason, compared to a similar number of stocks from different industries, the connectedness indexes for a group of stocks in the finance industry are likely to be higher. It is also likely to be higher compared to the connectedness index for a group of global markets, as these markets are not subject to common shocks as frequently as the stocks from the same industry and especially from the finance industry.²⁴

5.3 Dynamic (Rolling-Sample, Conditional) Analysis

The just-completed analysis of full-sample connectedness provides a good characterization of “average” or “unconditional” aspects of each of the connectedness measures, yet by construction it is silent as to connectedness *dynamics*. In this sub-section we provide a dynamic analysis by using rolling estimation windows. We include the same thirteen financial institutions that we included in our earlier full-sample analysis.²⁵ We start our dynamic analysis with total connectedness, and then we move to various levels of disaggregation (directional

²⁴We have in mind a comparison with the total connectedness indexes reported in Diebold and Yilmaz (2009) and Diebold and Yilmaz (2011).

²⁵In the next sub-section we specifically focus on the 2007-08 financial crisis and include the remaining four institutions (Bear Stearns, Lehman Brothers, Merrill Lynch and Wachovia), all of which ceased trading during the crisis.

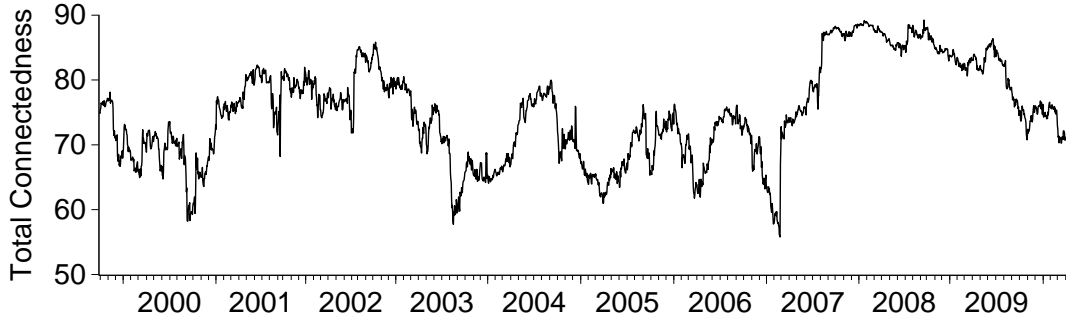


Figure 2: Rolling Total Connectedness. The rolling estimation window width is 100 days, and the predictive horizon for the underlying variance decomposition is 12 days.

and pairwise). We also provide a brief assessment of the robustness of our results to choices of tuning parameters alternative identification methods.

5.3.1 Total

In Figure 2 we plot total volatility connectedness over 100-day rolling-sample windows. From a bird's-eye perspective, the total connectedness plot in Figure 2 has some revealing patterns. It has two big waves; one starting in late-2000 and ending in mid-2003, whereas the second coincides with the development of the global financial crisis from early 2007 all the way to the end of 2009. The first big wave coincides with the burst of the dot-com bubble, followed by the downward spiral in Nasdaq and other stock exchanges and the 2001 recession. Even if the recession was over in early 2002, the MCI/World Com scandal of mid-2002 kept the volatility of the financial stocks and their connectedness high for another year. With the first signs of the sub-prime crisis, the total volatility connectedness index jumped up from a low of around 56% in February 2007 to reach all the way close to 90% in August 2007 and stayed above 80 % until mid-2009.

In between the two big waves of the total connectedness lie three smaller, but not necessarily negligible, waves. We will discuss each of these waves of the total connectedness plot, along with the events that possibly led to these movements. Before doing so, let's us point to another observation one can make in the total connectedness plots. From 1999 to 2007, whenever the total connectedness increased to reach a higher level reflecting higher volatility of some or all financial stock returns, it always came back down to the 55-70 % range as the sample-windows are rolled to leave that episode behind. Following the 2007-08 financial

crisis, as of the end of April 2010, the total connectedness measure stayed well above this range, even though the financial crisis ended almost a year ago.

Earlier on in our sample, the behavior of the total connectedness is affected by the developments in the tech-heavy Nasdaq composite stock index. In the late 1990s US stock markets experienced rapid price increases fueled in greater part by speculative demand for the internet stocks (also called dot-com). In the second half of 1990s, Nasdaq composite stock index increased more than fivefold, from around 1,000 in January 1996 to more than 5,000 in March 2000. Starting in March 2000, the dot-com bubble finally started to burst. After closing at a historical high of 5048 on 10 March 2000, Nasdaq composite index dropped all the way down to 3000 in two months. The drop in Nasdaq index had an impact on the performance of other stocks. As such, it also had some impact on the total connectedness. In March 2000 the index increased slightly (app. 7 percentage points). As the downward move of the Nasdaq composite index subsided temporarily and stabilized around 4000, the total connectedness measure started to subside down during the summer of 2000.

Despite short spells of recovery, troubles of the internet stocks continued for some time. As a result, the rapid decline in the Nasdaq composite index set in again in late September. Nasdaq index that was around 4000 in mid-September 2000 dropped all the way down to below 2000 by mid-March 2001. As the problems in the tech stocks worsened, they started to spread to the rest of the stock market and the economy. Solid signs of an imminent recession appeared in the horizon. The volatility in the bank stocks increased rapidly over this period, and so did the total connectedness. From a low of 60 % in early-September, the connectedness index increased to 75 % by mid-January 2001 and further to above 80 % by early May 2001.

The Federal Reserve's intervention, by way of lowering the Fed funds target rate by 2.5 percentage points in the first five months of 2001, helped stem the decline in the Nasdaq and other markets towards the second and third quarters of 2001. Total connectedness subsided all the way down to 71 % by early September. However, 9/11 terrorist attacks worsened the market sentiment again. Even though the markets were closed for a week after the terrorist attacks, the total connectedness among the financial stocks jumped 10 percentage points in the week it was reopened. The total connectedness stayed around 80 % as long as the data for 9/11 were included in the rolling-sample windows.

As the volatility of financial stocks stayed high in the first half of 2002, the total connectedness stayed unchanged slightly above 75%. After the Enron scandal of late 2001, another corporate scandal rocked the U.S. financial markets towards the end of June 2002. This time

around it was the bankruptcy of MCI WorldCom, which was once the second-largest long distance phone company in the U.S. As the news about the scandal were revealed WorldCom's stock price fell drastically and on July 21 the company had to file for bankruptcy, which at the time happened to be the third largest corporate bankruptcy in U.S. history.

Unlike the Enron scandal, MCI WorldCom scandal had an impact on major bank stocks. All major U.S. banks had credit positions with MCI WorldCom, hence when the company declared bankruptcy they all suffered losses. In addition, 17 commercial banks including Citigroup, J.P. Morgan and other major banks ended up paying a total of \$6 billion dollars to investors for underwriting WorldCom bonds. However, being an isolated source of loss for the banks, the scandal's impact on the financial system as a whole could be contained.

Following the bankruptcy, the total connectedness among the major financial institutions jumped from 72% to reach 85% in July 2002, the highest level achieved from the beginning of the sample. However, the scandal involved a single company only, by the end of 2002 total connectedness subsided very quickly to pre-July 2002 levels. By early-2003 the total connectedness declined to below 70 %. After a brief increase following the invasion of Iraq in May 2003, the total connectedness declined to 58 % in August 2003.

From 2003 all the way to 2007 the total connectedness fluctuated between 60 and 80%, going through 3 episodes of tighter connectedness. The first upward movement in the total connectedness took place in early to mid-2004, as the Fed signaled a change in its policy stance and decided to increase the policy rate from June 2004 onwards. Total connectedness among the bank stocks increased from 65 % at the end of 2003 all the way to almost 80 % by August 2004.

Finally, following the Federal Open Market Committee's May 2006 decision to increase Fed funds rate target by 25 basis points in May 2006, accompanied with a statement that announced that there will be further increases in the Fed funds rate in the future, international stock markets and especially the emerging market stock markets were substantially affected. The increased volatility in the U.S. stock markets are also reflected in the financial stocks. Total connectedness among the thirteen financial stocks increased by approximately 12 percentage points during this episode.

5.3.2 Total Directional

In Figure 3 we show time series of total directional connectedness ("to" and "from" degrees) separately for each firm. The total directional connectedness "to" others plots are presented in the upper panel, the total directional connectedness "from" others plots are in the middle

panel, and the “net” total directional connectedness to others plots are in the lower panel.

Looking at Figure 3 the first thing that one notes is the substantial difference between the “To” and “From” connectedness plots. The “From” connectedness plots are much smoother compared to the “To” connectedness plots.

The difference between the two directional connectedness measures is not hard to explain. When there is a shock to the return volatility of an individual stock or a couple of stocks, this volatility shocks are expected to be transmitted to other stocks. Since individual institutions’ stocks are subject to idiosyncratic shocks some of these shocks are very small and negligible, while others can be quite large. Furthermore, irrespective of the size of the shock, if it is the stock of a larger institution that received the volatility shock, we can expect this volatility shock to have even a larger spillover effect on stocks of other institutions. As the size of the shocks vary as well as the size of the institutions in our sample, the directional connectedness “to” others varies substantially across stocks over the rolling-sample windows.

We have already emphasized that the institutions in our sample are the largest ones in the U.S. financial industry. As a result, none of the stocks in our sample of thirteen institutions would be insulated from volatility shocks to stocks of other institutions. In other words, they are expected to be interconnected. As a result, each one will receive, in one form or the other, the volatility shocks transmitted by other institutions. While the volatility shocks transmitted “to” others by each individual stock may be large, when they are distributed among twelve other stocks the size of the volatility shock received by each stock will be much smaller. That is why there is much less variation in the directional connectedness “from others” compared to the directional connectedness “to others” in Figure 3.

The difference between the directional connectedness “to” and “from” others is equal to the “net” directional connected to others presented in the lower panel of Figure 3. As the connectedness “from” others measure is smooth over the rolling-sample windows, the variation in the “net” connectedness to others plots over the rolling-sample windows resembles the variation in the connectedness “to” others plots.

When we focus on the behavior of the directional connectedness measures over time, we observe that even though “from” others measures for each stock reached the highest levels during the 2007-08 crisis, we do not observe such a level shift in the “To” other and “net” to others measures during crisis period. This is so, perhaps because idiosyncratic shocks have always hit individual stocks and these shocks have been transmitted to other stocks. During the 2007-08 crisis these shocks tend to become larger was are transmitted

In the initial phase of the sub-prime crisis (February-July 2007), as the mortgage compa-

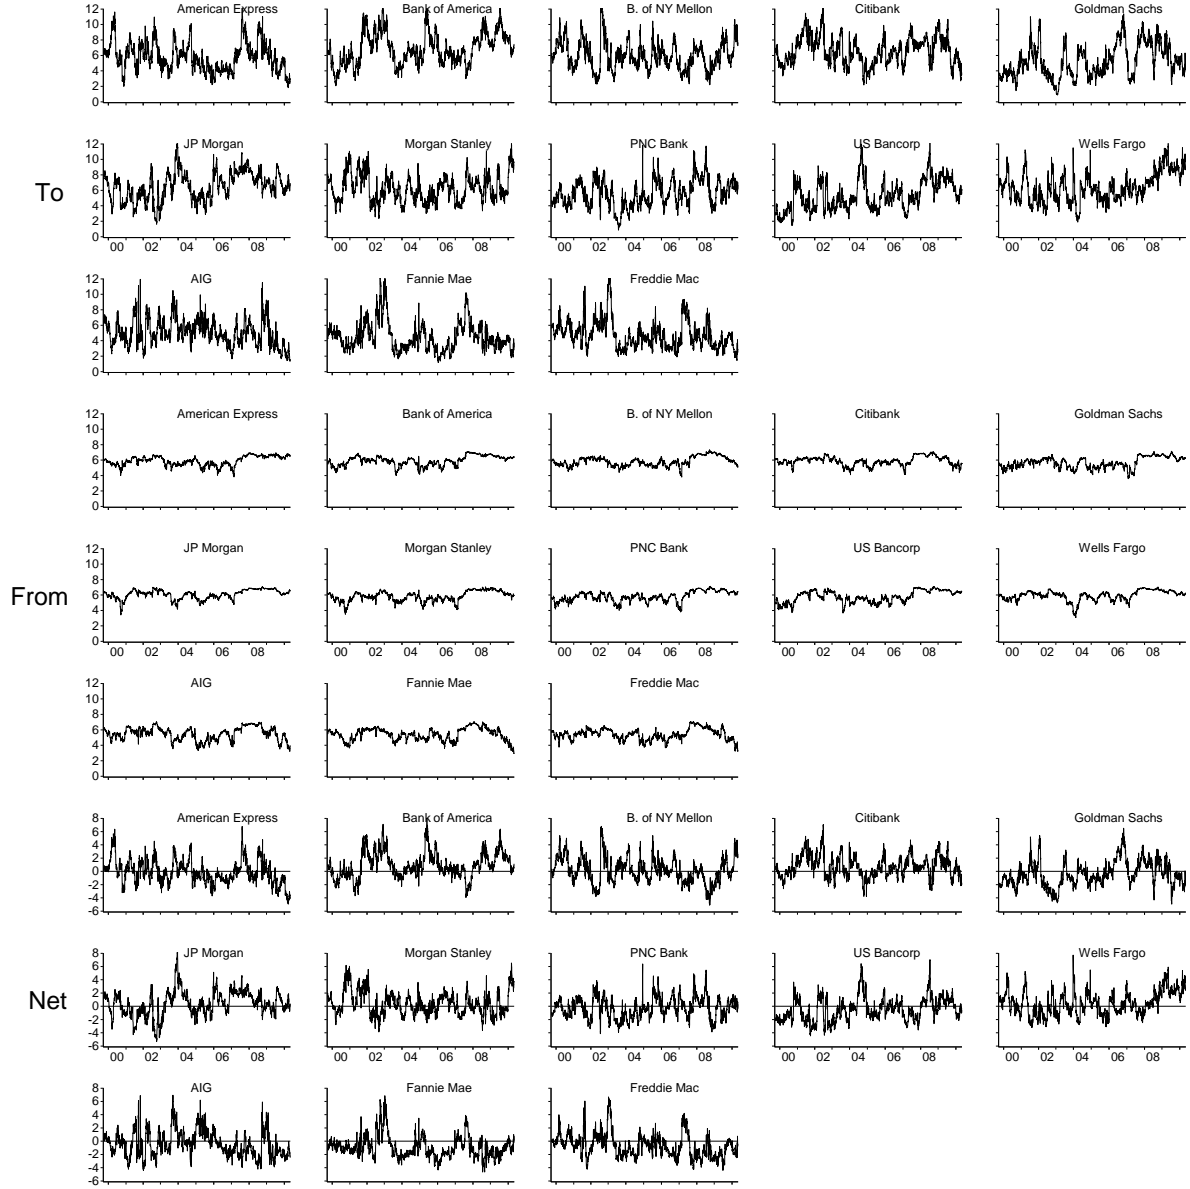


Figure 3: Rolling Total Directional Connectedness. The rolling estimation window width is 100 days, and the predictive horizon for the underlying variance decomposition is 12 days.

nies start to collapse the mean "to others" connectedness increased (late February 2007) and the distribution became a more fat-tailed. In the spring 2007 (April-May) the distribution widened further, followed by the summer lull and a tighter distribution until the end of July. With the sub-prime crisis turned into an international liquidity crisis at the end of July, the mean-degree increased and the distribution became wider, indicating a differentiation taking place among the stocks of major financial institutions. Lehman Brothers was closer to become an outlier in August 2007. Similarly, signs of trouble in Wachovia Bank and Merrill Lynch were already there as of the summer of 2007. In addition to these three banks, American Express Bank had very high (ranked among the top four) directional connectedness "to others". While American Express Bank was not in big trouble, it had a weaker balance sheet than many commercial banks in our sample, because most of its loans were high risk consumer credit loans financed not by deposits but rather by long-term borrowings.

As the liquidity crisis subsided down the distribution of the "to others" connectedness became tighter without any decline in the total connectedness. All major banks had to accept their untenable positions and searched ways to raise capital even though that would have diluted the incumbent shareholders' interests. In the first few weeks of 2008 it was JP Morgan stock that had high "to others" volatility connectedness. The problems of Bear Stearns had rather negligible impact on other financial institutions, as can be seen in low values of gross and net pairwise connectedness measures Bear Stearns had with other stocks. The imminent collapse of Bear Stearns and its implications for the greater market were thwarted by its last minute sale to the JP Morgan Chase in a deal engineered within the guises of the Federal Reserve Bank of New York. Even though the slight increase in the median connectedness showed the stress among several banks, especially among the investment banks, the handling of the Bear Stearns within the system calmed the markets for a while.

Thus far we have focused on time series of individual node degrees. In closing this section, however, we now show in Figure 4 the evolution of the entire "to" and "from" degree *distributions*. Although, by definition, the *mean* "To" and "From" directional connectedness measures are both equivalent to the total connectedness measure presented in Figure 2, each financial institution has rather different "To" and "From" directional connectedness. This implies that even though their means are the same, "To" and "From" connectedness are distributed quite distinctively. As emphasized earlier, the variation in the "From" connectedness is much lower than the variation in "To" connectedness. Even the first and second quartile band for the "To" connectedness is wider than the min-max range for the "From" connectedness.

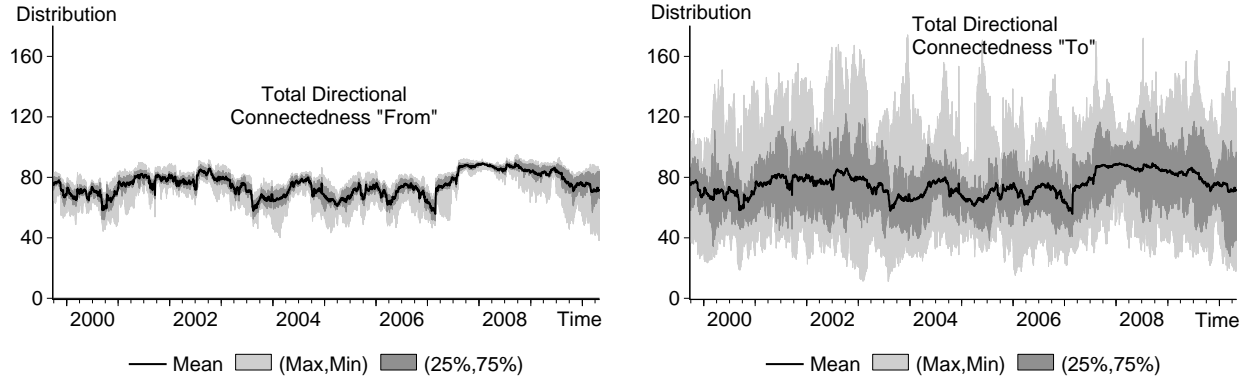


Figure 4: Rolling Distribution of Total Directional Connectedness. We plot the time series of daily min, 25%, mean, 75%, and max of the distributions of “to” and “from” total directional connectedness. The rolling estimation window width is 100 days, and the predictive horizon for the underlying variance decomposition is 12 days.

Temporal changes in the dispersion and skew of “To” and “From” connectedness may contain useful information. For example, it appears that “From” connectedness gets not only more dispersed but also more left-skewed during crises, and simultaneously that “To” connectedness gets more right-skewed. That is, during crisis times relatively more than non-crisis times, there are a few firms receiving very little, and a few firms transmitting very much. One might naturally want to identify firms that are simultaneously “small receipts” and “big transmissions” – those are the distressed firms potentially poised to wreak havoc on the system.

5.3.3 Pairwise Directional

In the analysis of the full-sample volatility connectedness in Section 5.2, we discussed the importance of pairwise volatility connectedness measures. In particular, we emphasized the importance of pairwise connectedness as a measure of how volatility shocks are transmitted across financial institution stocks. The relevance of the pairwise connectedness measures carries over to the rolling sample windows. Indeed, the analysis of pairwise connectedness measures are even more crucial in the rolling sample windows case, because it helps us identify how the connectedness across financial institution stocks vary over time. During times of crises, individual stocks are likely to be subject to frequent volatility shocks. How these shocks led to volatility connectedness across pairs of stocks is very crucial for any

analysis of crises. Unfortunately, given that there are 13 institutions in our sample from 1999 to 2010, presenting plots of the volatility connectedness (for each of the 156 pairwise directional measures, and 78 net pairwise directional measures) is an almost impossible task to accomplish in the confines of this article. Instead, when we are discussing the development of the global financial crisis over time and the volatility connectedness of the most troubled financial institutions during the crisis, we will present and discuss the net pairwise connectedness measures during the most critical days of the crisis.

5.3.4 Robustness Assessment

Finally, we conclude this section with a discussion of the robustness of our results to the choice of the parameters of the model. In particular, we plot the total connectedness for two alternative identification methods (namely, the Cholesky factor identification and the generalized identification), for alternative values of the window width (in addition to $w = 100$ days, we consider 75 and 125-days long sample windows), and for alternative forecast horizons (in addition to $H = 12$ days, we consider 6 and 18 days). The results are presented in Figure 5. In each plot, the solid line is the total connectedness measure obtained through the generalized identification for each value of H and w . In the case of Cholesky factor identification, we calculate the connectedness index for 100 random orderings of the realized stock return volatilities. The gray band in each plot corresponds to the (10%,90%) interval based on these 100 randomly-selected orderings.

In all subgraphs, the solid-line that corresponds to the generalized identification based total connectedness measure runs higher than the gray band that corresponds to the Cholesky identification. As the generalized identification treats each variable to be ordered as the first variable in the VAR system, the total connectedness obtained from generalized identification is never less than one obtained from the Cholesky-based identification. Nevertheless, in all subgraphs of Figure 5, the two series move very much in accordance over time, a strong indication of the robustness of our total connectedness measures based on generalized identification. It is also important to note that the (10%,90%) interval based on 100 random orderings of the Cholesky-based total connectedness is quite narrow. The ordering of the financial stocks in the VAR do not really matter much to follow the dynamic behavior of total connectedness.

As the window length, w , is increased, the gap between total connectedness based on the generalized identification and the one based on the Cholesky identification increases. Both connectedness measures are more wiggly when the window width is set to 75 days, but

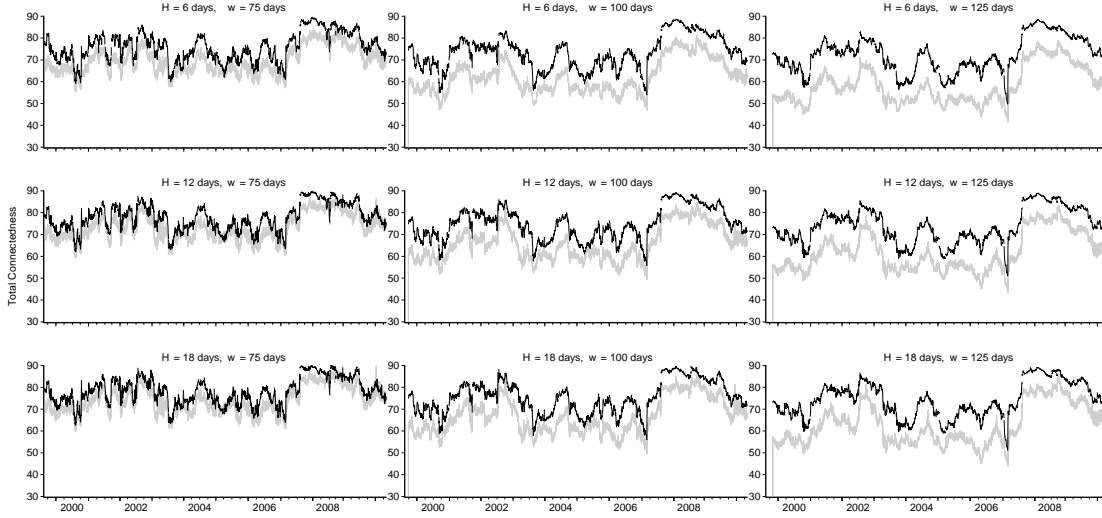


Figure 5: Robustness of Total Connectedness We explore estimation window widths w of 75, 100 and 125 days, predictive horizons H of 6, 12 and 18 days, and a variety of Cholesky orderings. In each sub-graph, the solid line corresponds to our benchmark ordering, and the gray band corresponds to a [10%, 90%] interval based on 100 randomly-selected orderings.

become smoother as we increase the window width to 125 days. Similarly, given the window length, a shorter forecast horizon, H , implies a smaller gap between the generalized- and Cholesky-based total connectedness measures.

To summarize, our robustness checks show that the dynamic behavior of the total connectedness measures over the rolling-sample windows is robust to the choice of alternative sample window lengths, forecast horizons, identification methods and orderings of stocks in the VAR system.

5.4 The Financial Crisis of 2007-2008

Having analyzed the dynamics of the various connectedness measures over time, in this section we focus on the global financial crisis, from 2007 through 2009. First, we discuss the stages of the crisis to show how the total volatility connectedness measure behaved as the crisis developed from the sub-prime and liquidity crises of 2007 in the U.S. into a global financial crisis in October 2008. In the second part of the analysis we focus on the most-troubled financial institutions and their total and pairwise directional volatility connectedness during the critical stages of the crisis.

5.4.1 Total Connectedness at Various Stages of the Crisis

As of the end of 2006 there were already some, albeit weaker, signs of slowdown in the real-estate markets.²⁶ In late-February 2007, the New Century Financial Corporation was reported to have troubles in servicing its debt. It was followed by the bankruptcy of three small mortgage companies. These in turn worsened the expectations about the real estate markets, and the mortgage based assets (MBAs) as well as the stock market. On the last day of February 2007 the total connectedness measure jumped by more than 17 points on a single day. The increase in the total connectedness was not due to a volatility shock to the stock of a single financial institution, rather all bank stocks were affected by the recent developments in the mortgage-based asset markets.

The churning in the MBA market continued from February until early June. New Century declared bankruptcy in April. Following the downgrading of different tranches of mortgages by credit rating agencies, in June and July the markets became aware that big financial institutions were not insulated from the debacle in the MBAs. Bear Stearns had to liquidate two of its hedge funds in July, leading to billions of dollars losses for Bears Stearns and the investors in these funds. From early March to late June the total volatility connectedness index climbed gradually from 73 to 80 % (See Figure 2).

In July 2007, the market for asset backed commercial paper (ABCP) showed signs of drying up market, which eventually led to the liquidity crisis of August 2007. As From July 25th to August 10, the index climbed twelve percentage points, to reach 88 % (See Figure 2). Reflecting the developments over the period, the total connectedness index doubled in the first eight months of 2007. After the liquidity crisis of August 2007, it was obvious that the whole financial system will be badly bruised by the collapse of the ABCP market.

After seven months of learning about the problems in MBA markets and the ensuing liquidity crisis, next came the months of reckoning with the consequences as nearly all U.S. banks started to announce huge losses. Many European banks that invested in U.S. MBAs also suffered billions of dollars of losses. In October 2007, some of the worst-affected banks replaced their CEOs and immediately engaged in search for new capital around the world. Even though it has already reached its historical maximum, volatility connectedness index continued its upward move by a couple of points.

Being the weakest of the five major investment banks, Bear Stearns' problems intensified throughout 2007. With the collapse of its two hedge funds the company lost billions of

²⁶The Case-Shiller home price index for 20 metropolitan regions was 2 percent lower in January 2007 compared to its historical high level reached in July 2006.

dollars. As the MBA markets continued their descent Bear Stearns financial position became untenable, amid widespread rumours of an eventual bankruptcy its stock price declined rapidly in mid-March, briefly increasing the tensions and volatility in the markets. In an operation directed by the New York Fed, J.P. Morgan acquired Bear Stearns on 18 March 2007 with financial assistance from the Fed. During the final days of Bear Stearns, the total connectedness for the surviving thirteen banks showed an upward movement of only a couple of percentage points. However, as we will analyze below when we include the Bear Stearns' stock in the analysis we observe an increase in Bear Stearns' net directional connectedness "to others".

There were no major volatility shocks in the stock market after the Bear Stearns' takeover until early June. During this period the total connectedness index declined a couple of percentage points. Throughout the summer of 2008 the tension in the stock market has started to build up again. Wachovia's troubles and its stock's resulting high volatility connectedness were the most important development throughout the summer. As a result of the volatility originated from Wachovia, the total volatility connectedness index increased from 85% to 88.5% in the first two weeks of July 2008 (See Figure 2).

In the meantime, there were failures of regional banks smaller than Wachovia. These were followed by news about constantly deteriorating asset positions of Fannie Mae and Freddie Mac. Before they went bankrupt, the two GSEs were taken to the explicit government conservatorship in the first week of September.

Then came the most significant event in the unfolding of the crisis. Following the news that Lehman will announce huge losses in its latest financial statement, market participants started selling Lehman Brother stocks. Despite the overwhelming efforts over the weekend of 13-14 September, no viable takeover bid could be produced for Lehman Brothers by the interested institutions. The U.S. government did not want to step in to save the Lehman Brothers with taxpayer money. As soon as the Lehman Brothers declared bankruptcy in the morning of 15 September 2008 the hell broke out in financial markets around the world. The same day the weakest of the three remaining investment banks, Merrill Lynch, announced it was being acquired by the Bank of America. The total volatility connectedness index increased further to reach its maximum level of 89.2 % (See Figure 2).

The government had to step in when the crisis spread to AIG through the credit default swaps (CDS) it issued on Lehman's debt. The U.S. Secretary of Treasury Hank Paulson was criticized for saving the AIG while letting the Lehman Brothers go bankrupt. Saving Lehman Brothers would have cost much less. However, it was also true that, had the AIG

Institution	Ticker	Business	Market Cap. (12/29/2006)	Important Events	
Bear Stearns	BSC	I-Bank	19	Acquired by JPM	3/17/2008
Lehman Brothers	LEH	I-Bank	41	Bankruptcy	9/15/2008
Merrill Lynch	MER	I-Bank	82	Acquired by BAC	9/15/2008
Wachovia Bank	WB	C-Bank	115	Acquired by WFC	10/3/2008

Table 4: Detail for Financial Institutions Acquired or Bankrupted During the Crisis of 2007-2008. I-Bank denotes an investment bank, and C-Bank denotes a commercial bank. Market capitalizations are in billions of U.S. dollars.

gone under, no one could have figured out the consequences for the U.S. and the global financial system.

After months of gyrations in the U.S. financial system, the volatility connectedness started to subside down towards the end of the first quarter of 2009. After March 2009, total connectedness measure fluctuated between 80-85 % for a while. It started to fall only in the summer of 2009. By October 2009 the index was below down to the 70-75% range. However, the news coming from Greece and the EU's inability to handle the Greek debt crisis in an orderly manner led to further volatility in financial industry stocks in the EU and the U.S., which prevented the volatility connectedness index to decline any further. As of the end of our sample, the index was fluctuating between 70 and 75 %, a range which is above the levels the index attained during tranquil times (See Figure 2).

5.4.2 Pairwise Connectedness of Troubled Financial Institutions

So far we have discussed the behavior of the total connectedness and total directional connectedness measures for a group of thirteen institutions along with the background of the events that took place in the U.S. financial markets during the financial crisis of 2007-2008. Our analysis did not include four major banks that disappeared during the crisis through bankruptcy or acquisitions. In the remainder of this section, we analyze the total directional and pairwise directional connectedness measures for these four institutions as well as for AIG and Morgan Stanley, two other troubled yet survived institutions. In Table 4 we list the information on the four major banks that cease to exist, with information on their stock tickers, market capitalization before and after the crisis, and critical dates during the crisis.

In this section, we first present total directional connectedness plots for AIG, Wachovia, Merrill Lynch, Lehman Brothers, Morgan Stanley and Bear Stearns in Figure 6 and an-

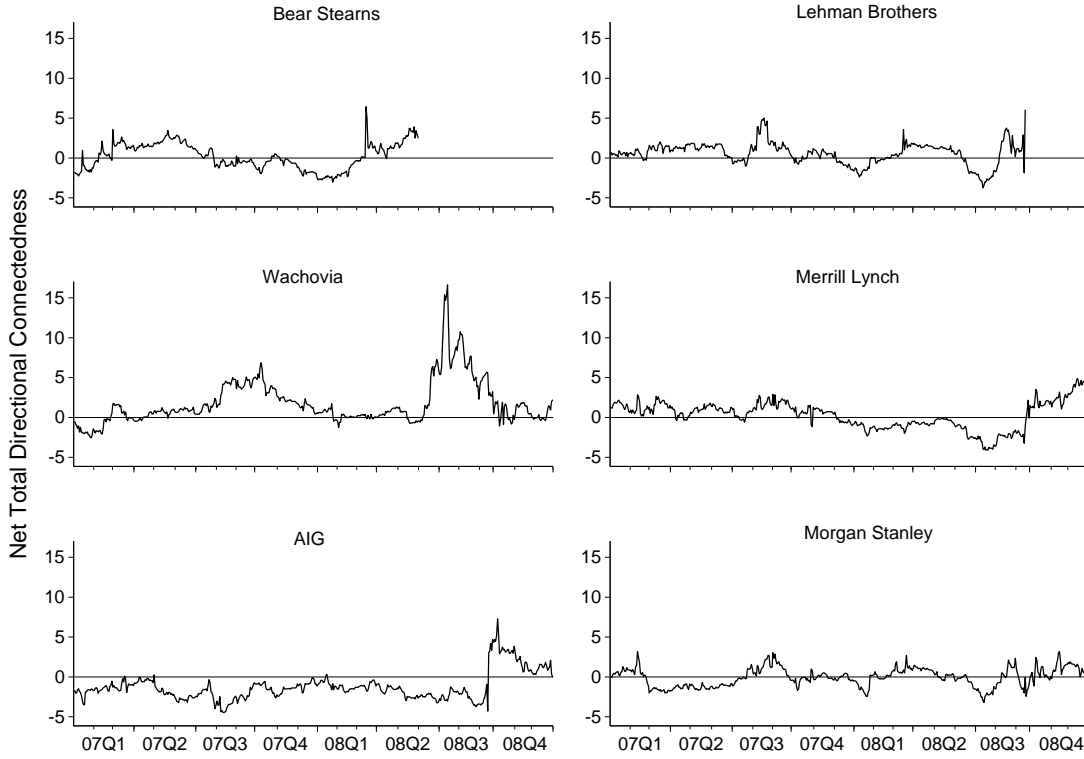


Figure 6: Net Total Directional Connectedness of Troubled Financial Firms. Net total directional connectedness for firm i is “from i to others” less “from others to i .”

analyze the dynamics of their net total directional connectedness with the stocks of other institutions during the financial crisis.²⁷ Let us spell out the most important observation in Figure 6: Even though it was the troubles of the investment banks that were followed the most throughout the crisis, Wachovia Bank is the one that had the highest net total and pairwise volatility connectedness in the climactic months of the second half of 2008.

Coming back to the four troubled investment banks, it was true that they had high net connectedness at several occasions as the global financial crisis unfolded steadily in 2007 and 2008. To start with the most vulnerable of the top 5 investment banks, in the run up to

²⁷It is worth noting that connectedness measurements generally will not, and should not, be robust to choice of reference universe. Hence, given a decision as to the x to be examined, a second important issue is precisely which (and hence how many) x 's to use. For example, in this paper's analysis of individual financial institution equity return volatilities, we intentionally use only the largest firms. In addition, note that our reference universe will change with “births” and “deaths” of financial firms. Births happen, for example, when a firm goes public as with Goldman Sachs in 1999, and deaths happen when firms go bankrupt as with Lehman Brothers in 2008.

its takeover by J.P. Morgan on 17 March 2008 the net volatility connectedness of the Bear Stearns stock was not sizable. Bear Stearns' net volatility connectedness was high only on Friday, March 14 (close 6.4 %) and Monday, March 17 (4.9 %). (See Figure 6). As we have already discussed above, Bear Stearns' net volatility connectedness in March through June 2007 was higher. However, in the three months prior to its demise its net total directional connectedness was negative, indicating that it was on the receiving end of the volatility shocks from other stocks.

Viewed as the most vulnerable investment bank after Bear Stearns, Lehman Brother's net directional connectedness during the liquidity crisis of 2007 was close to 5 percent. It also generated close to 4 percent net directional connectedness on the day Bear Stearns was taken over by J.P. Morgan (Figure 6). Furthermore, its net directional connectedness stayed around 2 percent for almost three months after the demise of Bear Stearns. As the news about Wachovia's troubles dominated the market from early June till early August and Lehman Brothers stayed as a net receiver of volatility shocks. This status, however, did not last for long. Lehman again became one of the front runners in terms of net directional connectedness (close to 4 %) in the first 20 days of August. As the focus shifted to Fannie Mae and Freddie Mac's troubles and their being undertaken to the government conservatorship, the net volatility transmission by Lehman declined in the first week of September.

On Friday, 12 September 2008, just one day before the critical weekend, Lehman Brothers was not at the center stage in terms of volatility connectedness; its net total directional volatility connectedness was less than one percent (see Figure 6). Its net pairwise connectedness with none of the other financial stocks was significant enough to make to the top ten percentile of all the net pairwise volatility connectedness took place between June 1 and December 31 of 2008 (see the stringball plot in Figure 7(a)). Only after the announcement of its bankruptcy in the morning of September 15, the Lehman Brothers' stock moved to the center stage in the crisis and generated substantial volatility connectedness. Its net total directional connectedness jumped to 6 % on September 15 (Figure 6). Its net pairwise connectedness with 5 financial stocks were in the top one (another five were in the top five percentile and two in the top ten percentile) of all the net pairwise volatility connectedness took place between June 1 and December 31 of 2008 (Figure 7(b)). Lehman Brothers' net pairwise directional connectedness increased substantially In the last two trading days of the stock, September 16 and 17 (see Figures 7(c) and 7(d)).

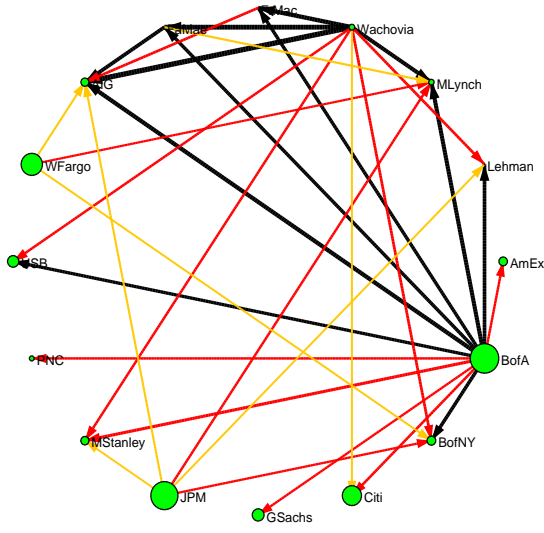
In order to bring forth the bank stocks that played a more central role in terms of the volatility connectedness in the last few days of the Lehman Brothers, in Figures 7(a)

through 7(d) we present the most significant net pairwise directional connectedness measures in an alternative arrangement proposed by Kamada and Kawai (1989). As can be seen in Figure 8(a), while Wachovia and Bank of America were at the center of the “star-shaped” connectedness graph on 12 September 2008, Lehman Brothers had no significant connectedness measure that day. On Monday, September 15, however, the graph changed dramatically. Lehman Brothers and AIG moved to the center of the graph, Lehman having more significant connectedness to others than AIG. The net pairwise directional connectedness of Lehman Brothers and AIG intensified even more on September 16 and 17. In both days, Goldman Sachs also had quite significant net pairwise connectedness to others. As Kamada and Kawai (1989) arrangement-based graphs and the string-ball graphs are quite complementary in terms of their emphases, in the rest of this section we will use two sets of graphs interchangeably.

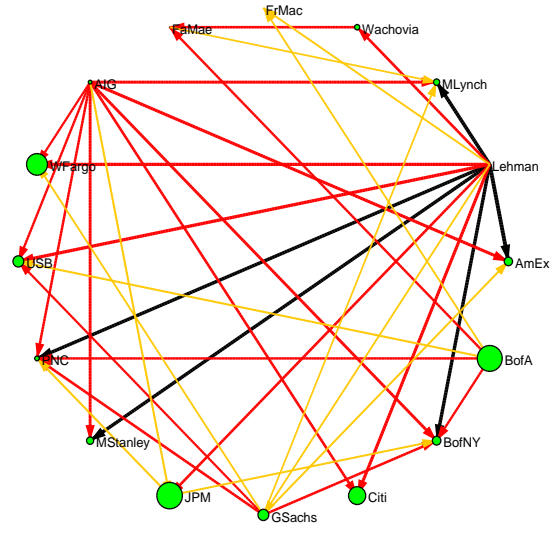
The other investment bank that was having troubles and hence would have definitely headed for bankruptcy after the Lehman Brothers was Merrill Lynch. However, the management of Merrill Lynch was able to sell the whole bank to Bank of America hours before the announcement of the bankruptcy of the Lehman Brothers on September 15, 2008. Merrill Lynch’s net directional connectedness from the beginning of the crisis until the September of 2008 never exceeded 2.5 %. Merrill Lynch’s net total connectedness increased only after it was sold to Bank of America and reached close to 4% in the final days of 2008 before the finalization of the the deal on December 31, 2008 (see Figure 6).²⁸

As we have already emphasized above in Figure 6, among the six troubled banks Wachovia was the one that had the highest net directional connectedness with other stocks. Wachovia’s problems had already been known in 2007. At the heart of its troubles lied its 2006 acquisition of the Golden West Financial, a large California-based mortgage lender specializing in option adjustable-rate mortgages. However, as an increasing number of adjustable-rate mortgage holders were unable to make scheduled mortgage payments, Wachovia’s balance sheet worsened much faster than expected. It made a loss of \$ 8.9 billion in the second quarter of 2008, approximately 80 percent of which was due to the nonperforming loans of the Golden West Financial. The board fired the CEO of the bank on 2 June 2008 as the pressure on the stock started to intensify. From June 11 onwards Wachovia started to transmit substantial net pairwise connectedness to other financial equities. As the markets started to worry about the future of Wachovia, the net volatility transmission by the bank

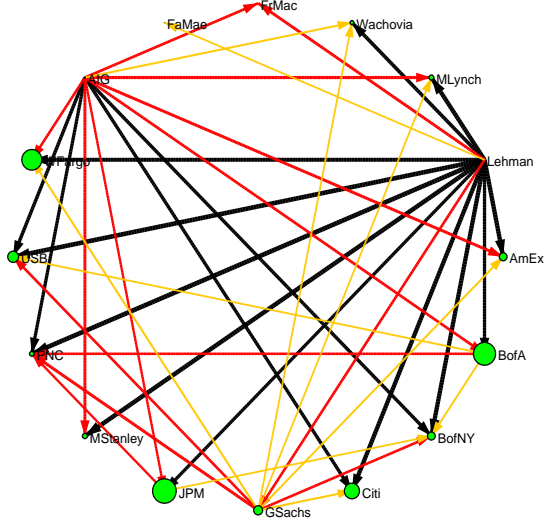
²⁸The uncertainty about the eventual closure of the acquisition process at the end of the year led to an increase in its volatility connectedness.



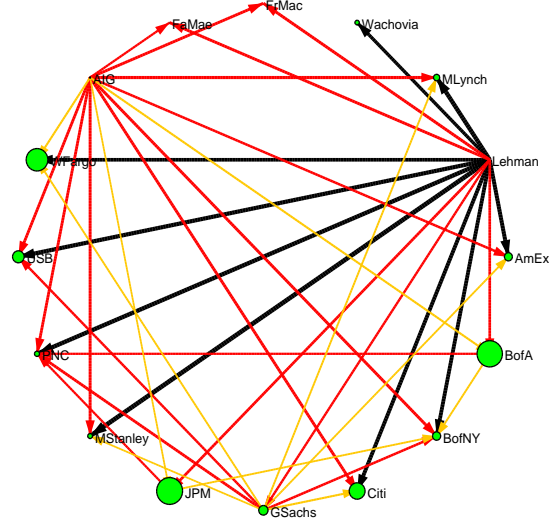
(a) Sep. 12, 2008



(b) Sep. 15, 2008

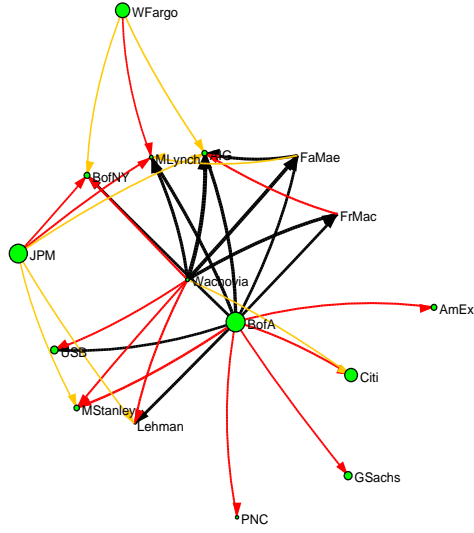


(c) Sep. 16, 2008

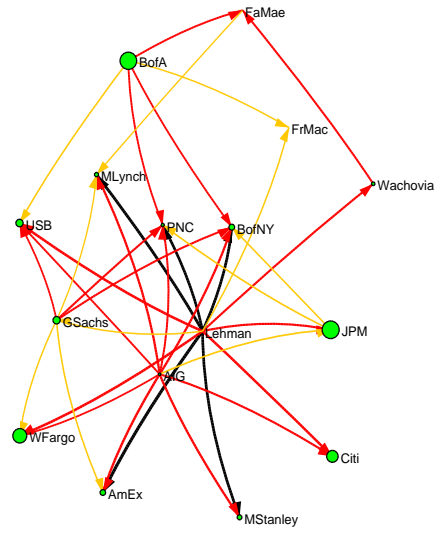


(d) Sep. 17, 2008

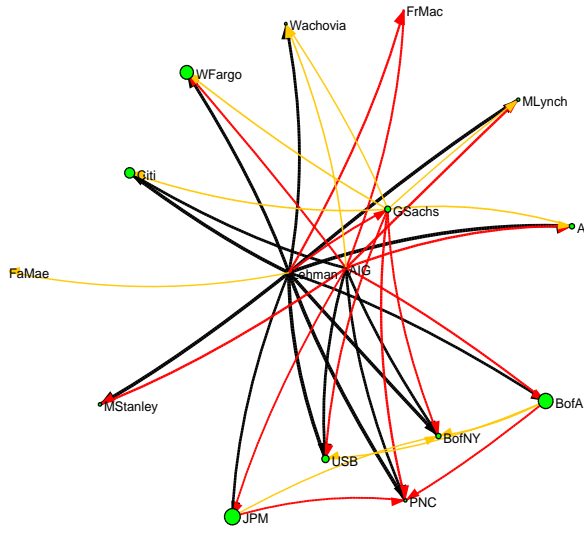
Figure 7: Net Pairwise Directional Connectedness During the Lehman Bankruptcy. Notes: We show the most important directional connections among the pairs of sixteen bank stocks on each day. Black, red and orange links (black, gray and light gray when viewed in grayscale) correspond to the first, fifth and tenth percentiles of all net pairwise directional connections from June 1 to December 31, 2008. Node size indicates stock market capitalization.



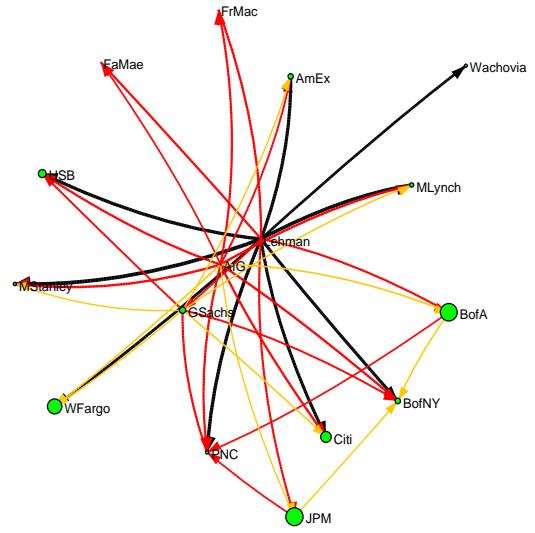
(a) Sep. 12, 2008



(b) Sep. 15, 2008

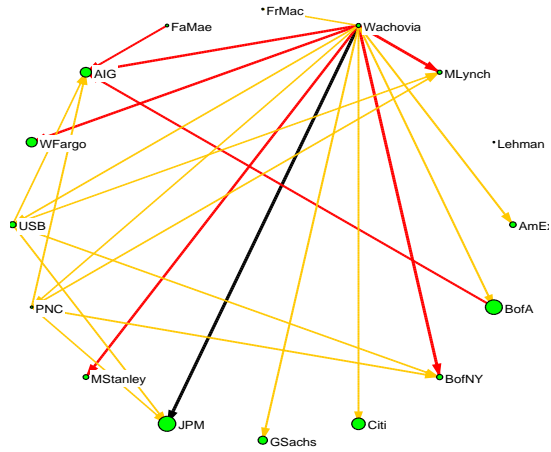


(c) Sep. 16, 2008

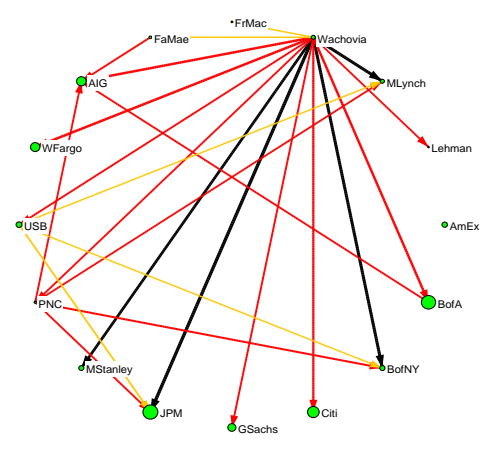


(d) Sep. 17, 2008

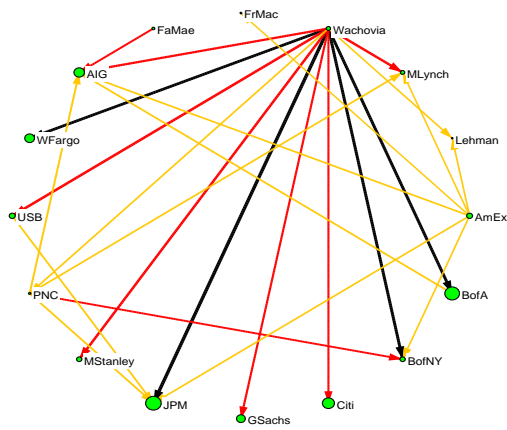
Figure 8: Net Pairwise Directional Connectedness During the Lehman Bankruptcy with Kamada and Kawai (1989) Node Arrangement. See Figure 7 for details.



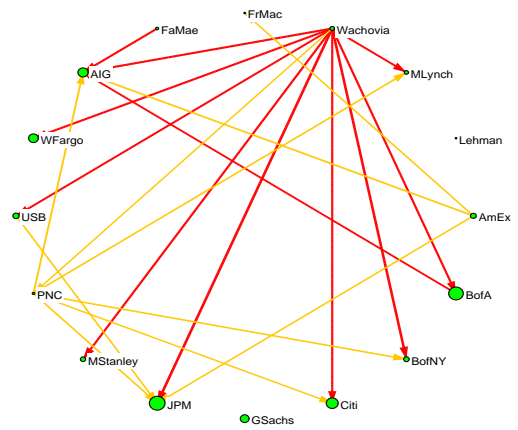
(a) June 18, 2008



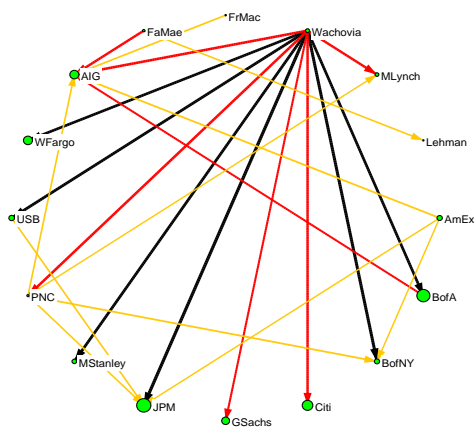
(b) June 19, 2008



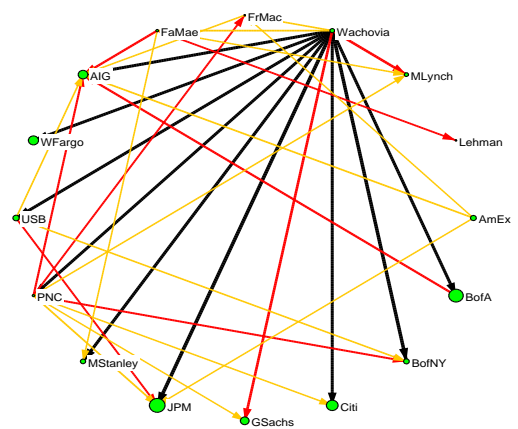
(c) June 20, 2008



(d) June 23, 2008

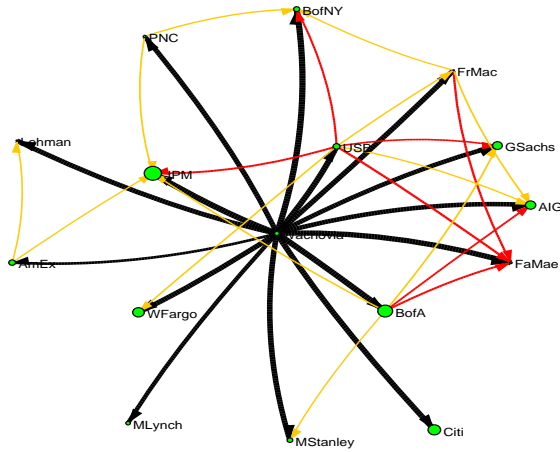


(e) June 24, 2008

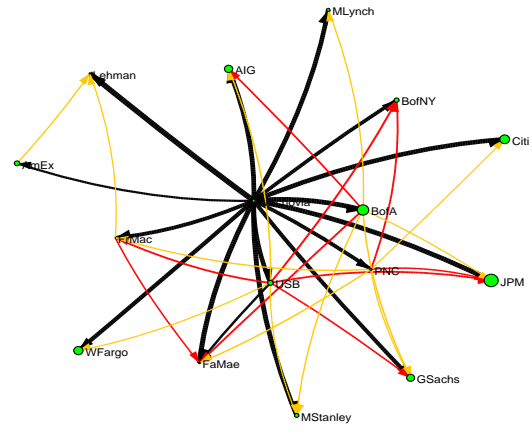


(f) June 25, 2008

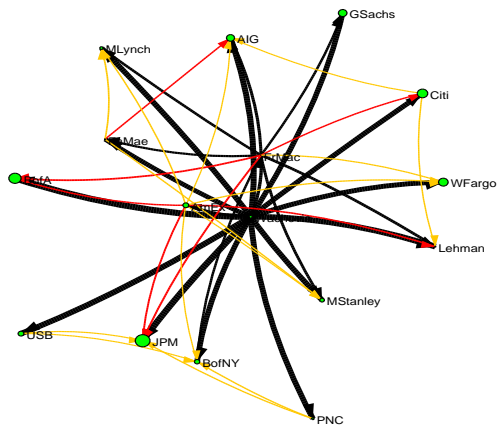
Figure 9: Pairwise Connectedness During June 2008. See Figure 7 for details.



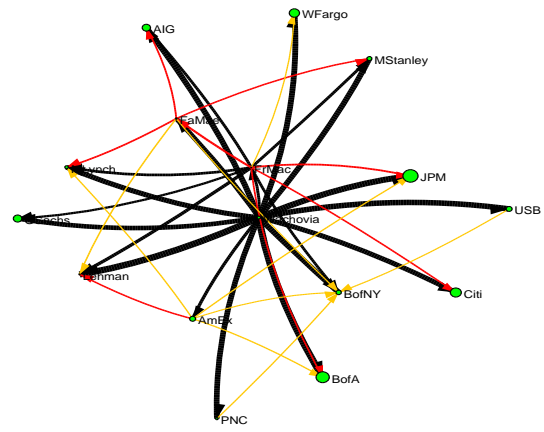
(a) July 9, 2008



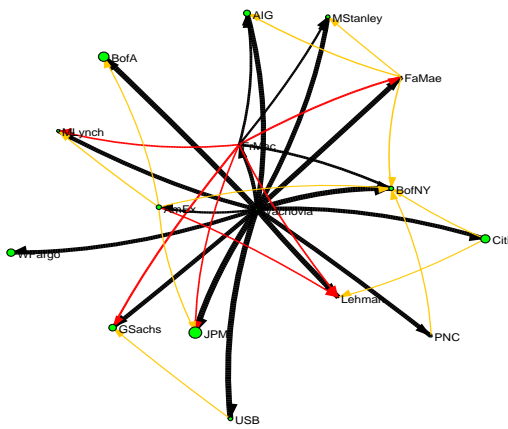
(b) July 10, 2008



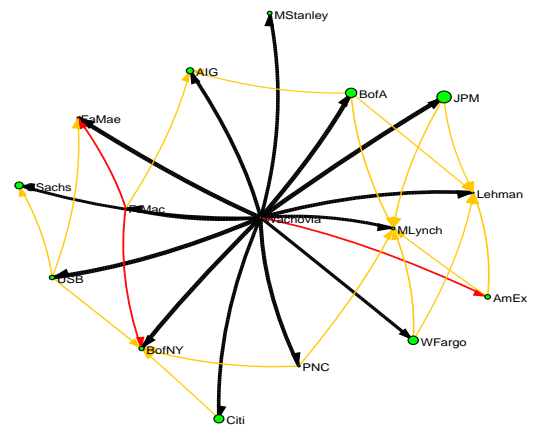
(c) July 11, 2008



(d) July 14, 2008

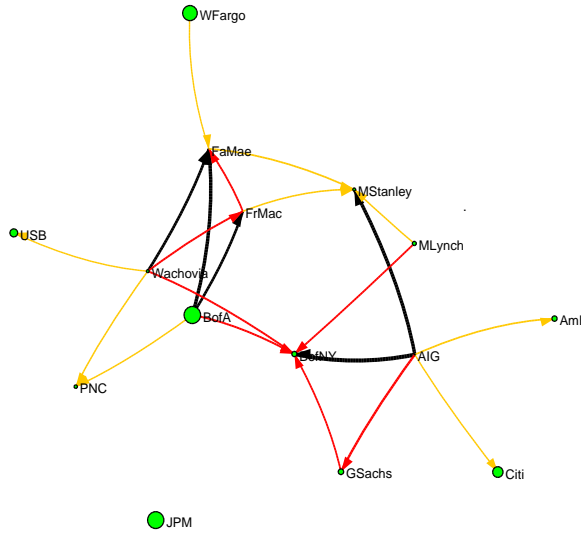


(e) July 15, 2008

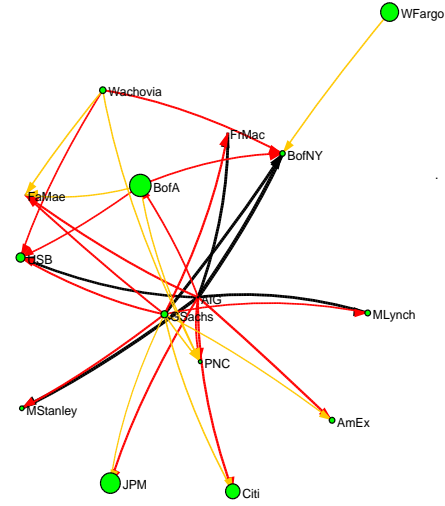


(f) July 16, 2008

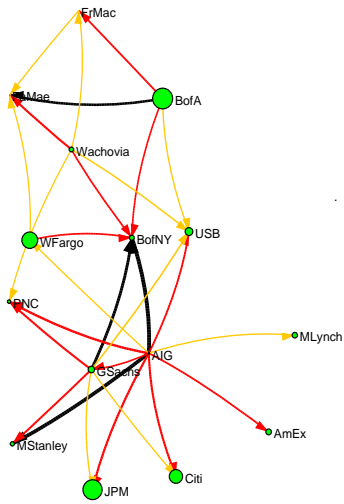
Figure 10: Pairwise Connectedness During July 2008. See Figure 7 for details.



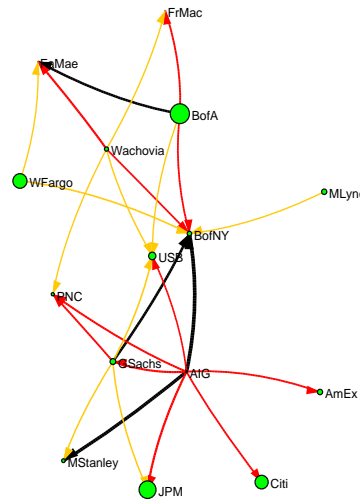
(a) Sep. 18, 2008



(b) Sep. 19, 2008



(c) Sep. 22, 2008



(d) Sep. 23, 2008

Figure 11: Pairwise Connectedness During September 2008. See Figure 7 for details.

indeed spread almost equally to all major banks.

Wachovia's problems got worse towards the end of the month in the second half of June. From June 18 through June 25, 2008, long before the climax month (September 15 – October 15) of the financial crisis, Wachovia's stock came under heavy pressure and its net pairwise connectedness with other financial stocks increased substantially to be ranked in the top one to top ten percentile of all net pairwise connectedness measures in the second half 2008 (See the string-ball plots in Figure 9). Wachovia's problems did not last a couple of days. As the news about the losses mostly related to the mortgages issued by the Golden West Financial continued to arrive, the stock continued to transmit volatility shocks to other bank stocks with increasing intensity, throughout July and August. As can be seen from the solid black colors of each graph, from July 9 through July 16, Wachovia was one of the most active net volatility transmitters among the financial stocks (see Figure 10). Furthermore, even a week before the collapse of the Lehman Brothers, the net pairwise volatility connectedness of Wachovia with other banks were quite significant.

After showing how a critical role Wachovia played during the summer of 2008 with its high directional volatility connectedness, we finally focus on the first few days in the post-Lehman bankruptcy period. The net pairwise volatility connectedness plots for September 18 through 23 show how central was AIG among the remaining stocks. Based on the evidence from the pairwise net volatility connectedness plots, it is possible to assert that the U.S. Treasury was right to prevent the collapse of the AIG after the collapse of the Lehman Brothers. Even though, the markets learned about the injection of \$85 billion (by the U.S. Treasury and the Fed) in the morning of September 16, AIG continued to be at the center of all net pairwise volatility connectedness plots for September 18 through 23 (Figure 11). Its net total connectedness continued to run high from late September through November 2008 (Figure 6).

6 Concluding Remarks

Schweitzer et al. (2009) provide an insightful description of the challenges of financial network modeling:

“In the complex-network context, ‘links’ are not binary (existing or not existing), but are weighted according to the economic interaction under consideration... Furthermore, links represent traded volumes, invested capital, and so on, and their weight can change over time.” [p. 423]

We hope to have successfully confronted the issues raised by Schweitzer et al., proposing connectedness measures at all levels – from system-wide to pairwise – that are rigorous in theory and readily implemented in practice, that capture the different strengths of different connections, and that capture time-variation in connectedness. Our approach effectively marries VAR variance-decomposition theory and network topology theory, recognizing that variance decompositions of VARs form weighted directed networks, characterizing connectedness in those networks, and in turn characterizing connectedness in the VAR.

We see our paper as part of a vibrant emergent literature using network perspectives in economic contexts, and introducing economic perspectives in network contexts. Leading examples include Acemoglu et al. (2010), Adamic et al. (2010), Allen et al. (2010), and Billio et al. (2010). Indeed Billio et al. (2010) is our closest relative, using pairwise Granger-causality to characterize network structure. The Granger-causal approach is in some respects less appealing than ours (e.g., it is directional but exclusively pairwise and unweighted (testing zero vs. nonzero coefficients, with arbitrary significance levels, and without tracking the magnitude of non-zero coefficients), and in other respects more appealing (e.g., there is no need for identifying assumptions, which are inescapable in variance-decomposition and impulse-response analyses), and the two are surely complements rather than substitutes. In any event it seems clear that the network and multivariate time series literatures have much to learn from each other, and that their blending may have much to contribute to the successful measurement of financial economic risks.

References

- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2010), “Cascades in Networks and Aggregate Volatility,” Manuscript, MIT.
- Acharya, V., L. Pedersen, T. Philippe, and M. Richardson (2010), “Measuring Systemic Risk,” Manuscript, Stern School, New York University.
- Adamic, L., C. Brunetti, J. Harris, and A. Kirilenko (2010), “Trading Networks,” Manuscript, University of Michigan, Johns Hopkins University, University of Delaware, and Commodity Futures Trading Commission.
- Adrian, T. and M. Brunnermeier (2008), “CoVaR,” Staff Report 348, Federal Reserve Bank of New York.
- Allen, F., A. Babus, and E. Carletti (2010), “Financial Connections and Systemic Risk,” NBER Working Paper 16177.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2006), “Practical Volatility and Correlation Modeling for Financial Market Risk Management,” In M. Carey and R. Stulz (eds.), *The Risks of Financial Institutions*, University of Chicago Press for NBER, 513-548.
- Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2011), “Risk Measurement for Market Risk Management,” In progress for M. Harris, G. Constantinides and R. Stulz (eds.), *Handbook of the Economics of Finance*, North-Holland.
- Andersen, T.G., T. Bollerslev, F. X. Diebold, and P. Labys (2003), “Modeling and Forecasting Realized Volatility,” *Econometrica*, 71, 579–625.
- Andersen, T.G., T. Bollerslev, and F.X. Diebold (2010), “Parametric and Nonparametric Volatility Measurement,” In L.P. Hansen and Y. Aït-Sahalia (eds.), *Handbook of Financial Econometrics*, North-Holland, 67-138.
- Barabási, A.-L. and R. Albert (1999), “Emergence of Scaling in Random Networks,” *Science*, 286, 509–512.
- Bech, M.L. and E. Atalay (2010), “The Topology of the Federal Funds Market,” *Physica A*, in press.

- Billio, M., M. Getmansky, A.W. Lo, and M. Pelizzon (2010), “Measuring Systemic Risk in the Finance and Insurance Sectors,” Manuscript, MIT.
- Diebold, F.X. and K. Yilmaz (2009), “Measuring Financial Asset Return and Volatility Spillovers, with Application to Global Equity Markets,” *Economic Journal*, 119, 158–171.
- Diebold, F.X. and K. Yilmaz (2011), “Better to Give than to Receive: Predictive Measurement of Volatility Spillovers,” *International Journal of Forecasting*, Forthcoming, with discussion.
- DiNicolò, G. and M. Lucchetta (2010), “Systemic Real and Financial Risks: Measurement, Forecasting, and Stress Testing,” Manuscript, IMF and European University Institute.
- Dufour, J.-M. and E. Renault (1998), “Short-Run and Long-Run Causality in Time Series: Theory,” *Econometrica*, 66, 1099–1125.
- Dufour, J.-M. and A. Taamouti (2010), “Short and Long Run Causality Measures: Theory and Inference,” *Journal of Econometrics*, 154, 42–58.
- Embrechts, P., C. Kluppelberg, and T. Mikosch (1997), *Modelling Extremal Events*, Springer-Verlag.
- Engle, R.F. and B.T. Kelly (2009), “Dynamic Equicorrelation,” Manuscript, New York University.
- Erdős, P. and A. Rényi (1959), “On Random Graphs,” *Publicationes Mathematicae*, 6, 290–297.
- Granger, C.W.J. (2008), “Non-Linear Models: Where Do We Go Next - Time Varying Parameter Models?” *Studies in Nonlinear Dynamics and Econometrics*, 12, 1–9.
- Hansen, P.R. and A. Lunde (2010), “Estimating the Persistence and the Autocorrelation Function of a Time Series that is Measured with Error,” Manuscript, Stanford University and University of Aarhus.
- Härdle, W.K., O. Okhrin, and Y. Okhrin (2011), “Time-Varying Hierarchical Archimedean Copulae,” Manuscript, Humboldt University, Berlin.
- Jackson, M.O. (2008), *Social and Economic Networks*, Princeton University Press.

- Kamada, T. and S. Kawai (1989), “An algorithm for Drawing General Undirected Graphs,” *Information Processing Letters*, 31, 7–15.
- Kelvin, W.T. (1891), *Popular Lectures and Addresses (three volumes)*, MacMillan and Company.
- Koop, G., M.H. Pesaran, and S.M. Potter (1996), “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74, 119–147.
- Lowenstein, R. (2010), *The End of Wall Street*, Penguin Press.
- Newman, M.E.J. (2010), *Networks: An Introduction*, Oxford University Press.
- Pesaran, H.H. and Y. Shin (1998), “Generalized Impulse Response Analysis in Linear Multivariate Models,” *Economics Letters*, 58, 17–29.
- Schweitzer, F., G. Fagiolo, D. Sornette, F. Vega-Redondo, A. Vespignani, and D.R. White (2009), “Economic Networks: The New Challenges,” *Science*, 325, 422–425.
- Sims, C.A. (1980), “Macroeconomics and Reality,” *Econometrica*, 48, 1–48.
- Taleb, N.N. (2007), *The Black Swan: The Impact of the Highly Improbable*, Random House.
- Watts, D. J. and S. H. Strogatz (1998), “Collective Dynamics of ‘Small-World’ Networks,” *Nature*, 393, 440–442.