

Matrix Metrics: Network-Based Systemic Risk Scoring

SANJIV RANJAN DAS

SANJIV RANJAN DAS is the William and Janice Terry professor of finance at Santa Clara University in Santa Clara, CA. srdas@scu.edu

Morpheus: Unfortunately, no one can be told what the Matrix is. You have to see it for yourself.

—The Matrix (Wachowski and Wachowski [1999])

This article proposes a new measure of aggregate systemic risk and additional system-wide and entity-specific metrics as a complement to existing measures of systemic risk. This measure provides a quantification of system-wide risk based on the level of vulnerability of each node in the system and the interconnectedness of all nodes in the network (see Alter, Craig, and Raupach [2014] for an approach that also uses these two quantities). This metric is easy to compute and has many appealing characteristics.

Systemic risk (as opposed to systematic risk) has become an important concern since the financial crisis of 2008. Measuring and managing such risk are two salient goals of this analysis. Although systemic risk is not always easy to define, there exist some universally accepted characteristics in the extant literature (discussed later), including (1) having a large impact, (2) being widespread (i.e., affecting a large number of entities or institutions), and (3) having a ripple effect that endangers the existence of the financial system. The mortgage/financial crisis of 2008 certainly had these three characteristics, but the market crash of 1987 affected only a small

set of assets (equities) and did not endanger the financial system. Definitions of systemic risk abound, however, and economists may not agree on any single one. We describe and discuss some popular measures that are related to our new measure.

There is a growing literature on systemic risk measurement in finance and although we mention some representative articles here, a range of similar articles exist. Much of this literature uses equity returns of financial institutions and the correlations of these returns to construct systemic risk measures. One such important article is by Billio et al. [2012], in which the authors applied return correlations and Granger causality regressions on returns to construct network maps and develop network measures of systemic risk. *Joint extreme tail risk*, such as the well-known conditional value at risk (CoVaR) metric of Adrian and Brunnermeier [2010], is also used as a systemic risk measure. The *systemic expected shortfall* (SES) measure of Acharya et al. [2011] examines the tail risk for a financial institution when the aggregate system is under stress. This is similar to the *distressed insurance premium* (DIP) metric of Huang, Zhou, and Zhu [2011]. Kritzman et al. [2011] developed the absorption ratio (AR) based on situations in which the co-movement of returns of assets in a principal components analysis becomes concentrated in a single factor. A modification of this approach by

Reyngold, Shnyra, and Stein [2013] denoted the *credit absorption ratio* (CAR), which extends AR to default risk data. Levy-Carciente et al. [2015] developed a methodology for stress testing banks for systemic risk using a bipartite graph of financial institutions and assets.

The systemic risk measure in this article is different from the ideas in these related articles. First, it does not depend only on equity returns because it is general and can be used with any measure of interconnectedness. For example, a network graph generated from interbank transactions (foreign exchange, credit default swap [CDS], loans, etc.) may be used as developed by Burdick et al. [2011], as could the network generated from the Granger causality analysis performed by Billio et al. [2012, 2014]. Second, the measure separates two aspects of overall risk: *compromise level* (i.e., the risk score at each node) and *connectivity* (i.e., the network graph) across nodes and explicitly uses the network configuration in scoring systemic risk. Third, an important property of this aggregate systemic risk measure is that it is additively decomposable into individual contributions to systemic risk, enabling regulators to impose a tax on financial institutions for individual institutional contributions to aggregate risk, as also suggested by Acharya et al. [2011].

In addition to these features of the systemic risk score, other useful attributes and applications of this measure are as follows. One, it may be used in combination with network *centrality* scores to manage the risk of the financial system (the *criticality* of a node in the financial system is defined as the product of its risk [compromise] level and its centrality). Two, we propose a measure of *fragility* that is related to concentration risk (i.e., resembles a Herfindahl–Hirschman Index). This enables assessment of the speed at which contagion can spread in the system. Three, we compute the risk *increments* of the aggregate systemic risk score (i.e., the extent to which each node in the system contributes to aggregate risk if its level of compromise increases by a unit amount), thereby enabling identification of critical nodes, even though they may not be compromised at the current time. Fourth, we also define a *normalized* systemic risk score that quantifies the network effect present in the system and complements the fragility score.

In addition to these stand-alone static metrics, we explore a few comparative statics in order to understand the dynamics of the network without a full-blown

dynamic analysis. First, we examine *cross risk* (i.e., the externality effect of one node's increase in risk on the risk contribution of other nodes). We explore this risk numerically and find that cross risk is low; that is, it is not easy for a badly performing node to impose large externalities on the other nodes in terms of our metric, thus making cross risk robust for practical use. Second, we examine whether breaking large banks into smaller banks helps reduce systemic risk and find that this remedy does not work. Instead, eliminating too-big-to-fail banks exacerbates systemic risk as it increases points of failure in the system.

This article proceeds as follows. In the next section we present the notation and structure of the new systemic risk score as well as related network measures. We then extend the measure to a normalized one and provide more examples. To set this metric in context, we provide discussion in a follow-up section that compares the new metric to other systemic risk measures, summarized in Appendix. The final section provides brief concluding discussion.

MODELING

Notation

Risk in a connected network arises from compromised nodes and their connections. We propose and define a parsimonious and intuitive metric for quantifying the aggregate risk in a network of related entities and explore its properties.

Assume that the network comprises n nodes and is formally defined as the graph $G(\mathbf{V}, \mathbf{E})$ where $\mathbf{V} \in \mathcal{R}^n$ is the vertex (node) set of entities or banks and $\mathbf{E} \in \mathcal{R}^{n \times n}$ is the edge (link) set comprising elements $\mathbf{E}(V_i, V_j) \equiv E_{ij} \in \{0, 1\}$, denoting which nodes are connected to each other. The graph may be assumed to be directed (i.e., $E_{ij} \neq E_{ji}$, and undirected graphs are special cases). Also, $E_{ii} = 1, \forall i$, which is necessary for computing the risk score shown later. The link E_{ij} in this network is to be interpreted as a flow/effect from node i to node j in the sense that if bank i is affected economically, it will then transmit this impact to bank j .

The network is represented by an $(n \times n)$ adjacency matrix with all elements in $\{0, 1\}$. However, one may imagine more complex networks in which the connectivity is not binary, but depends on the degree of interaction between nodes. These matrices may be normalized

such that the diagonal $E_{ii} = 1, \forall i$ and the off-diagonal elements are scaled to values $E_{ij}/\max(E_{ij}), \forall i, j, i \neq j$, thereby extending the adjacency matrix from the binary $\{0,1\}$ case to values in the range $[0,1]$. Higher values would denote a connection with greater influence.¹ The set up is simple, yet general (see Exhibit 1 for an example of the network and matrix).

For each node V_i we define the level of compromise as $C_i \geq 0$. The risk vector for all nodes is $\mathbf{C} = [C_1, C_2, \dots, C_n]^T \in \mathcal{R}^n$. Our risk score is agnostic as to how compromise is defined. For example, the inverse of the Altman [1968] Z-score would be a good measure of compromise to use. Another choice would be the expected loss measure for a financial institution as used by Acharya et al. [2011]. We may also use credit spreads or credit ratings.

In this formulation, there is no notion of the relative size of the nodes. For example, a small hedge fund could have the same credit score as a very large, low-rated investment bank. This is not, however, a severe limitation because the investment bank might have a greater influence on other banks than does the hedge fund, either through a greater number of links in the network or via stronger links using the more generalized version of the \mathbf{E} matrix, where values $E_{ij} \in [0,1]$.

Systemic Risk Score (S)

We now define a single systemic risk score for the aggregate system that accounts for the connections between institutions and the level of individual compromise at each node in the network, which is the main measure developed in this article. Note that bold font in the equations represents either a vector or matrix.

The risk score for the entire network is defined as:

$$S(\mathbf{C}, \mathbf{E}) = \sqrt{\mathbf{C}^T \mathbf{E} \mathbf{C}} \quad (1)$$

where scalar S is a function of the compromise level vector \mathbf{C} for all nodes and the connections between nodes, given by adjacency matrix \mathbf{E} .

The function $S(\mathbf{C}, \mathbf{E})$ has some useful mathematical properties. First, it is linear homogenous in \mathbf{C} , and this will be shown to be useful in the ensuing analytics in which we need to decompose the aggregate risk score into contributions from each node. Second, as long as all numbers in the \mathbf{C} vector and \mathbf{E} matrix are positive, the value of S remains positive as well. Third, the

metric is analogous to portfolio return standard deviation in which we have replaced portfolio weights with \mathbf{C} and the covariance matrix of returns with \mathbf{E} . The pre- and post-multiplication of the adjacency matrix \mathbf{E} with the credit score vector \mathbf{C} ensures that we obtain a scalar quantity as well as a linear homogeneity in \mathbf{C} . The difference between \mathbf{E} and the covariance matrix in the standard portfolio problem is that \mathbf{E} is not symmetric, although it is positive definite.

This measure is heuristic, but the economic motivation for the metric comes from two important economic underpinnings of systemic risk. This risk is characterized by (1) the interconnectedness of financial institutions (nodes), and (2) the credit quality of these nodes, as suggested by Billio et al. [2012, pp. 4-5]:


From a theoretical perspective, it is now well established that the likelihood of major financial dislocation is related to the degree of correlation among the holdings of financial institutions, how sensitive they are to changes in market prices and economic conditions (and the directionality, if any, of those sensitivities, i.e., causality), how concentrated the risks are among those financial institutions, and how closely linked they are with each other and the rest of the economy.

As an example, suppose we have 18 nodes in a network, depicted by the adjacency matrix (i.e., the directed, unweighted graph shown in Exhibit 1). The compromise vector is $\mathbf{C} = [0, 0, 1, 2, 2, 2, 2, 2, 1, 0, 2, 2, 2, 2, 1, 0, 1, 1]^T$, where 0 is no compromise, 1 is a low level of compromise, and 2 indicates a highly compromised node. We may think of these values as credit rating scores, in which the higher the score, the worse the credit quality of the financial institution.² The risk score using Equation (1) is $S = 11.62$. This can be interpreted as the systemic risk score of the financial system. The value does not connote any meaning, per se; changes in the systemic score S , however, may be tracked by a regulator over time. When a sudden spike in S occurs, regulators investigate which financial institution is most responsible for incrementing systemic risk, as computed using a risk decomposition metric which we develop in a later section of this article. Before turning to this metric, we present in the next section an older, very useful metric for the importance of a node, denoted as centrality.

EXHIBIT 1

Directed Network of 18 Nodes

Panel A: Network Schematic



Panel B: Adjacency Matrix

	[,1]	[,2]	[,3]	[,4]	[,5]	[,6]	[,7]	[,8]	[,9]	[,10]	[,11]	[,12]	[,13]	[,14]	[,15]	[,16]	[,17]	[,18]
[1,]	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
[2,]	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[3,]	1	1	1	1	0	0	0	1	0	0	0	0	0	0	1	1	1	1
[4,]	1	1	1	1	0	0	0	1	0	0	0	0	0	0	0	1	0	0
[5,]	1	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	0	0
[6,]	1	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0
[7,]	1	1	0	0	1	0	1	0	0	0	0	0	0	0	0	1	1	0
[8,]	1	1	1	1	1	0	0	1	0	0	0	0	0	0	0	1	1	0
[9,]	1	0	0	0	0	0	0	0	1	1	1	1	1	1	0	1	0	0
[10,]	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	1	0	0
[11,]	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	1	0	0
[12,]	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	1	0	0
[13,]	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	1	0	0
[14,]	1	1	0	0	0	0	0	0	1	0	0	0	0	1	0	1	0	0
[15,]	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0
[16,]	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0
[17,]	1	1	1	0	0	0	1	1	0	0	0	0	0	0	0	1	1	0
[18,]	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1

Notes: One-way arrows mean that risk flows in the direction of the arrow. Two-way arrows mean risk flows in both directions. The network is summarized in the adjacency matrix. Note that the diagonal values are all 1. The diameter of this network—the maximal shortest distance between any two nodes—is 2.

Centrality (x) and Criticality (y)

Although the main metric developed in this article is S , understanding systemic risk and drilling down into its constituents leads us to many other measures that relate to attributes of systemic risk. In this and subsequent sections, we explore these other measures as well.

We often wish to know which node in the network is most central, (i.e., which network has the most important location in terms of connection to other nodes in the network). In a complicated network, this is not a particularly easy question to answer. Here, we employ the well-known notion of centrality, developed by sociologists a few decades ago. We note that the influence of any node in a network (denoted x_i) stems from the other nodes to which j is connected (denoted by the edges in the adjacency matrix, E_{ij}); these nodes in turn are affected by the nodes to which they are connected, and so on. This circularity may be represented in the following system of n equations:

$$x_i = \sum_{j=1}^n E_{ij} x_j, \quad \forall i = 1, \dots, n$$

The left-hand side of this system of equations is a n -vector \mathbf{x} that provides a score for the influence or centrality of each node in the network. This leads to the following definition of centrality: Eigenvalue centrality is the normalized principal eigenvector $\mathbf{x} \in \mathcal{R}^n$ that, for scalar λ , satisfies the eigensystem:

$$\lambda \mathbf{x} = \mathbf{E} \mathbf{x} \quad (2)$$


Centrality was first defined by Bonacich [1987] and popularized more recently as Google's PageRank algorithm (Brin and Page [1998]).

We computed centrality for this network and plot it in Exhibit 2.

Centrality is normalized when the highest centrality node is set to value 1 and the other node values have relative centrality to this node. Neither centrality nor fragility (defined later in the paper) depend on the compromise vector \mathbf{C} since it is computed only using adjacency matrix \mathbf{E} . To expand the concept of centrality to also account for the compromise levels at each node, we define a new metric called *criticality*, which is compromise-weighted centrality. This new measure is defined as $\mathbf{y} = \mathbf{C} \times \mathbf{x}$ where \mathbf{y} , \mathbf{C} , and $\mathbf{x} \in \mathcal{R}^n$. Note that this is an element-wise multiplication of vectors \mathbf{C} and \mathbf{x} .

EXHIBIT 2

Centrality



Note: Normalized centrality for each node in the network shown in Exhibit 1, rank ordered for display.

Critical nodes need immediate attention either because they are heavily compromised, of high centrality, or both. Centrality offers a way for regulators to prioritize attention to critical financial institutions and preempt systemic risk from blowing up. We compute criticality for this network and plot it in Exhibit 3.

The node numbers in Exhibits 2 and 3 are the same nodes in our continuing example. In an application of the model to a financial system the nodes would be individual financial institutions. Note that the centrality scores in Exhibit 2 are ordered differently than the criticality scores in Exhibit 3. This is because centrality ordering does not depend on the credit quality of the banks (\mathbf{C}). Hence, node 1, which is the most connected, is the node with the highest centrality, and node 5 has low centrality (for a visual sense, see Exhibit 1). However, criticality depends on both centrality and credit quality; thus, node 1 has very low criticality as this bank has a high credit quality, whereas node 5 now has high criticality. Some nodes, such as 11 and 12, have moderate centrality and credit quality and hence remain high in terms of metrics of centrality and criticality, as shown in Exhibits 2 and 3.

Risk Decomposition (D)

We exploit the linear homogeneity of the function $S(\mathbf{C}, \mathbf{E})$ in \mathbf{C} using Euler's equation, which decomposes first-order homogenous functions, resulting in a representation of the aggregate systemic risk score into node-wise components.

Risk decomposition is the attribution of the aggregate network risk score S to each node's individual risk contribution D_i , $i = 1, 2, \dots, n$, such that $S = \sum_{i=1}^n D_i$. The risk contribution of each node is $D_i = \frac{\partial S}{\partial C_i} C_i$.

This decomposition formula is the result of applying Euler's theorem³ to the function $S(\mathbf{C}, \mathbf{E})$ and decomposes the system-wide risk score S into the contribution of each node to total risk. The formula also shows that the individual risk contributions sum up to the total systemic score S :


$$S = \frac{\partial S}{\partial C_1} C_1 + \frac{\partial S}{\partial C_2} C_2 + \dots + \frac{\partial S}{\partial C_n} C_n \quad (3)$$

When a node fails and exits the network, the systemic risk score for the network and the risk contribution of each node within the network will also change. *Ceteris paribus*, removal of a node will lower systemic risk; the effect is analogous to quarantining a node. When a node fails, it may also affect the credit quality of other nodes (i.e., some C_i will worsen, even though the adjacency matrix \mathbf{E} becomes smaller as some connections are removed). The overall effect on the systemic risk score S and risk contributions D_i is therefore indeterminate. To establish an expected change in these risk scores, the model here needs to be extended from a static model to a dynamic one.

We computed the risk decomposition of the network in Exhibit 1, and this is shown in Exhibit 4 where $\sum_{i=1}^n D_i = 11.62$. Note that the numbers D_i for each node i depend on both the compromise vector \mathbf{C} and the network adjacency matrix \mathbf{E} . In this risk network, nodes 5 and 8 contribute the most to system-wide risk. We note that even though these nodes are not central in the network, they have a high level of compromise (C_i , $i = 5, 8$) and therefore are the nodes to be monitored most closely.


This risk decomposition is especially useful for pinpointing the network effect when a sudden rise in systemic risk score S occurs. By examining the changes

EXHIBIT 3 Criticality



Note: Criticality for each node in the network shown in Exhibit 1, rank ordered for display.

EXHIBIT 4 Risk Decomposition



Notes: The risk contribution D_i for each node in the network shown in Exhibit 1, rank ordered for display. The aggregate risk is $\sum_{i=1}^{20} D_i = 11.62$.

in risk contribution for each node from one period to the next, critical causal nodes are quickly identified. Further, one may determine whether the increase in a node's risk contribution arises from an increase in its compromise level or from an increase in its connectivity.

Risk Increment (I)

A regulator may be interested in assessing how a node in the financial network is likely to affect the system should that node become excessively compromised. This is determined by computing *risk increments*, or the change in the aggregate network risk score S when the compromise score C_i of an asset changes, that is, $I_i = \frac{\partial S}{\partial C_i}$.

Given $S = [\mathbf{C}^\top \mathbf{E} \mathbf{C}]^{1/2}$, the derivative with respect to \mathbf{C} is the vector

$$\mathbf{I} = \frac{\partial S}{\partial \mathbf{C}} = \frac{1}{2S} [\mathbf{E} \mathbf{C} + \mathbf{E}^\top \mathbf{C}] \in \mathcal{R}^n$$

which is easy to compute even for large n .

We computed the risk increments of the network in Exhibit 1, and the results are shown in Exhibit 5. Note that the numbers I_i for each node i depend on both the compromise vector \mathbf{C} and the network adjacency matrix \mathbf{E} .


We see that although node 1 has a very low current risk contribution (as shown in Exhibit 4), it has the potential to be very risky as it has the highest risk increment (see Exhibit 5) because it is a highly connected node.

Both risk contribution and risk increment are useful in identifying the source of system vulnerabilities and in remediation. In assessing whether a node should be allowed to fail, we may disconnect it from the network and assess how these metrics are affected.

Fragility (R)

A more concentrated network is one in which a few nodes have many connections whereas most nodes have very few. A highly concentrated network tends to have a greater risk of transmission because once a highly central node is compromised, the malaise rapidly spreads to other nodes. This propensity for risk to spread through a network is denoted as fragility.

EXHIBIT 5 Risk Increment



Note: I_i for each node in the network shown in Exhibit 1, rank ordered for display.

Let d be the degree of a node (i.e., the number of connections it has to other nodes). We then define the fragility of the network to be:


$$R = \frac{\epsilon(d^2)}{\epsilon(d)}$$

where the function $\epsilon(\cdot)$ stands for the expectation of the random variable in the function.

Keeping $\epsilon(d)$ constant, an increase in concentration results in an increase in $\epsilon(d^2)$ with a corresponding increase in fragility R . This definition is intuitive and the fragility measure is similar to a normalized Herfindahl-Hirschman Index (which is the numerator). If the network's connections are concentrated in a few nodes, we obtain a *hub-and-spoke network* (also known as a *scale-free network*) on which the spread of a shock is rapid; once a node with many connections is infected, disease in a network spreads rapidly.

Consider, for example, a network with four nodes each with degree 2 and which is not fragile (i.e., the fragility score is low, $R = 2$) and the same network of four nodes with degrees $\{4, 2, 1, 1\}$, which has the same mean

EXHIBIT 6 Fragility Comparison



degree but is much more fragile ($R = 11$). Concentration of degree induces fragility. This metric is a useful complement to the systemic risk score S . The fragility of the example network used in this article is computed to be 7.94.

One may wish to simply look at the often-used Herfindahl Index $\epsilon(d^2)$, but in this case the normalization by $\epsilon(d)$ is relevant because it ensures that a smaller, more concentrated network with fewer connections is more fragile than a network with many connections but less concentration. For example, consider the following two networks shown in Exhibit 6.

Network A is the typically fragile hub and spoke network, but network B is less fragile as it does not have this structure. If we merely computed the Herfindahl Index $\epsilon(d^2)$ for A and B , these networks would have values of 7 and 7.67, respectively, indicating that network B is more fragile because B has more overall degree. Therefore, we normalized by $\epsilon(d)$ to yield 2.14 and 2.67, respectively. After normalization, fragility is higher for network A , equal to 3, whereas for B it is equal to 2.67.

Finally, although fragility is a measure for the entire network, centrality is a measure for each node; hence, they are different but still linked because a network that has concentrated centrality in a few nodes will likely be more fragile.

EXTENDED METRICS

The previous section introduced several new network-based systemic risk measures, such as the aggregate systemic risk score, risk decomposition, risk increment,

fragility, and criticality. In this section, we further modify and extend these metrics.

Normalized Risk Score (\bar{S})

The units of systemic risk score S are determined by the units of compromise vector \mathbf{C} . If \mathbf{C} is a rating, then systemic risk S is measured in rating units. If \mathbf{C} is a Z-score (for instance), then S is a system-wide Z-score, and if \mathbf{C} is expected loss, then S is in system-wide expected loss units.

To compare the network effect across systems, we extend the score S to normalized score \bar{S} :

$$\bar{S} = \frac{\sqrt{\mathbf{C}^\top \mathbf{E} \mathbf{C}}}{\|\mathbf{C}\|} = \frac{S}{\|\mathbf{C}\|} \quad (4)$$

where $\|\mathbf{C}\| = \sqrt{\mathbf{C}^\top \mathbf{C}}$ is the norm of vector \mathbf{C} . When no network effects are present, $\mathbf{E} = \mathbf{I}$, the identity matrix, and $\bar{S} = 1$ (i.e., the normalized baseline risk level with no network [system-wide] effects is unity). We can use this normalized score to order systems by systemic risk. For the system in our example, the normalized score is $\bar{S} = 1.81$.

We note that this normalized measure may mask high levels of risk (i.e., if all firms were rated CCC). It is always better for a regulator to only look at S and not at \bar{S} . Therefore, this measure is useful in separating out the network effect, but it is not to be used for measuring overall systemic risk.

Varying Risk or Connectivity

The addition of a link in the network will increase both S and \bar{S} , ceteris paribus, and a reallocation of risk among nodes in vector \mathbf{C} will also change S and \bar{S} . Limiting or setting constraints on entries in matrix \mathbf{E} is akin to controls on counterparty risk in an interbank system, and limiting each entry in vector \mathbf{C} constrains own risk. A network regulator may choose limits in different ways to manage systemic risk. Simulating changes to \mathbf{C} and \mathbf{E} allows for generating interesting test case scenarios of systemic risk.

For example, with increasing risk at a node, if we keep the example network unchanged but reallocate the compromise vector by reducing node 3's risk by 1 and increasing the risk of node 16 by 1, we find that the risk score S goes from 11.62 to 11.87, and the normalized risk

score \bar{S} goes from 1.81 to 1.85. This is because node 16 is marginally more central than node 3 (as may be seen in Exhibit 2).

In this manner, we may examine how adding a link to the network or removing a link from the network may help in reducing system-wide risk. Alternatively, we may examine how additional risk at any node leads to more systemic risk. A system regulator can run these analyses to determine the best way to keep system-wide risk in check.

Cross/Spillover Risk (ΔD_{ij})

An increase in the risk level at any node i not only affects its own risk contribution D_i , but also that of other nodes ($D_j, j \neq i$) and throughout the network matrix. A single financial institution mismanaging its own risk might impose severe externalities in terms of potential risk on other banks in the system through network effects. In a situation in which banks are taxed for their systemic risk contributions, for example, and required to keep additional capital based on their individual risk contributions (D_i), externalities may instigate retaliatory actions that result in escalation in the systemic score S .

Hence, it is important to compute the severity of cross risk. In their model of financial surveillance, Espinosa-Vega and Sole [2010] pointed out that spillover risk is an important motivation for proposed capital surcharges for systemic risk.


We analyzed our sample network by computing each node's effect on risk contribution if any other node has a unit increase in compromise level. We denote the cross risk of node i when node j has a unit increase in compromise level C_j as $\Delta D_{ij} = \frac{\partial D_i}{\partial C_j}$, keeping the network topology \mathbf{E} constant. The results are shown in Exhibit 7. It is apparent that cross risk is insignificant compared to own risk contribution. This suggests that regulators need not be overly concerned with moral hazard on networks in which one node can impose severe externalities on other nodes. It also means that the risk metric S is not easily gamed for externalities (i.e., if institutions are taxed based on their risk contributions, then any single institution cannot affect the taxes of another in a material way). To this extent, the measure is robust.

The analysis of cross risk assumes that network adjacency matrix \mathbf{E} does not change with \mathbf{C} . It is hard to say in what way network topology will change. It may be


EXHIBIT 7

Change in Risk Contribution When Any Node Experiences a Unit Increase in Compromise Level

Panel A: Change in Risk Contribution (bar form)



Panel B: Change in Risk Contribution (heat map)



Note: The impact from each node on every other node is shown.

that the worsening credit quality of a given bank reduces the number of connections as some other banks stop trading with it, thereby reducing systemic risk. On the other hand, that bank may reallocate trading in a manner in which trading becomes more concentrated in a few nodes, thus raising fragility and possibly systemic risk.

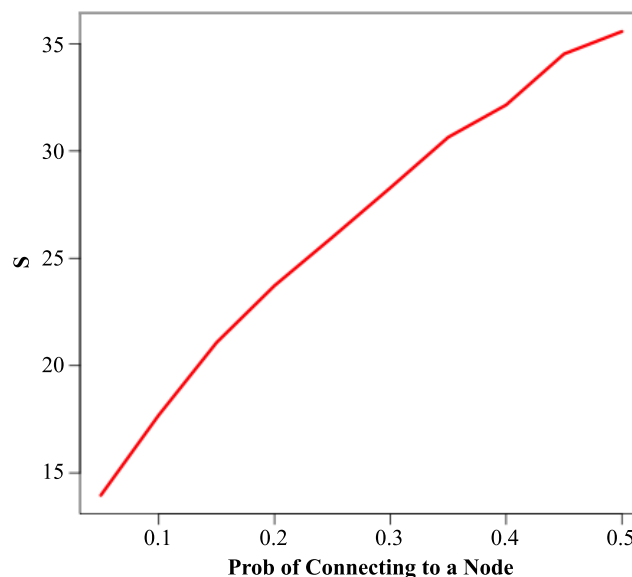
Risk Scaling and Real World Application

We assess three questions here to derive a deeper understanding of the properties of the systemic risk score S . These questions pertain to how the score changes when we scale the level of compromise, the level of interconnectedness, and the breaking down of connected nodes into less connected ones. We also provide a summary of the application of the model to real world data.

First, *ceteris paribus*, how does an across-the-board change in compromise vector C affect S ? The answer is simple: Since S is linear homogenous in C , this effect is purely linear.

Second, how does an increase in connectivity impact systemic risk S ? Is this a linear or non-linear effect? We ran a simulation of a 50-node network and examined S as the number of connections per node was increased. Simulation results are shown in Exhibit 8. The plot shows how the risk score increases as the probability of two nodes being bilaterally connected increases from 5% to 50%. For each level of bilateral probability, a random directed network of 50 nodes was generated.⁴ We then set the diagonal to 1 as required. The rest of the off-diagonal elements are 1 or 0 and were generated by the random graph function. This is the simulated E matrix. A compromise vector C was also generated with equally likely values $\{0,1,2\}$. Using C and E , we computed the systemic risk score S . This was repeated 100 times, and the mean risk score across 100 simulations is plotted on the y-axis against the bilateral probability on the x-axis. These results, based on random graph generation, show that the risk score increases with connectivity as expected, but in a less than linear fashion (i.e., the plot is mildly concave). This relates to the results obtained by Vivier-Lirimont [2006]; Blume et al. [2011]; and Gai, Haldane, and Kapadia [2011], who showed that dense interconnections destabilize networks as risk increases with increasing density. In our case, the systemic risk score also increases, but less than linearly, with the number of connections.⁵

EXHIBIT 8 The Increase in Risk Score S as the Number of Connections Per Node Increases




Notes: The graph shows how the risk score increases as the probability of two nodes being bilaterally connected increases from 5% to 50%. For each level of bilateral probability, a random network was generated for 50 nodes. A compromise vector was also generated with equally likely values $\{0,1,2\}$. This was repeated 100 times, and the mean risk score across 100 simulations is plotted on the y-axis against the bilateral probability on the x-axis.

Third, we examine whether partitioning nodes into more numerous, smaller entities reduces systemic risk (a question also addressed in very different models by Cabrales, Gottardi, and Vega-Redondo [2014] and Vivier-Lirimont [2006]). The idea here is to assess whether fracturing too-big-to-fail banks into smaller entities will result in a reduction in systemic risk. Whereas the first two questions did not consider an increasing or decreasing the number of nodes, in this case we explicitly increase the numbers of nodes and reduce the average number of connections per node so as to keep the overall connectivity unchanged while changing the structure of the network. The program logic is very much the same as in the previous simulation except that the C and E matrix are constructed differently and the x-axis in the exhibits is the number of nodes in the network. Exhibit 9 shows that the risk score S in fact increases. Thus, splitting large banks into smaller banks does not reduce systemic risk. Risk increases because the number of points of

EXHIBIT 9

The Change in Risk Score S as the Number of Nodes Increases While Keeping the Average Number of Connections Between Nodes Constant



Notes: This exhibit mimics the case in which banks are divided into smaller banks, each of which then contains part of the transacting volume of the previous bank. The graph shows how the risk score increases as the number of nodes increases from 10 to 100 while expected number of total edges in the network remains the same. A compromise vector was also generated with equally likely values $\{0,1,2\}$. This was repeated 5,000 times for each fixed number of nodes, and the mean risk score across 5,000 simulations is plotted on the y -axis against the number of nodes on the x -axis.

cascading failure increase when a large bank is split into smaller ones.

I would caution the reader to take this result on splitting banks with a grain of salt. The fact that splitting large banks into smaller ones increases S , thus keeping the number of connections overall constant, is a reduced-form result to evaluating a possible policy prescription to remedy the too-big-to-fail problem. It may be too simplistic an approach for the analysis of what is clearly a highly complex and controversial issue. This approach lacks a notion of size and implies that splitting small banks into even smaller ones will raise systemic risk, which may not be the case. Does this mean that it is best to collapse all banks into one large global bank? Obviously not, as other factors come into play, and the ceteris paribus nature of this analysis in which we assume that credit quality remains the same is not valid in extreme cases. The outcome of the simple analysis conducted

here seems related to the network result in which adding an extra road to a network to relieve traffic congestion has the effect of increasing congestion for all. Therefore, despite its simplicity and myriad assumptions, this result does offer one starting point for the analysis of policies around too-big-to-fail banks.

Fourth, the program code for systemic risk networks was applied to real world data in India to produce daily maps of the Indian banking network, as well as the corresponding risk scores.⁶ The credit risk vector \mathbf{C} was based on credit ratings for Indian financial institutions (FIs). The network adjacency matrix was constructed using ideas in the article by Billio et al. [2012], who created a network using Granger causality. This directed network comprises an adjacency matrix of values $(0,1)$ where node i connects to node j if the returns of bank i Granger cause those of bank j (i.e., edge $E_{ij} = 1$). This was applied to U.S. financial institution stock return data as well as in a follow-up article using CDS spread data from the United States, Europe, and Japan (see Billio et al. [2014]) in which the global financial system was also found to be highly interconnected. In applying the methodology of this article to an Indian context, the network matrix was created by applying this Granger causality method to Indian FI stock returns.

The system used here is available in real time and may be accessed directly through a browser. To begin, different selections of a subset of FIs may be made for analysis. See Exhibit 10 for screenshots of this step. Once these selections are made and the “Submit” button is clicked, the system generates the network and the various risk metrics, shown in Exhibits 11 and 12, respectively.

Discussion of Other Measures of Systemic Risk

As a practical matter, several measures of systemic risk have been proposed and each implicitly defines systemic risk as that risk being quantitatively determined by their measure. This is definition by quantification, or measurement as one sees it. In our setting of risk networks, the system-wide risk scores $\{S, \bar{S}\}$ capture systemic risk as a function of the compromise vector \mathbf{C} and the network of connected risk entities \mathbf{E} . Other research conducts this step differently. Some measures of systemic risk are network-based, but most of the measures are based on stock return correlations.

EXHIBIT 10

Screens for Selecting the Relevant Set of Indian FIs to Construct the Banking Network

Panel A: First Selection

Systemic Risk Dashboard

Segment: Banks

Firms: [Empty]

Parameter: Network Plot a

Date: [Empty]

[Submit]

- [Select all]
- Asset Management Cos.
- Banks
- Finance
- Financial Institutions
- Financial Services
- Insurance
- Investment Companies
- Misc

Panel B: Second Selection

Systemic Risk Dashboard

Segment: Banks

Firms: All selected

Parameter: Network Plot a

Date: [Empty]

[Submit]

- [Select all]
- STATE BANK OF INDIA
- IDBI BANK LTD.
- HDFC BANK LTD.
- KOTAK MAHINDRA BANK LT
- ORIENTAL BANK OF COMME
- FEDERAL BANK LTD.
- STATE BANK OF BIKANER &
- ING VYSYA BANK LTD.
- DENA BANK

Panel C: Selection Date

Systemic Risk Dashboard

Segment: Banks

Firms: All selected

Parameter: Network Plot, C

Date: 11/20/2015

[Submit]

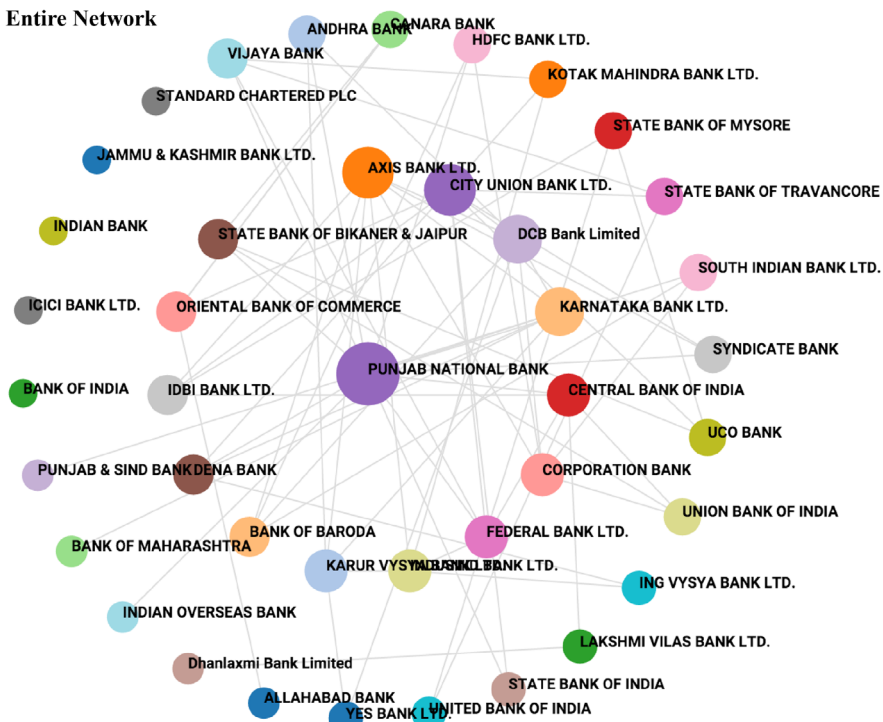
Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30					

Notes: Panel A shows selecting only banks. Panel B chooses within banks, and we select all of them. Panel C shows the date of the selection.

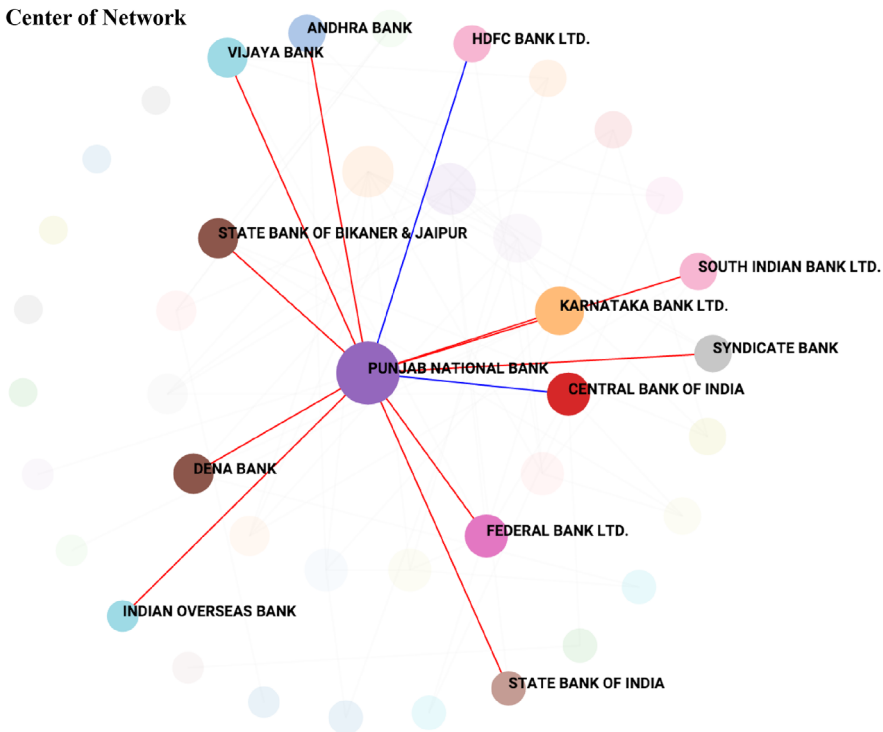
EXHIBIT 11

Screens for the Indian FIs Banking Network

Panel A: Entire Network



Panel B: Center of Network




Notes: Panel A shows the entire network. Panel B shows the network when we mouse over the bank in the middle of the plot. Lighter lines show that the bank is affected by the other banks, and darker lines indicate that the bank affects others in a Granger causal manner.

EXHIBIT 12

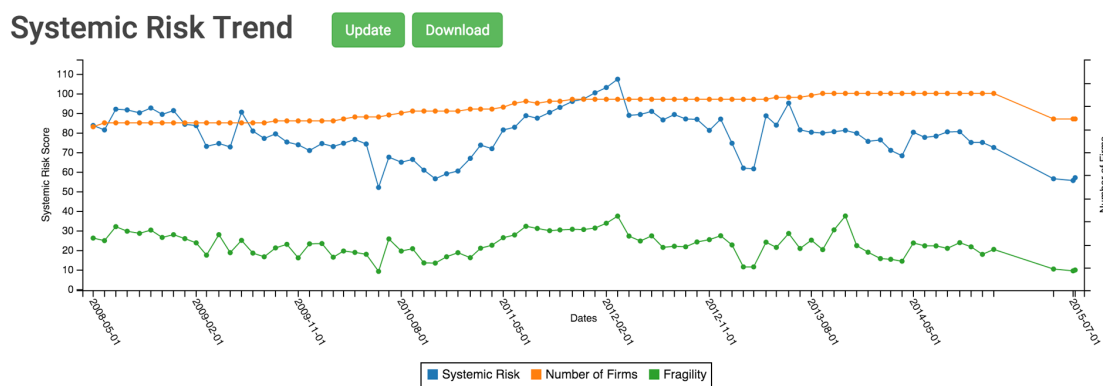
Screens for Systemic Risk Metrics of the Indian FI's Banking Network

Panel A: Current Risk Metrics

Fragility **2.91** Systemic Risk Score **15.75**



Panel B: Risk Trend History (2008–2015)



Notes: In Panel B, we see that the number of firms (flattish line) does not change much over time. Of the other two lines, the upper is the systemic score S and the lower is fragility. These lines track each other closely.

The Appendix provides a brief summary of some of the popular systemic-risk measures proposed in the literature, including the Granger causality-based network of Billio et al. [2012], the CoVaR metric of Adrian and Brunnermeier [2010], the AR of Kritzman et al. [2011], the CAR measure of Reyngold, Shnyra, and Stein [2013], and the SES measure of Acharya et al. [2011].

An important difference exists between the between the Granger causality-based network, AR, and CoVaR versus the SES measure. The three former measures assess the impact a single bank has on the system

whereas the latter measure assesses the reverse (i.e., the impact of system-wide risk on each bank). The new measures of system-wide risk $\{S, \bar{S}\}$ proposed in this article are akin to the first approach, and I believe that this is the more relevant view of systemic risk, which also offers an aggregate risk score. Both approaches, however, are relevant in computing extra systemic risk capital requirements.

There are some salient differences between these measures of systemic risk and the network score in this article. First, these measures focus on the effect a

given institution's failure has on others and thus they are pairwise and conditional. In contrast, network risk scores are system-wide and unconditional. Second, these measures are based on correlations, and correlations tend to be high in crisis periods but are not empirically established as early warning indicators of systemic risk. Relying on stock return correlations as an early warning indicator of network risk is likely to be futile, as correlation matrixes reflect systemic risk after, rather than before, the risk has arisen. Network-based measures may be better at identifying a systemic vulnerability prior to a system shock, but these measures need empirical validation as well. Third, correlation-based measures tend to be removed from the underlying mechanics of the system and are in the nature of implicit statistical metrics. Network-based measures directly model the underlying mechanics of the system because the adjacency matrix \mathbf{E} can be developed based on physical transaction activity. Furthermore, the compromise vector is a function of firm quality that may be measured in multidimensional ways. This separation of network effect (connectivity) and individual bank risk (compromise) and their combination into a single aggregate risk score offers a simple, practical, and general approach to measuring systemic risk.

This article is not only related to the growing literature on measures of systemic risk, but also to the network literature in economics in articles like those of Acemoglu, Ozdaglar, and Tahbaz-Salehi [2013]; Allen and Gale [2000]; Allen, Babus, and Carletti [2012], and the literature on risk in clearing systems (see Eisenberg and Noe [2001]; Duffie and Zhu [2011], and Borovkova and El Mouttali [2013]). Systemic risk measures based on dynamic conditional correlations are also proposed (see Brownlees and Engle [2010] and Engle, Jondeau, and Rockinger [2012]). Therefore, the novel framework in this article may be used as a complement to existing approaches. Whether the network is derived from a physical deal flow or from returns data, the risk score S may be computed, decomposed by node, and risk increments derived therefrom, along with many other metrics, to provide a useful dashboard for managing systemic risk.

CONCLUSION

This framework for network-based systemic risk modeling developed system-wide risk scores such as a

new aggregate systemic risk score (S), a normalized score (\bar{S}), a fragility score (R), and entity-specific risk scores—a risk decomposition (D_i), risk increments (I_i), criticality (γ_i), and a score for spillover risk (ΔD_{ij}). All of these metrics use simple data inputs: an institution specific compromise vector \mathbf{C} and the adjacency matrix of the network graph of financial institution linkages \mathbf{E} . The risk metrics are general (i.e., independent of the particular definitions of \mathbf{C} , \mathbf{E}) and complement and extend systemic risk measures in the extant literature.

Modeling extensions are also envisaged. In the current version of the model, the compromise vector \mathbf{C} is independent of the connectivity matrix \mathbf{E} . Making \mathbf{C} a function of \mathbf{E} (and vice versa) leads to interesting additional implications, and of course, fresh econometric questions. For example, \mathbf{C} may be an increasing function of \mathbf{E} , but then again \mathbf{E} may be a decreasing function of \mathbf{C} , making it unclear whether an increase in risk or transaction volume always leads to a higher level of potential systemic risk. Issues such as the structure of the network and the interaction of its components have been addressed in the models of Allen, Babus, and Carletti [2012], Glasserman and Young [2013], and Elliott, Golub, and Jackson [2014]. The welfare implications of over-linking have been discussed in the contagion model of Blume et al. [2011].

The question of how to construct composite connectivity matrixes across markets is also an interesting issue.⁷ One may obtain a network matrix from transactions in the CDS market (for example, see Getmansky, Girardi, and Lewis [2014]) and another from the bond markets, but the question of combining these two matrixes (call them \mathbf{E}_1 and \mathbf{E}_2) into one composite \mathbf{E} matrix requires a weighting scheme or other collapsing technical condition. One solution to this would be to construct \mathbf{E} from bilateral credit valuation adjustment (CVA) numbers because this directly measures the exposure of each financial institution to another across all products and asset classes. Using counterparty exposures as a device is also considered in the *10-by-10-by-10* systemic risk measurement approach recommended by Duffie [2011].

From a regulatory point of view, there are many applications for this framework. First, the imposition of additional capital requirements may be based on a composite score computed from risk decomposition numbers, taking into account additional informative metrics such as criticality, risk increments, and spillover risk.

(See a proposal for this by Espinosa-Vega and Sole [2010].) Second, this composite score may be used to allocate supervision money across various financial institutions. Third, the systemic score can be tracked over time, and empirical work will be needed to backtest whether the systemic score S is a useful early warning predictor of systemic risk events. Using a different approach, Kritzman et al. [2011] and Reyngold, Shnyra, and Stein [2013] found predictability of systemic risk. Fourth, an analysis of network robustness in addition to measuring systemic risk is a complementary analysis (see, for example, those conducted by Allen and Gale [2000] and Callaway et al. [2000]).

APPENDIX

OTHER SYSTEMIC RISK MEASURES

This appendix provides a brief summary of some of the popular systemic risk measures proposed in the literature and discussed in this article.

Systemic Risk

Billio et al. [2012] defined two measures of systemic risk across banks, hedge funds, broker/dealers, and insurance companies. Their idea was to measure correlations among institutions directly and unconditionally using principal components analysis (PCA) and Granger causality regressions and thereby to assess the degree of connectedness in the financial system.

In their framework, the total risk of the system is the variance of the sum of all financial institution returns, denoted σ_s^2 . PCA comprises an eigenvalue decomposition of the covariance matrix of financial institutions' returns, and systemic risk is higher when the number of principal components n that explain more than a threshold H of the variation in the system is small. Using the notation employed in their article:

$$h_n = \frac{\omega_n}{\Omega} > H \tag{A-1}$$

where h_n is the fraction of σ_s^2 that is explained by the first n components (i.e., $\Omega = \sum_{i=1}^N \lambda_i$ and $\omega_n = \sum_{i=1}^n \lambda_i$, where λ_i is the i -th eigenvalue). We note that σ_s is linear homogenous and therefore can be decomposed to obtain the risk contribution of each financial institution in the same manner as is done for our network risk measure S .

In addition to this covariance matrix-based measure of systemic risk, Billio et al. [2012] also created a network using Granger causality. This directed network is represented by an adjacency matrix of values (0,1) where node i connects to node j if the returns of bank i Granger cause (in a linear or nonlinear way) those of bank j (i.e., edge $E_{i,j} = 1$). This adjacency matrix is then used to compute connectedness measures of risk such as the number of connections, fraction of connections, centrality, and closeness. These measures correspond to some of those presented in the exposition earlier, and the first two measures report an aggregate measure of system-wide risk, different from the S measure developed in this article. Again, since system-wide risk is defined as a count of the number of connections, it is easy to determine what fraction is ascribable to any single financial firm. Billio et al. [2012] applied the metrics to U.S. financial institution stock return data and, in a follow-up article, to CDS spread data from the United States, Europe, and Japan (see Billio et al. [2014]) in which the global system is also found to be highly interconnected.

Overall, we note a strong complementarity between the analyses performed by Billio et al. [2012] and those in our article, and using the network matrix in their article, we may implement our systemic risk score S as well. Hence, this article can be extended to use the results found in this earlier work.

CoVaR

The CoVaR measure of Adrian and Brunnermeier [2010] estimates a bank or the financial sector's value at risk (VaR) given that a particular bank has breached its VaR. The authors used quantile regressions on asset returns (R) using data on market equity and book value of debt. Pairwise $\text{CoVaR}(j|i)$ for bank j given bank i is at VaR is defined implicitly as the quantile α satisfying

$$\Pr[R_j \leq -\text{CoVaR}_\alpha(j|i) | R_i = -\text{VaR}_\alpha(i)] = \alpha \tag{A-2}$$

where $\text{VaR}(i)$ is also defined implicitly as $\Pr[R_i \leq -\text{VaR}_\alpha(i)] = \alpha$. The actual measure of systemic risk is then

$$\Delta\text{CoVaR}_\alpha(j|i) = \text{CoVaR}_\alpha(j|i) - \text{VaR}_\alpha(j) \tag{A-3}$$

The intuition here is one of undercapitalization when a systemic event occurs; that is, extra capital is needed because capital needed for solvency at the time of a systemic event ($\text{CoVaR}_\alpha(j|i)$) is greater than capital needed in normal times ($\text{VaR}_\alpha(j)$). Replacing j with the system's value $\Delta\text{CoVaR}_\alpha(S|i)$ gives an aggregate measure of systemic risk. This is still, however, not an aggregate measure of risk (such as S in this article)

but rather one that assesses the systemic risk increment or contribution of the i -th financial institution.

AR

The AR metric of Kritzman et al. [2011] uses another approach to measure systemic risk. It calculates how many eigenvectors are needed to explain the variation in industry returns. We can infer that the fewer eigenvectors needed, the greater the systemic risk since the sources of risk are more unified. If the AR is low, then the sources of risk are disparate. The AR is computed as follows:

$$AR = \frac{\sum_{i=1}^n \sigma_{E_i}^2}{\sum_{j=1}^N \sigma_{A_j}^2}$$

where n is the number of eigenvectors used (in their article, Kritzman et al. [2011] used 1/5 the number of assets (N)). The variance of the eigenvectors is denoted $\sigma_{E_i}^2$ and that of the assets is $\sigma_{A_j}^2$. Reyngold, Shnyra, and Stein [2013] implemented a modified version of the AR ratio by using the covariance matrix of asset (not equity) returns only for financial firms, where asset values are derived from a structural credit model. They also only used the first eigenvector's variance since the data were restricted to a single industry. This measure is known as the CAR.

SES

The SES measure of Acharya et al. [2011] captures the amount by which an otherwise appropriately capitalized bank is undercapitalized in the event of a systemic crisis. It is related to marginal expected shortfall (MES), which is the average return of a financial institution for the 5% worst days in the market. Mostly, SES is analogous to CoVaR where value-at-risk is replaced with expected shortfall (ES), though the implementation details and variables used differ in the article of Acharya et al. We may think of SES as the equity shortfall a firm experiences when aggregate banking equity $e(S)$ is below a threshold H :

$$SES(j) = E[H(j) - e(j) | e(S) \leq H] \quad (\text{A-4})$$

where $H(j)$ is the desired threshold level of equity for bank j , with equity level $e(j)$. SES has useful properties in that it is in dollar terms and scales with institutional size and thus is easily aggregated. The DIP measure of Huang, Zhou, and Zhu [2011] is similar to the SES of Acharya et al. [2011] in that it also captures the expected losses of a financial institution conditional on losses being greater than a threshold level.

ENDNOTES

I am grateful for comments and feedback from editors Mila Getmansky and Roger Stein. Thanks to Adrian Alter, Ed Altman, Menachem Brenner, Amit Bubna, Jorge Chan-Lau, Nikhil Dighe, Marco Espinosa-Vega, Dale Gray, Levent Guntay, Raman Kapur, Nagpurnanand Prabhala, Sanjul Saxena, Shann Turnbull. Thanks also to participants at the Consortium for Systemic Risk Analytics at the Michigan Institute of Technology; the International Risk Management Conference, Poland; the International Monetary Fund; the Federal Deposit Insurance Corporation; Moody's Analytics; Quantitative Work Alliance For Applied Financial Education and Wisdom, San Francisco; Seoul National University; Hautes Etudes Commerciales, Montreal; Pan-Indian Institute of Management Meetings, Google Mountain View; R/Finance Conference Chicago; R Meetup Santa Clara; the Commodities and Futures Trading Commission Webinar; Center for Data Analytics and Risk (CDAR) at UC Berkeley; the University of Washington, Seattle; the INFORMS conference, Philadelphia. The Reserve Bank of India funded the implementation of this article on real time data in India, and the research firm InnovAccer collected the data and hosted the system. The author may be reached at srdas@scu.edu or by phone at 408-554-2776.

¹The networks in this article are not required to be *symmetric* ($E_{ij} = E_{ji}$) or *regular* ($\sum_{i \neq j} E_{ij} = \sum_{i \neq j} E_{ji} = \text{constant}$) as defined by Acemoglu, Ozdaglar, and Tahbaz-Salehi [2013].

²This is a static model, and at any point in time the credit scores for each bank are a given quantity. To the extent that banks' fortunes are correlated based on common economic factors, these credit scores are likely to reflect that correlation. In a static model, however, we do not require dynamics with an underlying correlation of credit scores. Also, the network adjacency matrix will capture some of the connections between banks' fortunes, and so correlation of credit quality may be implicit, despite no explicit modeling of correlations.

³Euler's theorem states that if a function $f(\mathbf{x})$, $\mathbf{x} \in \mathcal{R}^n$ is homogenous of degree 1, then it may be written as $\sum_{i=1}^n \frac{\partial f}{\partial x_i} x_i$.

⁴We used the R programming language and the package `igraph` for these analyses. Random networks were generated using the function `erdos.renyi.game`. Another more complex approach is to use the law of preferential attachment to generate only scale-free networks, although the results are likely to be the same as we are exploring the density of the network rather than its structure.

⁵The shape of the plot in Exhibit 7 is unsurprising in retrospect, as the metric S contains the adjacency matrix \mathbf{E} under the square root sign. As \mathbf{E} becomes denser, S will resemble a plot of the square root of increasing numbers, in mildly concave shape.

⁶Our thanks to the Reserve Bank of India for sponsorship and to InnovAccer (www.innovaccer.com) for collecting the data and hosting the site that runs the program code.

⁷Levy-Carciente et al. [2015] have an interesting model in which a network of banks intersects with a network of asset markets.

REFERENCES

- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi. "Systemic Risk and Stability in Financial Networks." Working paper, Michigan Institute of Technology, 2013.
- Acharya, V., L.H. Pedersen, T. Philippon, and M. Richardson. "Measuring Systemic Risk." Working paper, New York University, 2011.
- Adrian, T., and M. Brunnermeier. "CoVaR." Working paper, Princeton University, 2010.
- Allen, F., and D. Gale. "Financial Contagion." *Journal of Political Economy*, Vol. 108, No. 1 (2000), pp. 1-32.
- Allen, F., A. Babus, and E. Carletti. "Asset Commonality, Debt Maturity and Systemic Risk." *Journal of Financial Economics*, 104 (2012), pp. 519-534.
- Alter, A., B. Craig, and P. Raupach. "Centrality-Based Capital Allocations and Bailout Funds." Working paper, International Monetary Fund, 2014.
- Altman, E.I. "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy." *Journal of Finance*, Vol. 23, No. 4 (1968), pp. 189-209.
- Billio, M., M. Getmansky, A. Lo, and L. Pelizzon. "Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors." *Journal of Financial Economics*, Vol. 104, No. 3 (2012a), pp. 536-559.
- Billio, M., M. Getmansky, D. Gray, A. Lo, R. Merton, and L. Pelizzon. "Sovereign, Bank and Insurance Credit Spreads: Connectedness and System Networks." Working paper, International Monetary Fund, 2014.
- Blume, L., D. Easley, J. Kleinberg, R. Kleinberg, and E. Tardos. "Network Formation in the Presence Of Contagious Risk." *Proceedings of the 12th ACM Conference on Electronic Commerce*, 2011.
- Bonacich, P. "Power and Centrality: A Family of Measures." *American Journal of Sociology*, Vol. 92, No. 5 (1987), pp. 1170-1182.
- Borovkova, S., and H.L. El Mouttali. "Systemic Risk and Centralized Clearing of OTC Derivatives: A Network Approach." Working Paper, VU Amsterdam, 2013.
- Brin, S., and L. Page. "The Anatomy of a Large-Scale Hypertextual Web Search Engine." *Computer Networks and ISDN Systems*, 30 (1998), pp. 107-117.
- Brownlees, C.T., and R.F. Engle. "Volatility, Correlation, and Tails for Systemic Risk Measurement." Working paper, New York University, 2010.
- Burdick, D., S. Das, M.A. Hernandez, H. Ho, G. Koutrika, R. Krishnamurthy, L. Popa, I. Stanoi, and S. Vaithyanathan. "Extracting, Linking and Integrating Data from Public Sources: A Financial Case Study." *IEEE Data Engineering Bulletin* Vol. 34, No. 3 (2011), pp. 60-67.
- Cabrales, A., P. Gottardi, and F. Vega-Redondo. "Risk-Sharing and Contagion in Networks." Working paper, University College London, 2014.
- Callaway, D.S., M.J. Newman, S.H. Strogatz, and D.J. Watts. "Network Robustness and Fragility: Percolation in Random Graphs." *Physical Review Letters* Vol. 85, No. 25 (2000), pp. 5468-5471.
- Duffie, D. "Systemic Risk Exposures: A 10-by-10-by-10 Approach." In *Systemic Risk and Macro Modeling*, edited by M. K. Brunnermeier and A. Krishnamurthy. Chicago: University of Chicago Press, 2011.
- Duffie, D., and X. Zhu. "Does a Central Clearing Counterparty Reduce Counterparty Risk?" *Review of Asset Pricing Studies*, 1 (2011), pp. 74-95.
- Eisenberg, L. and T. Noe. "Systemic Risk in Financial Systems," *Management Science*, Vol. 47, No. 2 (2001), pp. 236-249.
- Elliott, M., B. Golub, and M. Jackson. "Financial Networks and Contagion." *American Economic Review*, Vol. 104, No. 10 (2014), pp. 3115-3153.
- Engle, R., E. Jondeau, and M. Rockinger "Dynamic Conditional Beta and Systemic Risk in Europe." Working paper, New York University, 2012.

Espinosa-Vega, M., and J. Sole. "Cross-Border Financial Surveillance: A Network Perspective." Working paper, International Monetary Fund (WP/10/105), 2010.

Gai, P., A. Haldane, and S. Kapadia. "Complexity, Concentration and Contagion," *Journal of Monetary Economics*, 58 (2011), pp. 453-470.

Getmansky, M., G. Girardi, and C. Lewis. "Interconnectedness in the CDS Market." Vanderbilt Owen Graduate School of Management Research Paper, 2014.

Glasserman, P., and H.P. Young. "How Likely Is Contagion in Financial Networks?" Working paper #0009, Columbia University, 2013.

Huang, X., H. Zhou, and H. Zhu. "Systemic Risk Contributions." Working paper, Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C., 2011.

Kritzman, M., Y. Li, S. Page, and R. Rigobon. "Principal Components as a Measure of Systemic Risk," *The Journal of Portfolio Management*, Vol. 37, No. 4 (2011), pp. 112-126.

Levy-Carciente, S., D.Y. Kenett, A. Avakian, H.E. Stanley, and S. Havlin. "Dynamical Macroprudential Stress Testing Using Network Theory." *Journal of Banking and Finance*, 59 (2015) pp. 164-181.

Reyngold, A., K. Shnyra, and R. Stein. "Aggregate and Firm-level Measures of Systemic Risk from a Structural Model of Default." *The Journal of Alternative Investments*, Vol. 17, No. 4 (2015), pp. 58-78.

Vivier-Lirimont, S. "Contagion in Interbank Debt Networks." Working paper, Reims Management School and CES, Paris I Pantheon Sorbonne University, 2006.

Wachowski, A., and L. Wachowski. *The Matrix*. DVD. Directed by A. Wachowski and L. Wachowski. Burbank, CA: Warner Home Video, 1999.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@ijournals.com or 212-224-3675.