WORKING PAPERS

RESEARCH DEPARTMENT

ON THE NETWORK TOPOLOGY OF VARIANCE DECOMPOSITIONS: MEASURING THE CONNECTEDNESS OF FINANCIAL FIRMS

Francis X. Diebold
University of Pennsylvania and
Visiting Scholar, Federal Reserve Bank of Philadelphia

Kamil Yilmaz
Koç University

September 30, 2011

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

Francis X. Diebold
University of Pennsylvania
and NBER

Kamil Yılmaz
Koç University

September 30, 2011

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 key measures of connectedness used in the network literature. 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

"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 largely wed to linear Gaussian thinking, making them of limited value in financial-market contexts. 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 elusive.

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 2.4 we relate our measures to some of those in the burgeoning network literatures; the relationships turn out to be direct and important. In section 3 we treat the sample estimation of connectedness, with detailed attention to allowing for its time-variation. Finally, in section 4, we apply our framework to study connectedness at all levels among a large set of return volatilities of U.S. financial institutions during the last decade, including during the financial crisis of 2007-2008. We conclude in section 5.

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 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 on the "non-own," or "cross," variance decompositions, d_{ij}^H, $i, j = 1, ..., N$, $i \neq j$. The key is $i \neq j$.

2.1   The Population Data-Generating Process

Consider an N-dimensional covariance-stationary data-generating process (DGP) with orthogonal shocks: $x_t = \Theta(L)u_t$, $\Theta(L) = \Theta_0 + \Theta_1 L + \Theta_2 L^2 + ...$, $E(u_t u_t') = I$. 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 $\{\Theta_1, \Theta_2, ...\}$. Nevertheless, attempting to understand connectedness via the potentially many hundreds of coefficients in $\{\Theta_0, \Theta_1, \Theta_2, ...\}$ is typically fruitless. One needs a transformation of $\{\Theta_0, \Theta_1, \Theta_2, ...\}$ that better reveals and more compactly summarizes connectedness. Variance decompositions achieve this.

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

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

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 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. Hence we define the *pairwise directional connectedness* from j to i as

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

Note that in general $C_{i \leftarrow j}^H \neq C_{j \leftarrow i}^H$, so 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, as apposed to gross, pairwise directional connectedness. We immediately define *net pairwise directional connectedness* as

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

There are $\frac{N^2-N}{2}$ net pairwise directional connectedness measures, analogous to bilateral trade balances.

Now consider not the individual elements of D^H, but rather its off-diagonal row or column sums. Take the first row, for example. The sum of its off-diagonal elements gives the share of

the H-step forecast-error variance of variable 1 coming from shocks arising in other variables (all other, as opposed to a single other). Hence we call the off-diagonal row and column sums, labeled "from" and "to" in the connectedness table, the total directional connectedness measures. That is, we define *total directional connectedness from others to i* as

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

and *total directional connectedness to others from j* as

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

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. We define *net total directional connectedness* as

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

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 have[1]

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

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

---

[1]Note 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.

(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 measures. 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 ($C^H_{i \leftarrow j}$, for various j), or how *all* others connect to it, $C^H_{i \leftarrow \bullet}$. 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, $C^H_{\bullet \leftarrow j}$, and they might also be more concerned with monitoring total (system-wide) connectedness C^H.

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 vector autoregression (VAR) identifications popularized by Sims (1980), or in any of the scores of subsequent "structural" VAR identifications. Sometimes it is less transparent, but equally true, as in the generalized variance decomposition (GVD) framework of Koop et al. (1996) and Pesaran and Shin (1998), or in more deeply-structural dynamic stochastic general equilibrium approaches like DiNicolo and Lucchetta (2010), which embody large sets of maintained assumptions.

Identifying assumptions are just that – *assumptions* – and 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. Many models, moreover, are exactly- as opposed to over-identified, so that the identifying restrictions cannot be tested. The upshot is that reasonable people may disagree as to their preferred assumptions, and they often do. We have nothing new to add; one must make an assumption and move forward conditional upon (and cognizant of) the assumption.

Our own preferences run toward Cholesky and related identifications such as GVD. 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-factor variance decompositions, GVDs rely on a largely data-based

identification scheme, but they are independent of ordering.[2]

GVDs were introduced in Pesaran and Shin (1998), which builds on Koop et al. (1996).[3] The H-step generalized variance decomposition matrix $D^{gH} = [d_{ij}^{gH}]$ has entries

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

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.[4] 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).[5] Hence we base our generalized connectedness indexes not on D^g, but rather on $\tilde{D}^g = [\tilde{d}_{ij}^g]$, where $\tilde{d}_{ij}^g = \frac{d_{ij}^g}{\sum_{j=1}^N d_{ij}^g}$. By construction $\sum_{j=1}^N \tilde{d}_{ij}^g = 1$ and $\sum_{i,j=1}^N \tilde{d}_{ij}^g = N$. Using $\tilde{D}^g$ 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 $\tilde{C}_{ij}$.

2.4 Relationship to the Network Literature

Networks are everywhere in modern life, from power grids to Facebook. Not surprisingly, research on networks has grown explosively in recent years.[6] A network $\mathcal{N}$ is composed of N nodes and L links between nodes. 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. $\mathcal{N}$ is *connected* if $s_{ij} \leq N-1, \forall i, j$, and one is naturally led to think about measures of the *strength* of network connectedness. That is, presumably two connected networks need not be equally *strongly* connected. But then deep questions arise. Just what *is* strength of network connectedness? Is it a pairwise or system-wide concept, or both, or neither? How, if at all, might it be related to the notion of connectedness that we have proposed independently and thus far emphasized, based on variance decompositions?

To approach the issue of measuring network connectedness, we need to analyze the math-

---

[2]GVDs of course make other assumptions, most notably normality of shock distributions. We will discuss this later in our application to equity return volatilities in section 4.

[3]We refer the reader to the original papers for motivation and background.

[4]Note 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}.

[5]We now drop the "H" superscripts, because from this point onward they are not needed for clarity.

[6]Newman (2010) and Jackson (2008) provide good general and economic introductions, respectively.

ematical structure of networks a bit more deeply. A network is simply an $N \times N$ adjacency matrix A of zeros and ones, $A = [A_{ij}]$, 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. Mathematically (i.e., algebraically), the adjacency matrix A *is* the network, and all network properties are embedded in A. Hence any sensible connectedness measure must be based on A. Nevertheless, there is no single, all-encompassing measure, and several have been proposed.[7] The most important and popular by far – as well as the most useful for our purposes – are based on the idea of node degree (and a closely-related concept, network diameter), to which we now turn.

2.4.1 Degree and Diameter

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

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

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 discrete univariate distribution with support $0, ..., (N-1)$, and aspects of its shape (location, scale, skewness, tail thickness, etc.) are closely linked to aspects of network behavior.[8] As regards the aspect of network behavior that concerns us – connectedness – the *location* of the degree distribution is obviously key, and the standard location measure is of course the mean. Hence the mean of the degree distribution ("mean degree") has emerged as a canonical benchmark measure of overall network connectedness. The larger the mean degree, the greater is overall network connectedness.

The just-described adjacency matrix and degree distribution might more precisely be called "1-step," as the links are direct. 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. The distinction between 1-step and multi-step adjacency emphasizes *distance*. Recall that, as introduced earlier, 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. Distance is a two-node property, in contrast to degree, which is a single-node property. Closely related to the idea of distance is the idea

---

[7] Bech and Atalay (2011) and Adamic et al. (2010) provide good reviews and financial applications of network-theoretic connectedness measures.

[8] The support of $0, ..., (N-1)$ stems from our adoption of the standard convention of writing $A_{ii} = 0, \forall i$.

of diameter. The *diameter* of a network is the maximum distance between any two nodes, $s_{max} = max_{i,j}s_{ij}$. Diameter is another canonical benchmark measure of overall network connectedness. The smaller the network diameter, the greater is overall connectedness.

A beautiful large-N approximation relates network diameter, network mean degree and network size in Erdős-Rényi random networks (Erdős and Rényi (1959)):[9]

$$s_{max} \approx \frac{\ln N}{\ln E(\delta)}.$$

This "network diameter grows only as $\ln N$" approximation is typically introduced as a mathematically-precise characterization of the "small-world" phenomenon, namely that diameters tend to be small even for huge networks.[10] For our purposes, however, it is useful because it emphasizes in a very precise way the centrality of the mean degree as a measure of network connectedness. As we shall now see, our earlier-proposed connectedness measures are intimately related to certain network node degrees and mean degree.

2.4.2 Variance Decompositions as Weighted, Directed Networks

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 aspects of network connectedness. Indeed we are now in a position to notice that *variance decompositions are networks.* More precisely, the variance decomposition matrix D, which defines our connectedness table and all associated connectedness measures, is a network adjacency matrix A. Hence network connectedness measures may be used in conjunction with variance decompositions 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

---

[9]See, for example, Newman (2010), p. 420. Erdős-Rényi random networks have the simplest imaginable probabilistic model of link formation: independent Bernoulli trials with fixed probability θ. Hence the degree distribution is binomial, $f(\delta) = \binom{N-1}{\delta}\theta^\delta(1-\theta)^{N-1-\delta}$. Erdős-Rényi random networks have emerged as a canonical benchmark, but they are sometimes poor descriptions of real-world networks, due for example to strategic aspects of link formation such as clustering, which 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. Interestingly, however, Watts and Strogatz (1998) have shown that the "network diameter grows only as $\ln N$" approximation nevertheless holds in networks with small clusters of linked nodes with just a few long-range links.

[10]For example, for $N = 300,000,000$ (roughly the U.S. population) and mean degree $E(\delta) = 20$, network diameter is still small ($s_{max} \approx 6$). That is, even if every person in the U.S. is linked to only twenty others on average, then the approximate maximum number of steps needed to link any two people is nevertheless only six.

(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 that of the ji link, so the adjacency matrix is generally not symmetric. Third, there are constraints on the row sums of A. In particular, each row must sum to 1 because the entries are variance shares. Hence we write the diagonal elements as $A_{ii} = 1 - \sum_{\substack{j=1 \\ j \neq i}}^{N} A_{ij}$. Note in particular that the diagonal elements of A are no longer 0.

Weighted, directed versions of the earlier-introduced network connectedness statistics are readily defined, including degrees, degree distributions, distances and diameters. For example, node degrees are now obtained not by summing zeros and ones, but rather by summing weights in $[0, 1]$. Moreover, there are now "to-degrees" and "from-degrees," corresponding to row sums and column sums.[11] The from-degree of node i is $\delta_i^{from} = \sum_{\substack{j=1 \\ j \neq i}}^{N} A_{ij}$. The from-degree distribution is the probability distribution of from degrees across nodes. It is a univariate distribution with support on $[0, 1]$. Similarly, the to-degree of node j is $\delta_i^{to} = \sum_{\substack{i=1 \\ i \neq j}}^{N} A_{ij}$. The to-degree distribution is the probability distribution of to degrees across nodes. It is a univariate distribution with support on $[0, N]$.

By now the relationships between our earlier-defined connectedness measures and those used in the network literature should be apparent. First, our total directional connectedness measures $C_{i \leftarrow \bullet}$ and $C_{\bullet \leftarrow j}$ are precisely the from-degrees and to-degrees, respectively, associated with the notes of the weighted directed network D. Second, our total connectedness measure C is simply the mean degree of the network D (to or from – it's the same either way, because the sum of all row sums must equal the sum of all column sums).

3 Sample 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 $C(x, H, A(L))$.[12] In reality $A(L)$ is unknown and must be approximated (e.g., using a finite-ordered vector autoregression). Recognizing the centrality of the approximating model adopted, we write $C(x, H, A(L), M(L; \theta))$, where $M(L; \theta)$ is a dynamic approximating model with finite-dimensional parameter θ. One hopes that $M(L; \theta))$ is in some sense close to the

---

[11] Our to-degrees and from-degrees are often called "out-degrees" and "in-degrees" in the network literature.
[12] The 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.

true population dynamics $A(L)$ for some pseudo-true parameter configuration θ_0, but there is of course no guarantee.

In addition, and crucially, 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 is 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 refers to the *population*, whereas in reality we have available only finite *samples* of observed data. That is, we must use *estimated* approximating models, so we write $\widehat{C}_t(x, H, A_t(L), M(\hat{\theta}_t))$, where the data sample runs from $t = 1, ..., T$. To economize on notation we henceforth drop $A_t(L)$, because it is determined by nature rather than a choice made by the econometrician, relying on the reader to remember its relevance and simply writing $\widehat{C}_t(x, H, M(\hat{\theta}_t))$. In what follows we successively discuss aspects of x, H and $M(\hat{\theta}_t)$.

3.1 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 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 the natural log of realized equity return volatility, the x choice is approximately fifteen major U.S. financial institutions, and the x frequency is daily.[13]

---

[13]Taking logs converts realized volatilities to approximate normality, as emphasized in Andersen et al. (2003).

3.2 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 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 larger for larger H.[14] 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 $H \to \infty$, we obtain an *un*conditional variance decomposition.

In this paper we anchor on a horizon of $H = 12$ days, but 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. Instead we view examination of a menu of H values simply as an interesting part of a phenomenological investigation.

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

A first issue is choice of approximating model class. As discussed previously, many choices are possible, ranging from traditional data-driven VAR approaches, to so-called "structural" VARs, to fully-articulated dynamic stochastic general equilibrium (DSGE) models.

A second issue is how to allow for time-varying connectedness, which is potentially of central interest for risk measurement and management (e.g., over the business cycle, or during financial crises). But connectedness is simply a transformation of model parameters, so allowance for time-varying connectedness effectively means allowance for time-varying parameters in the approximating model.[15] Linear models with time-varying parameters

---

[14]Such 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.

[15]Connectedness may be a highly nonlinear phenomenon, and time-varying parameters are an important

are actually very general nonlinear models, as emphasized in White's Theorem (Granger (2008)).[16]

As with choice of approximating model class, many choices are possible to allow for time-varying parameters. A simple and popular scheme involves use of a rolling estimation window. To track time-varying connectedness in real-time, for example, we might use a uniform one-sided estimation window of width w, sweeping through the sample, at each period using only the most recent w periods to estimate the approximating model and calculate connectedness measures.[17] We write $\hat{C}_t(x, H, M_{t-w:t}(\hat{\theta}))$. The rolling-window approach has the advantages of tremendous simplicity and coherence with a wide variety of possible underlying time-varying parameter mechanisms. Rolling windows do, however, require choice of window width w, in a manner precisely analogous to bandwidth choice in density estimation. In this paper we use a VAR(3) approximating model with a one-sided rolling estimation window of $w = 100$ days.[18]

4   Application to U.S. Financial Institution Return Volatilities, 1999-2010

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 real economic health.

We proceed in four steps. First, in section 4.1, we describe the data that we use to measure financial institution connectedness. Next, in section 4.2, we perform a full-sample

---

way to allow for nonlinearity. 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).

[16]Interestingly, it seems that "White's Theorem" was neither written nor published by White. Instead, White's Theorem appears in Granger (2008). Evidently Granger distilled the theorem from conversations with White, and he named it White's Theorem.

[17]Alternatively, we might use an expanding- as opposed to rolling-sample estimator, $\hat{C}_t(x, H, M_{1:t}(\hat{\theta}))$. Unfortunately, however, expanding sample estimates become very slow to adapt as sample size gets large. Another possibility is explicitly specifying a process for the dynamically evolving model parameters, as is commonly done in a state-space framework using the Kalman filter for estimation.

[18]We also explore robustness to alternative choices of w.

(static) analysis, in which we effectively characterize average, or unconditional, connectedness. This is of intrinsic interest, and it also sets the stage for section 4.3, 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 4.4, we "zoom in" on financial institution connectedness during the global financial crisis of 2007-2008.

4.1 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.

Volatility is latent and hence must be estimated. In this paper we use realized volatility, which has 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.

---

[19]Obviously 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 which may not be robust to violations of those assumptions.

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

We treat realized volatility as the object of direct interest, as in Andersen et al. (2003).[21] 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. This is helpful not only generally, as normality-inducing transformations take us into familiar territory, but also specifically as we use generalized variance decompositions (Koop et al. (1996), Pesaran and Shin (1998)), which invoke normality.

4.2   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

---

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

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

| 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
U.S. Bancorp	USB	C-Bank	64	43
Bank of New York Mellon	BK	C-Bank	30	34
PNC Group	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 into government conservatorship on September 7, 2008, and AIG began government ownership on September 17, 2008.

2007-2008, but also compare and contrast connectedness during other episodes.

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 17% 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. Many features are notable. Clear blocks of high pairwise directional connectedness appear, as with the GSEs (Freddie and Fannie) and various investment banks. The diagonal elements ("own connectednesses") tend to be the largest individual elements of the table, but total directional connectedness (from others or to others) tends to be much larger, and total connectedness is a very high 78%. In addition, the spread of the "from" degree distribution is noticeably less than that of the "to" degree distribution.

Let us discuss some of the features of the connectedness table at greater length, beginning 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

	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).

connectedness measure observed is from Freddie Mac to Fannie Mae ($\tilde{C}^H_{FNM \leftarrow FRE} = 22\%$). In return, the pairwise connectedness from Fannie Mae to Freddie Mac ($\tilde{C}^H_{FRE \leftarrow FNM} = 17.6\%$) 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, $\tilde{C}^H_{FRE,FNM} = 4.6\%$.

The next largest pairwise directional connectedness takes place from Morgan Stanley to Goldman Sachs ($\tilde{C}^H_{GS \leftarrow MS} = 13.3\%$), the two top investment banks that were able to survive the 2007-08 financial crisis.[23] While the connectedness from Goldman Sachs to Morgan Stanley is also high ($\tilde{C}^H_{MS \leftarrow GS} = 9.8\%$), in net terms the directional connectedness takes place from Morgan Stanley to Goldman Sachs stock ($\tilde{C}^H_{GS,MS} = 3.5\%$).

---

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

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 ($\tilde{C}^H_{BAC\leftarrow C} = \tilde{C}^H_{JPM\leftarrow C} = 10.6\%$). A high value of pairwise connectedness from Citigroup to either 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. They are well connected with AIG as well. Pairwise directional connectedness of the stocks of these three institutions with the stocks of other financial institutions tends to be much lower than connectedness of other bank stocks in our sample. We need to remind the reader 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 ranges 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 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 Group, the total directional connectedness from others is distributed rather tightly. The total directional connectedness to others, on the other hand, varies from a low of 53% for Fannie Mae, to all the way up to 106% for Citigroup: A range of 53 points for the connectedness to others compared to a
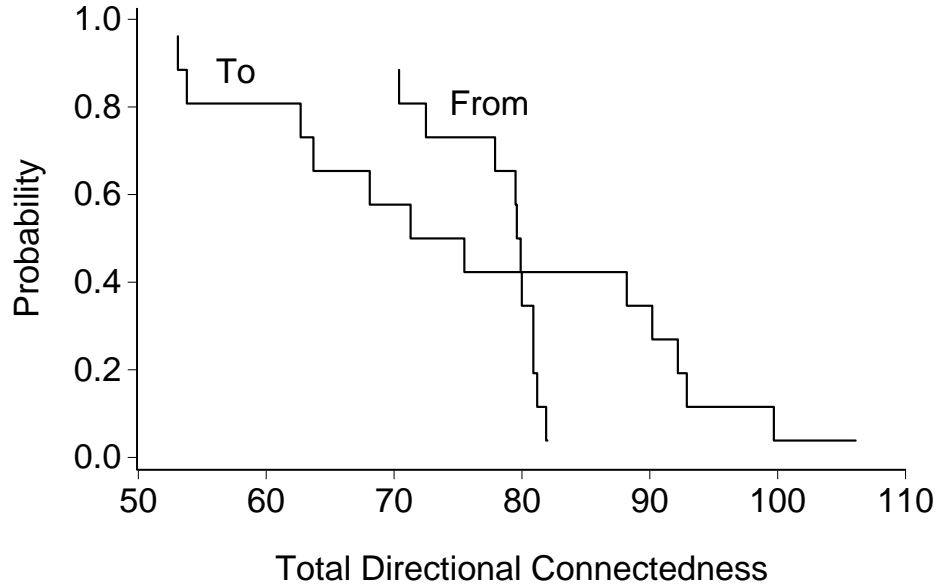
Figure 1: Full-Sample Total Directional Connectedness: Empirical Survivor Functions. We plot the empirical survivor functions for total directional connectedness "to" others and "from" others. The predictive horizon for the underlying variance decomposition is 12 days.

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, Citigroup generated a total directional connectedness measure of 106% to others. Besides the top four 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($\tilde{C}_i^H = \tilde{C}_{\bullet \leftarrow i}^H - \tilde{C}_{i \leftarrow \bullet}^H$). In terms of the net total directional connectedness Citigroup (26.5%) leads the way, followed by Bank of America (18.8%), American Express (13%), and J. P. Morgan (8.9%). AIG (-19%), PNC Group (-18%), Fannie Mae (-17%), Goldman Sachs (-15%) and Bank of New York Mellon (-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 a high degree of connectedness even during tranquil times. There is another reason for the total connectedness for a set of financial stocks to be higher than for a set of major national stock markets around the world or 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 and their stocks are more vulnerable to external and/or industry-wide shocks than others, they are likely to be transmitting these shocks to other financial stocks, 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 for a group of stocks in the finance industry is likely to be higher. It is also likely to be higher compared to the connectedness for a group of global markets, as these markets are not subject to common shocks as frequently as the stocks from the finance industry.[24]

4.3 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]

In contrast to our theoretical discussion in section 2, as well as our static empirical analysis in section 4.2, in which we progressed from "micro to macro" – that is, from pairwise connectedness, to total directional connectedness, to total connectedness – here it proves useful to proceed in reverse order, from macro to micro. We start our dynamic analysis with total connectedness, and then we move to various levels of disaggregation (total directional and pairwise directional). Finally, we also provide a brief assessment of the robustness of our results to choices of tuning parameters and alternative identification methods.

---

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

[25]In 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.

4.3.1 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 cycles; 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 cycle coincides with the burst of the dot-com bubble, followed by the downward spiral in the Nasdaq and other stock exchanges and the 2001 recession. Even if the recession was over in early 2002, the MCI WorldCom scandal of mid-2002 kept the volatility of the financial stocks and their connectedness high for another year. The second cycle started at the end of February 2007. 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 close to 90% in August 2007 and stayed above 80% until mid-2009.

In between the two big cycles of the total connectedness lie three smaller, but not necessarily negligible, cycles. We will discuss each of these cycles along with the events that possibly led to them. Before doing so, let us point to another fact that emerges from the total connectedness plot. From 1999 to 2007, whenever the total connectedness increased to a higher level, it always came back down to the 55-65% range as the sample-windows are rolled to leave that episode behind. Following the 2007-08 financial crisis, the total connectedness index stayed well above this range as of the end of April 2010, even though the financial crisis had ended almost a year before.

Earlier on in our sample, developments in the tech-heavy Nasdaq stock exchange influenced the behavior of the total volatility connectedness among the financial stocks. Starting in March 2000, the so-called dot-com bubble finally started to burst. The bursting of the

20

dot-com bubble had a serious impact on the total volatility connectedness of financial stocks. In March 2000, the volatility connectedness index increased by 7 percentage points. Despite short spells of recovery, troubles of the internet stocks continued for some time and solid signs of an imminent recession appeared on the horizon. The volatility in the bank stocks increased rapidly over this period, and so did the total volatility connectedness. From a low of 60% in early September, the connectedness index increased to 75% by mid-January 2001 and further to surpass 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 toward the second and third quarters of 2001. Total connectedness declined to 71% by early September 2001. However, 9/11 terrorist attacks worsened 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.

After the Enron scandal of late 2001, which did not have much impact on financial stocks, another corporate scandal rocked the U.S. financial markets toward the end of June 2002. This time around it was the bankruptcy of the MCI WorldCom, which was once the second-largest long distance phone company in the U.S. Unlike the Enron scandal, the MCI WorldCom scandal had a serious impact on major bank stocks. All major U.S. banks had credit positions with MCI WorldCom and hence they all suffered losses when the company declared bankruptcy.

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, being an isolated source of loss for the banks, the scandal's impact on the financial system as a whole could be contained. As of the end of 2002 total connectedness subsided very quickly to pre-July 2002 levels. After a brief increase following the invasion of Iraq in March 2003, the total connectedness declined to 58% in August 2003.

From August 2003 to February 2007, the total connectedness index went through three smaller cycles, during which it moved within the 55–80% range. The first cycle lasted from August 2003 to March 2005; the second from April 2005 to February 2006; and the third from March 2006 to February 2007. The three cycles mostly coincide with the tightening of monetary policy and its impact on the behavior of long-term interest rates.[26]

---

[26]The link between the volatility connectedness and the long-term rates is directly a result of the choices

4.3.2   Total Directional

The dynamic analysis of total connectedness gave us a clear understanding of the factors influencing the volatility connectedness across major U.S. financial stocks over the 1999-2010 period. Keeping this analysis in the back of our minds, we can now focus on the dynamics of directional connectedness over time.

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

One of the first things one notices in Figure 3 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 shock is 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. Irrespective of the size of the volatility shock, if it is the stock of a larger institution or a highly central institution (which has strong balance-sheet and off-balance-sheet connections with other banks) that received the volatility shock, then one 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 and centrality 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 are insulated from the volatility shocks to other institutions' stocks. 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"

---

of the investors. Rising long-term interest rates reflect optimism about the future economic performance. As they expect the growth to pick up, investors sell more defensive stocks such as the financial stocks and instead invest in manufacturing, energy and airlines sector stocks that are likely to benefit most from an economic recovery.

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.

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 connectedness to others presented in the lower panel of Figure 3. As the connectedness "from" others measure is smoother over the rolling-sample windows, the variation in the plots for "net" connectedness to others over the rolling-sample windows resembles the variation in the plots for connectedness "to" others.

In Figure 3 we observe that even though for each stock the "from" connectedness reached the highest levels during the 2007-08 crisis, we do not observe such a level shift in the "to" and "net" connectedness measures over the same 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 became more frequent and each time hit more stocks than before the crisis and hence were transmitted to others in larger amounts than before.

In order to better evaluate the differences between the "to" and "from" directional connectedness, in Figure 4 we plot 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 measures 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.

Temporal changes in the dispersion and skew of the "to" and "from" connectedness in Figure 4 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 "recipients of small" and "transmitters of big" – those are the distressed firms potentially poised to wreak havoc on the system.
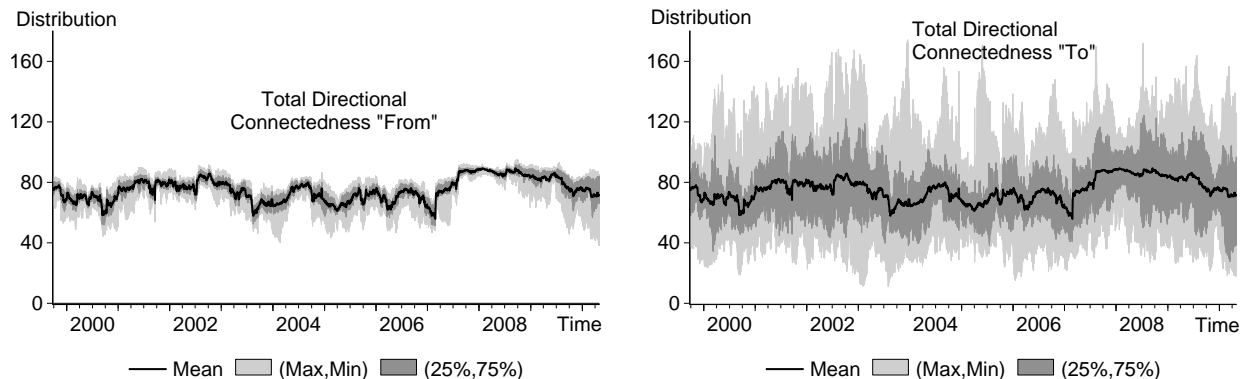
Figure 4: Rolling Distribution of Total Directional Connectedness. We plot the time series of daily min, 25%, mean, 75%, and max of the distributions of "to" and "from" total directional connectedness. The rolling estimation window width is 100 days, and the predictive horizon for the underlying variance decomposition is 12 days.

4.3.3 Pairwise Directional

In the analysis of the full-sample volatility connectedness in Section 4.2, 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 is even more crucial in the rolling sample windows case, because it helps us identify how the connectedness measures 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.

4.3.4 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 sample windows of 75 and 125 days), 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.


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.

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 the Cholesky-based

identification is the lower bound of the one obtained from the generalized 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 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.

4.4   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 the end of 2008. The analysis of this section shows how the measurement and daily monitoring of connectedness can help us understand the developments at each stage of the global financial crisis.

4.4.1   Total Connectedness at Various Stages of the Crisis

As of the end of 2006 there were already some, albeit weak, signs of slowdown in the U.S. real estate market.[27] 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, the mortgage-backed securities (MBS) markets as well as the stock market, and on the last day of February 2007, the total connectedness measure jumped by more than 17 points, the biggest increase on a single day. The increase in the total connectedness was not

---

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

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 MBS markets.

The churning in the MBS markets continued from February until early June. New Century declared bankruptcy in April. In June and July the markets became aware that big financial institutions were not insulated from the debacle in the MBS. Bear Stearns had to liquidate two of its hedge funds in July, leading to billions of dollars of losses. 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, which eventually led to the liquidity crisis of August 2007. From July 25 to August 10, the index climbed 12 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 would be badly bruised by the collapse of the ABCP market.

After seven months of learning about the problems in MBS 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. Even though it had already reached its historical maximum, in late 2007 the volatility connectedness index continued its upward move by several points.

As the MBS markets continued their descent in early 2008 Bear Stearns' financial position became untenable. Amid widespread rumors 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 March 17, 2008 with financial assistance from the Fed. As a result of the timely rescue operation, in the final days of Bear Stearns the total connectedness of the surviving thirteen banks showed an upward movement of only a couple of percentage points.

In the summer of 2008 the tension in the stock market had started to build up again as a result of Wachovia Bank's troubles. Thanks to Wachovia's high volatility, the total volatility connectedness index increased, reaching to 88.5% in mid-July (see Figure 2).

In the meantime, regional banks smaller than Wachovia failed. These were followed by news about the constantly deteriorating asset positions of Fannie Mae and Freddie Mac. Before going bankrupt, these two "government-sponsored enterprises" were taken into government conservatorship in the first week of September.

Then came the most significant event in the unfolding of the crisis. Following the news

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.

that Lehman would announce huge losses in its latest financial statement, market participants started selling Lehman Brothers' stock. Despite the overwhelming efforts over the weekend of September 13-14, 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 Lehman Brothers with taxpayers' money. As soon as Lehman Brothers declared bankruptcy on the morning of September 15, 2008 all hell broke out in financial markets around the world. That same day, the weakest of the three remaining investment banks, Merrill Lynch, announced it was being acquired by Bank of America. The total volatility connectedness index increased further to reach its maximum level of 89.2% (see Figure 2).

After months of gyrations in the U.S. financial system, the volatility connectedness started to subside toward the end of the first quarter of 2009. In March and April 2009, the total connectedness measure fluctuated between 80-85% for a while. It started to fall only after the announcement of the stress test results in May 2009. By October 2009 the index was down to 70%. 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 from declining any further. As of the end of our sample, the index was fluctuating between 70 and 75%, a range that is above the levels the index attained during tranquil times (See Figure 2).

4.4.2 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 institutions. In Table 4 we list the information on the four major banks that ceased to exist, with information on their stock tickers, market capitalization before and after the crisis, and critical dates during the crisis.


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."

Net total directional connectedness plots for AIG, Wachovia, Merrill Lynch, Lehman Brothers, Morgan Stanley and Bear Stearns are presented in Figure 6.[28] Let us spell out the most important observation in Figure 6 upfront: Even though it was the troubles of the investment banks that were followed the most throughout the crisis, Wachovia Bank is

---

[28]It is worth noting that connectedness measurements generally will not, and should not, be robust to the 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 the "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.

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 on several occasions as the global financial crisis unfolded steadily in 2007 and 2008. To start with the most vulnerable of the top five investment banks, the net volatility connectedness of Bear Stearns' stock was not sizable in the run-up to its takeover by J.P. Morgan on March 17, 2008, but it increased substantially to 6.4% on March 14 and 4.9% on March 17 (see Figure 6).

Viewed as the most vulnerable investment bank after Bear Stearns, Lehman Brothers' net directional connectedness during the liquidity crisis of August 2007 was close to 5%. It also generated close to 4% 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% for almost three months after the demise of Bear Stearns. From early June till early August 2008 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.

On Friday, September 12, 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 1% (see Figure 6). Markets were still expecting another government-orchestrated rescue operation. Only after the announcement of its bankruptcy on the morning of September 15 did Lehman Brothers' stock move to center stage in the crisis and generated substantial volatility connectedness, with a net total directional connectedness of 6% (see Figure 6). Its net pairwise connectedness with five financial stocks was in the top percentile (another five were in the top five percentiles and two in the top ten percentiles) of all the net pairwise volatility connectedness that 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)).

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, yet the markets learned that they were actually worse than previously known when the bank announced that it incurred a loss of $8.9 billion in the second quarter of 2008. In the month of June, long before the climax month (September 15 – October 15) of the financial crisis, Wachovia's stock came under heavy pressure and its

(a) Sep. 12, 2008

(b) Sep. 15, 2008
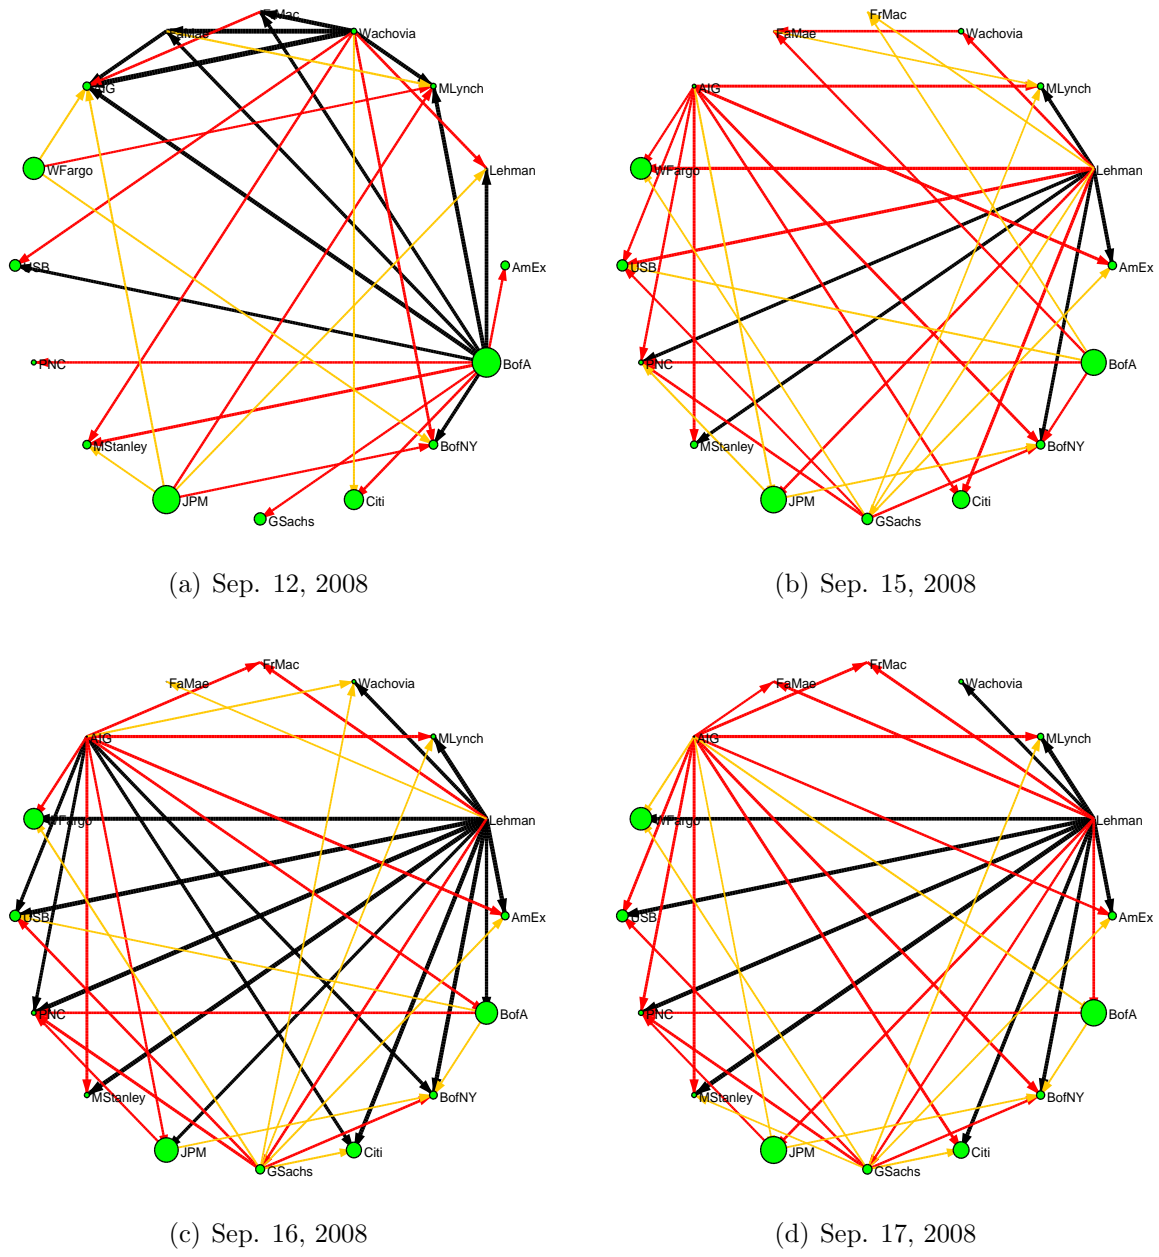
(c) Sep. 16, 2008

(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.

total directional connectedness to others (see Figure 6) increased substantially. On October 3, Wachovia was sold to Wells Fargo.

5   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.

33

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 Michagan, 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. Constantinedes 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*, Elsevier, 67-138.

Bech, M.L. and E. Atalay (2011), "The Topology of the Federal Funds Market," *Physica A: Statistical Mechanics and its Applications*, 389, 5223–5246.

Billio, M., M. Getmansky, A.W. Lo, and M. Pelizzon (2010), "Measuring Systemic Risk in the Finance and Insurance Sectors," Manuscript, MIT.

34

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.

DiNicolo, 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.

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.

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.

Watts, D. J. and S. H. Strogatz (1998), "Collective Dynamics of 'Small-World' Networks," *Nature*, 393, 440–442.