

NBER WORKING PAPER SERIES

CLUSTERED NETWORK CONNECTEDNESS:
A NEW MEASUREMENT FRAMEWORK, WITH
APPLICATION TO GLOBAL EQUITY MARKETS

Bastien Buchwalter
Francis X. Diebold
Kamil Yilmaz

Working Paper 34796
<http://www.nber.org/papers/w34796>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

We gratefully acknowledge useful comments from the Co-Editor (Viktor Todorov), the Guest Editor, and two referees, as well as Peter Hansen, Oscar Jorda, Chang-Jin Kim, Mikkel Plagborg-Moller, and Neil Shephard. All remaining errors are ours alone. Jacob Broussard provided outstanding research assistance. Yilmaz gratefully acknowledges the support of the Scientific and Technological Research Council of Turkey (TUBITAK) under grant number 121C271. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Bastien Buchwalter, Francis X. Diebold, and Kamil Yilmaz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Clustered Network Connectedness: A New Measurement Framework, with Application to
Global Equity Markets

Bastien Buchwalter, Francis X. Diebold, and Kamil Yilmaz

NBER Working Paper No. 34796

February 2026

JEL No. G1

ABSTRACT

Network connections, both across and within markets, are central in countless economic contexts. In recent decades, a large literature has developed and applied flexible methods for measuring network connectedness and its evolution, based on variance decompositions from vector autoregressions (VARs), as in Diebold and Yilmaz (2014). Those VARs are, however, typically identified using full orthogonalization (Sims, 1980), or no orthogonalization (Koop, Pesaran, and Potter, 1996; Pesaran and Shin, 1998), which, although useful, are special and extreme cases of a more general framework that we develop in this paper. In particular, we allow network nodes to be connected in "clusters", such as asset classes, industries, regions, etc., where shocks are orthogonal across clusters (Sims style orthogonalized identification) but correlated within clusters (Koop-Pesaran-Potter-Shin style generalized identification), so that the ordering of network nodes is relevant across clusters but irrelevant within clusters. After developing the clustered connectedness framework, we apply it in a detailed empirical exploration of sixteen country equity markets spanning three global regions.

Bastien Buchwalter
SKEMA Business School
bastien.buchwalter@skema.edu

Kamil Yilmaz
Koc University
Department of Economics
kyilmaz@ku.edu.tr

Francis X. Diebold
University of Pennsylvania
Department of Economics
and NBER
fdiebold@sas.upenn.edu

1 Introduction

Network connectedness is of interest in numerous economic contexts, from financial markets, to business cycles, to international trade. In this paper we propose a flexible approach to empirical network connectedness measurement and interpretation, working in the Diebold-Yilmaz (DY) framework (Diebold and Yilmaz, 2009), which has been developed and applied extensively in recent decades (Diebold and Yilmaz, 2023).

In the DY framework, one proceeds in several steps. First, one fits a dynamic approximating model to the N -dimensional set of objects (network nodes) y whose connectedness is to be measured. Vector autoregressions (VARs) are widely used as approximating models, because of their appealing blend of approximation accuracy and simplicity. Hence we will refer to the approximating model as a “VAR” throughout this paper, but other models, including structural models, may be used.

Second, having estimated an approximating VAR, one uses it to produce impulse-response functions (IRFs) or variance decompositions (VDs). DY-style empirical connectedness measurement seeks consistency with a wide variety of unknown underlying data-generating processes, and it therefore seeks to impose only minimal restrictions when identifying IRFs/VDs.

Two approaches to IRF/VD identification, namely the “Cholesky-factor approach” and the “generalized approach”, are dominant in the empirical DY connectedness measurement literature.¹ In the Cholesky-factor approach, the IRFs/VDs are obtained from Cholesky-factor orthogonalizing transformations of the reduced-form VAR shocks (Sims, 1980). Because shocks are orthogonal after Cholesky transformation, one can obtain the causal impact of a shock to y_j on y_i (the IRF object of interest), or the corresponding fraction of the optimal forecast error variance of y_i due to shocks originating with y_j (the VD object of interest), for any i and j .² But this benefit comes at a cost: the resulting IRFs/VDs can depend importantly on the variable ordering, and the number of possible orderings grows massively (indeed factorially) with the VAR dimension N . Hence N must be very small when using Cholesky-factor IRFs/VDs if one hopes, for example, to check robustness to ordering.

In the generalized approach to IRF/VD identification, the IRFs/VDs are obtained directly from the reduced-form VAR shocks, without orthogonalization (Koop, Pesaran, and

¹We also note the connectedness measurement potential of VDs obtained not via traditional Cholesky or generalized VAR-based methods, but rather via direct local-projection methods, as in Jorda (2005) and Montiel Olea and Plagborg-Møller (2021), and surveyed for example in Jorda and Taylor (2024). We leave exploration of that avenue to future research.

²Rambachan and Shephard (2025) provide a deep analysis of the conditions under which IRFs and VDs may be given causal interpretation, and how those conditions relate to VAR identification.

Potter, 1996; Pesaran and Shin, 1998). The benefits and costs of the generalized approach are precisely opposite those of the orthogonalized approach. On the one hand, the generalized approach effectively treats each variable as if it were first in a Cholesky ordering, so that IRF/VD results do not depend on ordering, which allows for inclusion of many variables in the VAR without worry about robustness to ordering. On the other hand, the generalized IRF/VD results capture co-movement (correlation) but not contagion (causality).

Finally, having obtained VDs, one uses network theory to draw connectedness implications. In particular, a VD matrix may be interpreted as the adjacency matrix of a weighted, directed network (Diebold and Yilmaz, 2014), so that connectedness is characterized by the VD network in- and out-degrees, the degree distribution and its moments (particularly the mean degree), etc.

In this paper we extend the DY connectedness measurement framework to simultaneously incorporate *both* orthogonalized and generalized identification, by allowing for shock clusters (or blocks, or groups, corresponding to different asset classes, industries, regions, etc.), such that shocks are orthogonal across clusters but correlated within clusters. This allows us to bridge the divide between Forbes-Rigobon (2002) “contagion” (empirically captured by orthogonalized IRFs and VDs – that is, N “clusters” of size 1) and “co-movement” (empirically captured by generalized IRFs and VDs – that is, 1 cluster of size N), which emerge as very special cases. Moreover, it pragmatically allows for incorporation of causal ordering while simultaneously keeping the number of possible orderings small, because ordering is relevant only across clusters, not within them.

We proceed as follows. In section 2 we present orthogonalized, generalized, and clustered VAR identifications for IRFs and VDs. The IRF perspective is more effective for introducing the relevant concepts and issues, so we begin with it in section 2.1 in orthogonalized and generalized identification contexts. However, the VD perspective is more effective for actual connectedness measurement, so we move to it in section 2.2, where we introduce specific connectedness measures in clustered contexts. In section 3 we maintain the VD connectedness perspective, illustrating and comparing various connectedness measures, and comparing results from clustered vs generalized identifications, in a detailed empirical exploration of equity markets for sixteen countries spanning three global regions. We conclude in section 4.

2 Measuring Network Connectedness

2.1 Background

In this section we introduce our network framework in several steps. We will refer to network nodes as “assets” or “asset returns”. This serves two purposes. First, it lends economy and concreteness to the discussion, just as with our use of “VARs” rather than “network approximating models”. And second, it sets the stage for our subsequent exploration of global equity markets in section 3. We emphasize, however, that our application to global equities is just one example (albeit a very important example), and that even within the global equity space our framework could be applied not only to returns, but also to return volatilities, liquidities, etc.³

2.1.1 Basic VAR Framework

Consider a covariance-stationary N -variable P^{th} -order VAR,

$$\mathbf{x}_t = \sum_{p=1}^P \Phi_p \mathbf{x}_{t-p} + \mathbf{u}_t = \sum_{i=0}^{\infty} \mathbf{A}_i \mathbf{u}_{t-i}, \quad t = 1, 2, \dots, T, \quad (1)$$

where $\mathbf{x}_t = [x_{1t} \ x_{2t} \ \dots \ x_{Nt}]'$ is an $[N \times 1]$ vector of asset returns and Φ_p is an $[N \times N]$ parameter matrix for lag p . Further, $\mathbb{E}[\mathbf{u}_t] = 0$ and $\mathbb{V}[\mathbf{u}_t] = \Sigma$ for all t , where $\Sigma = \{\sigma_{ij}, i, j = 1, 2, \dots, N\}$ is an $[N \times N]$ symmetric positive semi-definite matrix. Finally, $\mathbf{A}_0$ is an $[N \times N]$ identity matrix, and $\mathbf{A}_i = \Phi_1 \mathbf{A}_{i-1} + \Phi_2 \mathbf{A}_{i-2} + \dots + \Phi_p \mathbf{A}_{i-p}$, $i = 1, 2, \dots$ (with $\mathbf{A}_i$ an $[N \times N]$ zero matrix for $i < 0$). For lower-triangular non-singular $[N \times N]$ matrix $\mathbf{Q}_C$, we can rewrite the moving average representation in (1) without loss of generality as

$$\mathbf{x}_t = \sum_{i=0}^{\infty} \mathbf{A}_i \mathbf{u}_{t-i} = \sum_{i=0}^{\infty} (\mathbf{A}_i \mathbf{Q}_C) (\mathbf{Q}_C^{-1} \mathbf{u}_{t-i}) = \sum_{i=0}^{\infty} (\mathbf{A}_i \mathbf{Q}_C) \boldsymbol{\epsilon}_{t-i},$$

where $\mathbb{E}[\boldsymbol{\epsilon}_t] = \mathbb{E}[\mathbf{Q}_C^{-1} \mathbf{u}_t] = 0$, and $\mathbb{V}[\boldsymbol{\epsilon}_t] = \mathbb{V}[\mathbf{Q}_C^{-1} \mathbf{u}_t] = \boldsymbol{\Omega}_C$ with elements $\{\omega_{C,ij}, i, j = 1, 2, \dots, N\}$.

Assuming linearity of conditional expectations as in Koop, Pesaran, and Potter (1996),

³As an example of a different application area, consider product prices within vs across industry sectors. One can easily imagine a scenario where structural shocks to within-sector prices are correlated, but cross-sector prices are approximately orthogonal.

the N -vector of h -step responses of the elements of $\mathbf{x}$ to a δ_j shock in ϵ_{jt} is

$$\psi_j^C(h) = (\mathbf{A}_h \mathbf{Q}_C) \mathbb{E}[\epsilon_t | \epsilon_{j,t} = \delta_j] = \frac{(\mathbf{A}_h \mathbf{Q}_C) \boldsymbol{\Omega}_C \mathbf{e}_j \delta_j}{\omega_{C,jj}},$$

where $\mathbf{e}_j$ is a $[N \times 1]$ selection vector with one in the j^{th} position and zero elsewhere. Setting $\delta_j = \sqrt{\omega_{C,jj}}$ gives the responses to a one standard deviation shock in ϵ_{jt} ,

$$\psi_j^C(h) = \frac{(\mathbf{A}_h \mathbf{Q}_C) \boldsymbol{\Omega}_C \mathbf{e}_j}{\sqrt{\omega_{C,jj}}}. \quad (2)$$

In concluding this section, let us say a bit more about the nonsingular matrix $\mathbf{Q}_C$, which plays a crucial role in what follows but has not yet been discussed. For now, suffice it to say that C will denote the number of clusters, which determines the structure of $\mathbf{Q}_C$ and hence $\boldsymbol{\Omega}_C$, both of which impact the impulse response (2). While any number of clusters $C \in [1, N]$ could be operative, the literature has so far only focused on two very special cases: $C = N$, where each network node is its own cluster, and $C = 1$, where all nodes are grouped into a single cluster. These two settings translate respectively to orthogonalized and generalized impulse response functions, to which we now turn.

2.1.2 Orthogonalized Impulse Responses ($C = N$)

As the name indicates, in this approach the structural VAR shocks are uncorrelated. This means that each asset represents a cluster and there is no correlation across clusters, i.e. there are as many clusters, C , as there are assets, N . Mathematically, this is achieved by setting $C = N$, which translates into $\mathbf{Q}_N = \mathbf{M}$, where $\mathbf{M}$ is the unique lower-triangular matrix that satisfies the Cholesky decomposition, $\mathbf{M}\mathbf{M}' = \boldsymbol{\Sigma}$. Further, in this approach the variance-covariance matrix of the underlying structural VAR shocks, $\boldsymbol{\Omega}_N = \mathbb{V}[\mathbf{Q}_N^{-1} \mathbf{u}_t] = \mathbb{V}[\mathbf{M}^{-1} \mathbf{u}_t]$, is an $[N \times N]$ identity matrix $\mathbf{I}_N$, so that $\sqrt{\omega_{N,jj}} = 1$ and $\sqrt{\omega_{N,ij}} = 0$ for $i \neq j$. This yields the well known orthogonalized impulse response functions introduced by Sims (1980),

$$\psi_j^N(h) = \frac{\mathbf{A}_h \mathbf{Q}_N \boldsymbol{\Omega}_N \mathbf{e}_j}{\sqrt{\omega_{N,jj}}} = \frac{\mathbf{A}_h \mathbf{M} \mathbf{I}_N \mathbf{e}_j}{\sqrt{1}} = \mathbf{A}_h \mathbf{M} \mathbf{e}_j = \psi_j^o(h), \quad (3)$$

where the superscript o indicates “orthogonalized” and it is understood that $C = N$.

2.1.3 Generalized Impulse Responses ($C = 1$)

Generalized impulse responses go to the opposite extreme. That is, rather than putting each asset in its own cluster and removing all correlation across shocks, as in the orthogonalized approach, generalized impulse responses allow for correlated shocks. In particular, the generalized approach uses the reduced-form VAR shocks depicted in equation (1). Intuitively, this translates to the idea of having only one cluster, i.e. setting $C = 1$, and allowing for correlated shocks within that single cluster. From a mathematical perspective, this translates to $\mathbf{Q}_1 = \mathbf{I}_N$. It follows that $\mathbf{\Omega}_1 = \mathbb{V}[\mathbf{Q}_1^{-1}\mathbf{u}_{t-i}] = \mathbb{V}[\mathbf{I}_N^{-1}\mathbf{u}_{t-i}] = \mathbb{V}[\mathbf{u}_{t-i}] = \mathbf{\Sigma}$ and that $\omega_{1,jj} = \sigma_{jj}$, $\forall j$. This yields the generalized impulse responses introduced by Koop, Pesaran, and Potter (1996),

$$\psi_j^1(h) = \frac{\mathbf{A}_h \mathbf{Q}_1 \mathbf{\Omega}_1 \mathbf{e}_j}{\sqrt{\omega_{1,jj}}} = \frac{\mathbf{A}_h \mathbf{I}_N \mathbf{\Sigma} \mathbf{e}_j}{\sqrt{\omega_{1,jj}}} = \frac{\mathbf{A}_h \mathbf{\Sigma} \mathbf{e}_j}{\sqrt{\sigma_{jj}}} = \psi_j^g(h), \quad (4)$$

where the superscript g indicates “generalized” and it is understood that $C = 1$.

2.1.4 Discussion

The benefits of orthogonalized IRFs stem from their ability to quantify the causal impacts of shocks. But those benefits come at a cost: The outcome can depend crucially on the variable ordering, and the number of possible orderings grows factorially with the number of assets, which prohibits checking IRF robustness to ordering except in very low-dimensional VARs. The benefits and costs of generalized IRFs are precisely opposite. The generalized approach avoids the issue of ordering, but that benefit comes at the cost of quantifying only co-movement, not contagion.

Against this background, in what follows we build on the seminal work of Forbes and Rigobon (2002), who ask whether asset-return connections are better characterized as co-movement or contagion (“spillovers”). We progress by effectively allowing for *both* co-movement and contagion, integrating the orthogonalized and generalized IRF approaches via *clustering*, to which we now turn.

2.2 Clustering

2.2.1 Cluster-Orthogonalized Impulse Responses ($C \in [1, N]$)

Consider the reduced-form VAR given by equation (1), and suppose that its variables can be grouped into C known clusters with similar characteristics, such as asset classes (e.g.,

stocks, bonds, commodities, etc.) or regions (e.g., North America, Europe, East Asia, etc.). When $C = 3$, for example, we might have reduced-form VAR shock covariance matrix

$$\mathbf{\Sigma}_3 = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} & \sigma_{14} & \sigma_{15} & \sigma_{16} & \sigma_{17} & \dots & \sigma_{1N} \\ \sigma_{21} & \sigma_{22} & \sigma_{23} & \sigma_{24} & \sigma_{25} & \sigma_{26} & \sigma_{27} & \dots & \sigma_{2N} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} & \sigma_{34} & \sigma_{35} & \sigma_{36} & \sigma_{37} & \dots & \sigma_{3N} \\ \sigma_{41} & \sigma_{42} & \sigma_{43} & \sigma_{44} & \sigma_{45} & \sigma_{46} & \sigma_{47} & \dots & \sigma_{4N} \\ \sigma_{51} & \sigma_{52} & \sigma_{53} & \sigma_{54} & \sigma_{55} & \sigma_{56} & \sigma_{57} & \dots & \sigma_{5N} \\ \sigma_{61} & \sigma_{62} & \sigma_{63} & \sigma_{64} & \sigma_{65} & \sigma_{66} & \sigma_{67} & \dots & \sigma_{6N} \\ \sigma_{71} & \sigma_{72} & \sigma_{73} & \sigma_{74} & \sigma_{75} & \sigma_{76} & \sigma_{77} & \dots & \sigma_{7N} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \sigma_{N1} & \sigma_{N2} & \sigma_{N3} & \sigma_{N4} & \sigma_{N5} & \sigma_{N6} & \sigma_{N7} & \dots & \sigma_{NN} \end{pmatrix} = \begin{pmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} & \mathbf{\Sigma}_{13} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} & \mathbf{\Sigma}_{23} \\ \mathbf{\Sigma}_{31} & \mathbf{\Sigma}_{32} & \mathbf{\Sigma}_{33} \end{pmatrix} \quad (5)$$

where the clustering is indicated by dashed boxes.⁴

Note that we are not assuming block diagonality of $\mathbf{\Sigma}$, as its off-diagonal blocks are generally non-zero, as in equation (5) immediately above. Rather, we are effectively assuming block diagonality of the corresponding underlying *structural* shock covariance matrix $\mathbf{\Omega}$, which neither generalized nor orthogonalized impulse responses can accommodate, because they both place overly-restrictive structure on $\mathbf{Q}_C$. As we have seen, generalized IRFs consider all assets to be part of a single cluster ($C = 1$) and set $\mathbf{Q}_1 = \mathbf{I}_N$, which forces the underlying structural shock covariance matrix to be $\mathbf{\Omega}_1 = \mathbf{\Sigma}$. Alternatively, orthogonalized IRFs consider each asset to represent a single uncorrelated cluster ($C = N$) and set $\mathbf{Q}_N = \mathbf{M}$, which forces the underlying structural shock covariance matrix to be $\mathbf{\Omega}_N = \mathbf{I}_N$.

What is needed is a $\mathbf{Q}$ matrix consistent with block-diagonality of the underlying structural shock covariance matrix $\mathbf{\Omega}$; that is, a $\mathbf{Q}$ matrix such that $\mathbf{Q}^{-1}$ orthogonalizes reduced-form residuals across, but not within, clusters. Returning to our $C = 3$ example, we need a

⁴We use $C = 3$ clusters purely for illustrative concreteness; there is no loss of generality, as the framework that we will soon discuss holds for any number C of potential clusters (as long as $C \leq N$) and with any number of assets N_c in a given cluster c (as long as $\sum_{c=1}^C N_c = N$).

$\mathbf{Q}_3$ such that $\mathbf{Q}_3^{-1}\mathbf{u}_t$ has covariance matrix

$$\mathbf{\Omega}_3 = \begin{pmatrix} \boxed{\begin{matrix} \omega_{3,11} & \omega_{3,12} & \omega_{3,13} \\ \omega_{3,21} & \omega_{3,22} & \omega_{3,23} \\ \omega_{3,31} & \omega_{3,32} & \omega_{3,33} \end{matrix}} & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & \boxed{\begin{matrix} \omega_{3,44} & \omega_{3,45} \\ \omega_{3,54} & \omega_{3,55} \end{matrix}} & 0 & 0 & \dots & 0 \\ 0 & 0 & \boxed{\begin{matrix} \omega_{3,66} & \omega_{3,67} & \dots & \omega_{3,6N} \\ \omega_{3,76} & \omega_{3,77} & \dots & \omega_{3,7N} \\ \vdots & \vdots & \ddots & \vdots \\ \omega_{3,N6} & \omega_{3,N7} & \dots & \omega_{3,NN} \end{matrix}} & \dots & 0 \end{pmatrix} = \begin{pmatrix} \boxed{\Omega_{3,11}} & 0 & 0 \\ \mathbf{0} & \boxed{\Omega_{3,22}} & 0 \\ \mathbf{0} & 0 & \boxed{\Omega_{3,33}} \end{pmatrix}. \quad (6)$$

In Appendix A we show by sequential linear projection that the relevant $\mathbf{Q}^{-1}$ matrix is

$$\mathbf{Q}_3^{-1} = \begin{pmatrix} \mathbf{I} & 0 & 0 \\ -\Sigma_{21}\Sigma_{11}^{-1} & \mathbf{I} & 0 \\ -\Sigma_{31}\Sigma^{11} - \Sigma_{32}\Sigma^{21} & -\Sigma_{31}\Sigma^{12} - \Sigma_{32}\Sigma^{22} & \mathbf{I} \end{pmatrix}, \quad (7)$$

where

$$\begin{pmatrix} \Sigma^{11} & \Sigma^{12} \\ \Sigma^{21} & \Sigma^{22} \end{pmatrix} \equiv \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}^{-1}, \quad (8)$$

and we provide formulae for the Σ^{ij} 's.

2.2.2 Cluster-Orthogonalized Variance Decompositions

Thus far we have focused exclusively on IRFs, where basic issues and identification concepts are most easily introduced, but forecast error variance decompositions (VDs), which are simple transformations of IRFs, turn out to be more appealing for constructing and applying actual connectedness measures. First, like IRFs, VDs make obvious intuitive sense and answer a key connectedness question, namely (at the most granular pairwise level) ‘‘How much of the H -step-ahead uncertainty in asset return i is due to shocks originating from return j ?’’ Second, VDs also easily allow for levels of cross-sectional aggregation beyond pairwise, answering broader questions like ‘‘How much of the future uncertainty in one return is due to

shocks from *all other* returns?”⁵ Third, VDs easily allow not only for cross-sectional aggregation, but also for temporal aggregation, via different connectedness strengths at different horizons H , facilitating examination of a variety of horizons (and selection of a preferred horizon if desired). Finally, the matrix of VDs can be viewed as the adjacency matrix of a weighted directed network, as emphasized in Diebold and Yilmaz (2014), bringing powerful network perspectives and tools in touch with connectedness measurement.

We denote the H -step-ahead VD by $\tilde{\theta}_{ij}^C(H)$:

$$\tilde{\theta}_{ij}^C(H) = \frac{\sum_{h=0}^{H-1} (\mathbf{e}_i' \boldsymbol{\psi}_j^C(h))^2}{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \boldsymbol{\Sigma} \mathbf{A}_h' \mathbf{e}_i)^2} = \frac{\omega_{C,jj}^{-1} \sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \mathbf{Q}_C \boldsymbol{\Omega}_C \mathbf{e}_j)^2}{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \boldsymbol{\Sigma} \mathbf{A}_h' \mathbf{e}_i)^2}, \quad (9)$$

where $\tilde{\theta}_{ij}^C$ is the share of the H -step-ahead forecast error variance of asset i due to shocks from asset j . In parallel to the IRF equations (3) and (4), the VD equation (9) nests both orthogonalized and generalized versions:

$$\tilde{\theta}_{ij}^o(H) = \tilde{\theta}_{ij}^N(H) = \frac{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \mathbf{M} \mathbf{e}_j)^2}{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \boldsymbol{\Sigma} \mathbf{A}_h' \mathbf{e}_i)^2}$$

and

$$\tilde{\theta}_{ij}^g(H) = \tilde{\theta}_{ij}^1(H) = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \boldsymbol{\Sigma} \mathbf{e}_j)^2}{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \boldsymbol{\Sigma} \mathbf{A}_h' \mathbf{e}_i)^2}.$$

We note that $\sum_{j=1}^N \tilde{\theta}_{ij}^o(H) = 1$, while generally $\sum_{j=1}^N \tilde{\theta}_{ij}^C(H) \neq 1$, and indeed $\sum_{j=1}^N \tilde{\theta}_{ij}^g(H) \neq 1$. This is due to the non-zero covariance of residuals in the case of clustered and generalized shocks. However, in line with Diebold and Yilmaz (2012), we can normalize to produce $\theta_{ij}^C(H) = \frac{\tilde{\theta}_{ij}^C(H)}{\sum_{j=1}^N \tilde{\theta}_{ij}^C(H)}$ and $\theta_{ij}^g(H) = \frac{\tilde{\theta}_{ij}^g(H)}{\sum_{j=1}^N \tilde{\theta}_{ij}^g(H)}$, so that $\sum_{j=1}^N \theta_{ij}^C(H) = 1$, and $\sum_{j=1}^N \theta_{ij}^g(H) = 1$.

A graphical VD illustration that matches the three-cluster structure of $\boldsymbol{\Sigma}$ in equation (5) and $\boldsymbol{\Omega}$ in equation (6) is

⁵In contrast, aggregative connectedness measurement is trickier with IRFs, which as routinely studied have a pairwise orientation. Hence, for example, attempts at IRF aggregation must confront the fact that positive and negative responses can offset, unlike variance shares, all of which must be positive.

$$\Theta^C(H) = \begin{pmatrix} \begin{array}{|c|c|c|} \hline \theta_{11} & \theta_{12} & \theta_{13} \\ \hline \theta_{21} & \theta_{22} & \theta_{23} \\ \hline \theta_{31} & \theta_{32} & \theta_{33} \\ \hline \end{array} & \begin{array}{|c|c|} \hline \theta_{14} & \theta_{15} \\ \hline \theta_{24} & \theta_{25} \\ \hline \theta_{34} & \theta_{35} \\ \hline \end{array} & \begin{array}{|c|c|c|c|} \hline \theta_{16} & \theta_{17} & \dots & \theta_{1N} \\ \hline \theta_{26} & \theta_{27} & \dots & \theta_{2N} \\ \hline \theta_{36} & \theta_{37} & \dots & \theta_{3N} \\ \hline \end{array} \\ \hline \begin{array}{|c|c|c|} \hline \theta_{41} & \theta_{42} & \theta_{43} \\ \hline \theta_{51} & \theta_{52} & \theta_{53} \\ \hline \end{array} & \begin{array}{|c|c|} \hline \theta_{44} & \theta_{45} \\ \hline \theta_{54} & \theta_{55} \\ \hline \end{array} & \begin{array}{|c|c|c|c|} \hline \theta_{46} & \theta_{47} & \dots & \theta_{4N} \\ \hline \theta_{56} & \theta_{57} & \dots & \theta_{5N} \\ \hline \end{array} \\ \hline \begin{array}{|c|c|c|} \hline \theta_{61} & \theta_{62} & \theta_{63} \\ \hline \theta_{71} & \theta_{72} & \theta_{73} \\ \hline \vdots & \vdots & \vdots \\ \hline \theta_{N1} & \theta_{N2} & \theta_{N3} \\ \hline \end{array} & \begin{array}{|c|c|} \hline \theta_{64} & \theta_{65} \\ \hline \theta_{74} & \theta_{75} \\ \hline \vdots & \vdots \\ \hline \theta_{N4} & \theta_{N5} \\ \hline \end{array} & \begin{array}{|c|c|c|c|} \hline \theta_{66} & \theta_{67} & \dots & \theta_{6N} \\ \hline \theta_{76} & \theta_{77} & \dots & \theta_{7N} \\ \hline \vdots & \vdots & \ddots & \vdots \\ \hline \theta_{N6} & \theta_{N7} & \dots & \theta_{NN} \\ \hline \end{array} \end{pmatrix} \quad (10)$$

where the dark gray, light gray, and white boxes denote, respectively, own variance shares, co-movement shares, and contagion shares. The own variance shares capture the fraction of the forecast-error variance of asset i due to shocks from asset i itself; the co-movement shares capture the fraction due to shocks from other assets in the same cluster; and the contagion shares capture the fraction due to shocks from other assets in other clusters. In contrast, computing VDs with orthogonalized impulse responses yields only own variance and contagion shares, and computing VDs with generalized impulse responses yields only own variance and co-movement shares.

The distinction between co-movement and contagion stems from the correlation of residuals. Within a cluster, the structural residuals are correlated, so we are unable to pinpoint the shock to a given asset; instead we observe only co-movement (light gray areas). However, due to the constraints imposed via $\mathbf{Q}_C$, the structural residuals are uncorrelated across clusters. This means that we are able to narrow down the origin of a shock to a given cluster, and quantify the reverberation across clusters. The absence of cross-cluster correlation of the structural residuals and the resulting causality of spillovers translate into a quantifiable contagion across clusters (white areas).

2.2.3 Connectedness Measurement Within and Across Clusters

Following Diebold and Yilmaz (2012), we now define empirical average connectedness measures that parallel the theoretical concepts sketched above. We define the own variance share for a cluster c , which captures how much of the forecast error variance of cluster c is due to shocks specific to that same cluster, as $\Theta_c^{own} = \frac{1}{N_c} \sum_{i \in c} \theta_{ii}$, where N_c is the number

of members of cluster c .

In addition to own variance shares, we also measure connectedness across assets, but unlike Diebold and Yilmaz (2012), who use generalized IRFs to obtain VDs, our clustering framework allows us to distinguish between two types of connectedness: Co-movement (within-cluster) and contagion (cross-cluster). Co-movement shares capture the extent to which the forecast-error variance of cluster c is driven by co-movements among assets in that cluster, $\Theta_c^{comove} = \frac{1}{N_c} \sum_{\substack{i,j \in c \\ i \neq j}} \theta_{ij}$. Alternatively, contagion shares capture the extent to which the forecast-error variance of cluster c is driven by shocks from another cluster k , $\Theta_{c \leftarrow k}^{contag} = \frac{1}{N_c} \sum_{i \in c} \left(\frac{1}{N_k} \sum_{j \in k} \theta_{ij} \right)$. The total contagion received by cluster c is then $\Theta_{c \leftarrow \bullet}^{contag} = \sum_{k \neq c} \Theta_{c \leftarrow k}^{contag}$. It will also prove useful to consider cross-cluster averages of the above own, co-movement, and total contagion measures. We write $\Theta^{own} = \frac{1}{C} \sum_{c=1}^C \Theta_c^{own}$, $\Theta^{comove} = \frac{1}{C} \sum_{c=1}^C \Theta_c^{comove}$, and $\Theta^{contag} = \frac{1}{C} \sum_{c=1}^C \Theta_{c \leftarrow \bullet}^{contag}$. Note that $\Theta^{own} + \Theta^{comove} + \Theta^{contag} = 1$.

3 Clustered Connectedness in Global Equity Markets

Global equity markets are likely connected both locally (within regions) and globally (across regions), but the strengths and directions of connectedness are generally unknown, and moreover, they may be time-varying. Simultaneously, improved quantitative characterization of market network connectedness would be of value not only to academic economists (of course), but also to a variety of financial-market participants, including private-sector agents (e.g., for improved portfolio allocation, risk management, and business planning), policymakers (e.g., for improved anticipation and tracking of cross-market spillover episodes as, for example, in the financial crises of 2007-9), and regulators (e.g., for improved monitoring of the effects of balance sheet and other linkages among financial institutions and trading exchanges). Additional beneficiaries include those not directly involved in financial markets, but who may nevertheless want to use the markets to help assess the effects of non-financial policies, such as imposition of tariffs or sanctions).

Before proceeding, however, we emphasize that our approach is intentionally non-structural. Rather, it is meant to be a flexible reduced-form description of a set of covariance-stationary series, subject only to the assumed block-diagonality of the underlying structural shock covariance matrix, *consistent with whatever (generally unknown and potentially time-varying) structural DGP is operative*. Relative to a more structural approach, our approach has both potential benefits (no damage from imposing false asset-pricing theory restrictions) and potential costs (less efficiency from failing to impose true asset-pricing theory restrictions). We

feel that the benefits of our approach significantly outweigh the costs, particularly given the strikingly poor empirical performance of structural asset-pricing models in equity premium prediction, as in the classic work of Goyal and Welch (2008).

Factor structure, for example, is one such popular asset-pricing theory restriction (assumption), which again may or may not be valid. It implies restrictions on a reduced-form VAR and is therefore compatible with, and encompassed by, a VAR approach, as shown by Stock and Watson (2005). If linear factor structure holds, then a VAR will capture it, and if one *knows* that such structure holds, then imposing it on the VAR may produce some efficiency gains, but one never knows. And even within the narrow world of linear factor models there are many possibilities – single-factor, multi-factor, strong, weak, exact, approximate, etc.

Against this background, in this section we use our clustering framework to study connectedness in sixteen country equity markets spanning three global regions. We proceed as follows. In section 3.1 we discuss estimation of the network-approximating VAR and provide full-sample analyses. In section 3.2 we provide rolling-sample analyses. Finally, in section 3.3 we discuss differences in results under clustered vs generalized identification.

3.1 Full-Sample Connectedness

Here we characterize country equity market connectedness using the full data sample. If the structure of connectedness is fixed over time, then full-sample estimation is of immediate and unique interest, and even if it varies over time, the full-sample estimates provide a “time-averaged” or “unconditional” summary. Later, in section 3.2, we will explicitly allow for time-variation in conditional connectedness via rolling-sample analysis.

3.1.1 Country Return Data and Network VAR Estimation

Our sample includes sixteen countries spanning three global regions: North America (U.S. and Canada), Europe (Germany, France, United Kingdom, Portugal, Italy, Ireland, Greece, and Spain), and East Asia (Japan, China, South Korea, Taiwan, Hong Kong, and Singapore). The sample period is 10 July 2002 (corresponding roughly to the full launch of the Euro) through 29 December 2021.

We construct the sixteen series of weekly nominal local-currency equity market returns. Our use of local returns is intentional and reflects common practice among portfolio managers, who often invest internationally while intentionally failing to fully hedge currency

exposure. The reason is that although currency risk is potentially an important factor in long-horizon risk-adjusted returns, the hedging decision also depends on a variety of other factors, including hedging costs, risk tolerance, investment objectives, and investment horizon. In addition, as a pragmatic matter, although currencies obviously fluctuate significantly over long horizons, their fluctuations have little impact on higher-frequency (daily or weekly) stock returns.

We proceed by taking daily local-currency equity market indices, P_t , from the Wharton Research Data Services (WRDS) database, and converting them to daily (log) returns r_t , using $r_t = \Delta \log P_t$.⁶ Next, we convert the returns from daily to weekly by cumulating the daily returns from Thursday through Wednesday each week. We use Thursday to Wednesday to avoid distortions due to beginning-of-week and end-of-week trades.

In Table 1 we provide return summary statistics, grouped by region, which in this application we naturally take to be the relevant continent. Mean returns across markets are sometimes positive and sometimes negative, but generally near zero, with standard deviations much larger. Skewnesses, like means, are small and of mixed sign, whereas kurtoses are generally larger and well above three, consistent with the well-known fat tails in high- and medium-frequency asset returns.

Following Demirer et al. (2018), we proceed with equation-by-equation estimation of a 16-variable VAR(3) using an adaptive elastic net (Zou and Zhang, 2009) for regularization (shrinkage and selection), which is generally preferred to related penalized estimators (e.g., lasso (Tibshirani, 1996)) in environments like ours with correlated predictors and grouped structure. In particular, for each equation we solve

$$\hat{\beta} = \operatorname{argmin}_{\beta} \left(\sum_{t=1}^T \left(y_t - \sum_i \beta_i x_{it} \right)^2 + \lambda \sum_{i=1}^K w_i \left(\frac{1}{2} |\beta_i| + \frac{1}{2} \beta_i^2 \right) \right), \quad (11)$$

where $w_i = 1/|\hat{\beta}_{i,OLS}|$ and λ is selected equation-by-equation by 10-fold cross-validation.⁷ Once all equations of the VAR have been estimated, we obtain the residuals for each equation, from which we obtain in the usual way an estimate $\hat{\Sigma}$ of the reduced-form shock covariance matrix.

We assume that clustering is by region. Market ordering within the three clusters is

⁶See <https://wrds-www.wharton.upenn.edu/>.


⁷The adaptive elastic net penalty averages the LASSO penalty with a ridge penalty, and moreover it weights the average by inverse OLS parameter estimates, thereby shrinking the smallest OLS-estimated coefficients most heavily toward zero.

Table 1: Summary Statistics
Sixteen Weekly Country Equity Market Returns

Region	Country	Label	Mean	Std	Info	Skew	Kurt
North America	All		6.96	15.60	0.45	-1.42	12.74
North America	United States	USA	8.25	16.65	0.50	-1.02	10.18
North America	Canada	CAN	5.68	16.09	0.35	-1.52	15.14
Europe	All		1.17	18.44	0.06	-0.98	9.07
Europe	France	FRA	4.91	19.75	0.25	-0.69	10.24
Europe	Germany	GER	4.43	19.41	0.23	-1.00	9.97
Europe	United Kingdom	GBR	2.85	16.47	0.17	-0.79	9.40
Europe	Portugal	PRT	-0.36	19.60	-0.02	-1.03	8.66
Europe	Ireland	IRL	3.76	23.79	0.16	-0.95	10.86
Europe	Italy	ITA	0.39	21.18	0.02	-0.76	7.71
Europe	Greece	GRC	-8.32	31.06	-0.27	-0.66	8.89
Europe	Spain	SPA	1.69	20.56	0.08	-0.47	6.70
East Asia	All		4.79	15.52	0.31	-0.83	8.97
East Asia	Japan	JPN	3.74	19.78	0.19	-0.70	8.13
East Asia	China	CHN	5.04	24.59	0.20	-0.69	8.08
East Asia	South Korea	KOR	6.71	20.48	0.33	-0.60	10.02
East Asia	Taiwan	TWA	5.61	19.05	0.29	-0.58	7.49
East Asia	Hong Kong	HKG	4.97	20.52	0.24	-0.46	7.12
East Asia	Singapore	SGP	2.66	15.85	0.17	-0.76	10.87
Global	All		3.25	15.60	0.21	-1.16	10.06

Notes: We present summary statistics for annualized weekly local-currency nominal equity returns for sixteen country markets. The weekly returns are from Thursday to Wednesday, and the sample period is 10 July 2002 through 29 December 2021. “Mean” denotes the (sample) mean, “Std” denotes standard deviation, “Info” denotes the information ratio (Mean/Std), “Skew” denotes skewness, and “Kurt” denotes kurtosis. See text for details.

Figure 1: Estimated Global Equity Market Return Network Graph, Clustered Identification



Notes: We show the estimated network spring graph obtained from the ForceAtlas2 algorithm. See text for details.

irrelevant for estimated VD network structure, but cluster ordering is potentially relevant, and there are $3! = 6$ possible cluster orderings. For a given cluster ordering, the appropriate “clusterizing” (as opposed to “orthogonalizing”) transformation matrix $\hat{Q}_3^{-1}$ is

$$\hat{Q}_3^{-1} = \begin{pmatrix} \mathbf{I} & 0 & 0 \\ -\hat{\Sigma}_{21}\hat{\Sigma}_{11}^{-1} & \mathbf{I} & 0 \\ -\hat{\Sigma}_{31}\hat{\Sigma}^{11} - \hat{\Sigma}_{32}\hat{\Sigma}^{21} & -\hat{\Sigma}_{31}\hat{\Sigma}^{12} - \hat{\Sigma}_{32}\hat{\Sigma}^{22} & \mathbf{I} \end{pmatrix}, \quad (12)$$


as per equation (7) and Appendix A. We then obtain VDs from the clusterized VAR, using a VD horizon of $h = 12$ weeks. There are of course six cluster orderings that can be used to produce clusterized VDs. In an effort to remain agnostic regarding orderings, we calculate VDs for all possible orderings and report averages.

3.1.2 The Market Network Graph with Clustered Identification

We visualize estimated networks using “spring graphs” obtained from the ForceAtlas2 algorithm of Jacomy et al. (2014), as implemented in the open-source Gephi software.⁸ The


⁸See <https://gephi.github.io/>.

Figure 2: Cross-Country Total Directional Connectedness Densities, Clustered Identification




Notes: We show full-sample kernel density estimates for three total directional connectedness measures (to, from, and net) across sixteen country equity markets, using clustered identification. The “to” density is solid black, the “from” density is dashed black, and the “net” density is red. See text for details.

Figure 3: Estimated Global Equity Market Return Network Graph, Generalized Identification



Notes: We show the estimated network spring graph obtained from the ForceAtlas2 algorithm. See text for details.

Figure 4: Cross-Country Total Directional Connectedness Densities, Generalized Identification



Notes: We show full-sample kernel density estimates for three total directional connectedness measures (to, from, and net) across sixteen country equity markets, using generalized identification. The “to” density is solid black, the “from” density is dashed black, and the “net” density is red. See text for details.

algorithm finds a steady state in which repelling and attracting forces exactly balance, where nodes repel each other like similar poles of two magnets, while edges (links), attract their nodes like springs, with the attracting force proportional to average pairwise directional connectedness “to” and “from.”⁹

There are six associated graph components: Node label, node size, node color, edge thickness, edge color, and edge arrow size (two per edge, because the network is directed).¹⁰ Node label indicates the country as shown in Table 1. Node size indicates the total capitalization of the country’s equity market, obtained from the World Bank’s World Development Indicators database. Node color, very importantly, indicates the net directional connectedness to others, ranging from bright green, the weakest, to vivid yellow, to brick yellow, to bright red, to dark red, the strongest. Edge thickness indicates the average directional pairwise connectedness between two nodes. As it is not always easy to discern the thickness difference between two edges, we also use edge color to indicate the average directional pairwise connectedness between two nodes. Edge color follows a similar scale to node color, starting with light gray, the weakest, followed by vivid yellow, brick yellow, and red, the strongest. Edge arrow size, also very importantly, indicates the pairwise directional connectedness from

⁹Steady state node locations depend on initial node locations and are therefore not unique, but that is largely irrelevant for us, as we are interested in relative, not absolute, node locations in equilibrium.

¹⁰For details see Demirer, Diebold, Liu, and Yilmaz (2018).

one node to the other.

In Figure 1 we present the network spring graph for our clustered identification. Several aspects of the network are apparent. First, North America clusters together, Europe clusters together, and East Asia clusters together, but with two major outliers: GRC for Europe and CHN for East Asia. Alternatively, another three-cluster interpretation could be: Anglo-American (USA, CAN, GBR), “Core Europe” (GER, FRA, SPA, ITA, IRL, PRT), and “Core East Asia” (SGP, HKG, KOR, TWA, JPN), again with GRC and CHN as outliers.

Second, regardless of which interpretation one adopts, GBR plays a key role in linking North America and Europe. In particular, both USA and CAN have strong pairwise directional connectedness to GBR, which then links strongly to FRA and onward to the rest of Europe (except GRC). (There are also strong directional links (large arrows) from USA to GER, FRA, and ITA.)

Finally, most obviously and importantly, North America sits squarely in the network graph center, with red nodes indicating very high net directional connectedness to others. That is, on balance North America sends large amounts of 12-week-ahead uncertainty to others. We have mentioned already the strong pairwise directional connectedness from North America to the European countries of GBR, FRA, GER, and ITA, and there is similarly strong directional connectedness (large, if not red, arrows) from North America to all East Asian countries except CHN.

In closing this section, we highlight an additional important aspect of the country equity market network graph: Total directional connectedness “to” others and “from” others (and their difference, “net” total directional connectedness). The total directional measures are of course implicit in the spring graph of Figure 1, which provides a complete network characterization, but it is impossible to extract them visually. Hence in Figure 2 we supplement the spring graph with estimates of the cross-country densities of total directional connectedness (to, from, and net). The “to” density has a similar mean but is more dispersed than the “from” density; that is, “uncertainty transmissions” range more widely across countries than do “uncertainty receipts”. In addition, both the “to” and “from” densities are skewed left; indeed the “from” density has a small second left mode. Hence the “net” (“to” minus “from”; that is, net transmissions) density is centered near zero but skewed right, with a second right mode corresponding to a few countries with large net transmissions (USA, Canada, and France).

3.1.3 Benchmarking Clustered Identification

For comparison to Figure 1, we show the network graph for generalized as opposed to clustered identification in Figure 3. The two graphs have both similarities and differences. Let us begin with similarities. In both graphs there are clear connectedness clusters for Europe and East Asia, with countries in both regions located closely together. In both graphs GRC and CHN are outliers, indicating relatively weak connections to other equity markets, including those in their own regions of Europe and East Asia, respectively. Finally, in both graphs GRC and CHN feature bright green nodes, indicating that they are net recipients of future uncertainty from other countries.

Now let us consider differences between the generalized and clustered network graphs. Most importantly and obviously, North America is located on the outskirts of the generalized identification graph, which indicates that it has only weak links to other equity markets and, hence, a marginal role in the global equity market return network. Furthermore, its brown node also indicates that its net connectedness to other markets is rather low. Recall that, in contrast, in the clustered identification graph North America sits at the center, with a red node, indicating that it plays a crucial role in the global return network – indeed it is the most significant generator of net return connectedness among all markets.

Finally, we show total directional connectedness densities under generalized identification in Figure 4, for comparison to the densities under clustered identification in Figure 2. The situation under generalized identification differs greatly. In particular, the “to” density (“uncertainty transmissions”) and “from” density (“uncertainty receipts”) are very similar under generalized identification, which makes the “net” (transmissions) density tightly centered around zero. That is, under generalized identification no countries are identified as disproportionately large net transmitters, in contrast to the clear identification of USA, CAN, and FRA as large net transmitters under clustered identification.

3.2 Rolling-Sample Connectedness

We now allow for time-variation in connectedness by changing from full-sample estimation to rolling-sample estimation. The window width for rolling-sample estimation is 2 years (104 weeks). In addition, it will now prove useful to present results for various aspects of clustered connectedness (CC) and generalized connectedness (GC) simultaneously, rather than sequentially as we did for full-sample estimation.

3.2.1 System-Wide Connectedness and its Components

In Figure 5 we show system-wide connectedness (the sum of all off-diagonal VD matrix elements), together with its within-cluster and cross-cluster components (sums of all off-diagonal VD matrix elements inside clusters and outside clusters, respectively).

Let us first discuss the CC results in the upper panel of Figure 5. System-wide CC has two prominent movements, first a large and multi-year increase in 2007-2009 during the global financial crisis, and then a large sharp increase in early 2020 as the global pandemic emerged. Examination of the within- and cross-cluster components reveals that the system-wide CC movements are driven largely by clear and pronounced movements in the cross-cluster component. The within-cluster component, in contrast, is quite stable.

Now let us compare the just-discussed upper-panel CC results to the lower-panel GC results. The system-wide CC and GC movements are clearly very similar, with CC always below GC. This is expected, because the GC approach allows for simultaneous shocks to all variables, both within and across regions, whereas the CC approach imposes uncorrelated shocks across regions, as we discuss in greater detail in section 3.3 below. System-wide CC and GC are closest during the global financial crisis (following the collapse of Lehman Brothers in mid-September 2008) and the spread of the COVID-19 pandemic to the West (in early March 2020). Closer inspection, however, reveals a key difference between the CC and GC measures: Movements in CC are sharper and more pronounced than those of GC, with system-wide CC (and its key driver, cross-cluster CC) varying over a wider range.

3.2.2 Regional Net Directional Connectedness


In Figure 6 we show regional net directional connectedness “to” others (that is, net transmissions of future uncertainty from one region to the other two – the sum of all out-of-region VD matrix elements in the region’s columns, minus the sum of all out-of-region VD matrix elements in the region’s rows), for the North America, Europe, and East Asia regions.¹¹ The top panel of the figure is based on cluster identification (regional net CC), and the bottom panel is based on generalized identification (regional net GC).

Let us first consider the regional net CC shown in the top panel of Figure 6, starting with North America. Throughout the sample, North American net transmissions to Europe and East Asia are positive and typically very large (and often huge) relative to European and East Asian net transmissions to North America. Key episodes include:

¹¹We normalize by the number of countries in the transmitting region.

Figure 5: System-Wide Connectedness and its Components

(a) Clustered Identification



(b) Generalized Identification






Figure 6: Regional Net Directional Connectedness

(a) Clustered Identification



(b) Generalized Identification



1. The financial crisis of 2007-2009. Following the summer 2007 escalation of tensions in the U.S. mortgage and financial markets, North American net CC climbed significantly by the late 2007 and surged following the collapse of Lehman Brothers in late 2008, peaking in early 2009 before dropping.
2. The financial crises of 2010-2014. As the North American crisis moved to Europe, it created a hump-shaped North American net transmissions trajectory that started in 2010, peaked in 2011-2012, and subsided by 2015, linked to the series of European crises in Greece, Portugal, Ireland, Iceland, Italy, and Spain, with the North American transmissions absorbed almost exclusively by Europe.
3. The onset of the COVID-19 pandemic in 2020. COVID-19 burst into the Western Hemisphere in early 2020, producing a huge increase in North American net transmissions.

Europe and East Asia, to which we now turn, had very different net CC experiences.

In contrast to the typically large, positive, and fluctuating values of North American net CC, European and East Asian net CC are typically small, negative, and stable. Europe and East Asia are largely net recipients of transmissions from North America (i.e., they have negative net CC). European net receipts, for example, increase sharply (i.e., European net CC decreases sharply, becoming even more negative) during the major North American net transmissions episodes sketched above.


Now let us compare the just-discussed regional net CC results to the GC results in the lower panel of Figure 6. In general the movements in GC are less pronounced than those of CC, particularly for North America, just as was the case earlier for system-wide connectedness and its components in Figure 5. The 2020 pandemic outbreak, for example, is hardly noticeable in GC North American net transmissions.

Moreover, there are important GC vs CC differences in regional net transmissions well beyond the lower resolution of North American GC movements. In particular, North American transmissions fluctuate around zero under the GC, Europe's are consistently positive, and East Asia's are consistently negative.

3.3 On Connectedness Under Clustered vs Generalized Identification


Having provided results for both full-sample and rolling-sample clustered connectedness, and having emphasized differences under clustered vs generalized identification, we now pro-

Figure 7: Correlation Matrix,
Sixteen Weekly Country Equity Market Returns



Notes to figure: Shading indicates strength of correlation, with brighter shades indicating higher correlation.

Figure 8: Differences in Connectedness
Generalized Minus Clustered Identification



Notes to figure: We show the changes in connectedness (total system-wide, within-cluster, and cross-cluster) when moving from clustered identification to generalized identification.

vide some additional insight into the reasons for the differences, the essence of which is that generalized identification captures co-movement across nodes, but not contagion. As such, both within- and cross-cluster connectedness under generalized identification summarize aspects of correlation, but not causality, so that the clustered and generalized approaches can produce very different results.

In the absence of orthogonalization across regions (i.e., under generalized identification), shocks are subject to a feedback loop that smooths them across the system, so that shocks cannot be properly attributed to their origin.¹² To see why, note that USA appears to be less central to the global equity markets than GER, FRA, SPA, and ITA in the generalized identification network graph of Figure 3, and simultaneously that those four European countries have the highest pairwise correlations in the dataset as shown in the correlation matrix of Figure 7. What happens is that, in the absence of orthogonalization across regions, the

¹²Closely related, in the absence of orthogonalization across regions, the densities of the “from” and “to” directional connectedness measures become diffuse and difficult to distinguish, as is clear from comparing their shapes under clustered identification in Figure 2 to those under generalized identification in Figure 4.

high correlations in Europe capture the otherwise causal connectedness of USA. In contrast, the orthogonalization across regions embedded in the clustered approach links co-movement to within-cluster connectedness and contagion to cross-cluster connectedness, in which case USA emerges, as expected, as the key player in global equity markets as shown in Figure 1.

We show in Figure 8 that total connectedness is always greater under generalized identification; that is, the total connectedness difference – “total generalized minus total clustered” – is always positive. The reason is that shocks reverberate more across the system when not orthogonalized by region, translating into greater connectedness. Most of the total difference stems from the underlying cross-cluster difference, which again is always positive, because cross-cluster connectedness under generalized identification reflects co-movement in addition to contagion. In contrast, the within-cluster difference is typically near zero, with two key exceptions: the European debt crises of 2010-2014 and the emergence of COVID-19 in 2020. In each case, the within-cluster difference rises but the cross-cluster difference falls.

4 Summary and Directions for Future Research

Network connectedness and its evolution are central in economics and finance, and a large literature has arisen that explores connectedness measurement based on variance decompositions from VARs. However, those VARs are typically identified using full orthogonalization or no orthogonalization, which, although useful, are special and extreme cases of the more general and empirically-realistic “clustered orthogonalization” approach developed in this paper, which allows for correlated structural shocks within clusters (e.g., asset classes, industries, regions, etc.) while maintaining orthogonality across clusters, thereby facilitating a nuanced empirical exploration of the “contagion vs co-movement” distinction emphasized by Forbes and Rigobon (2002).

We used our clustered-connectedness framework to explore linkages in global equity returns for sixteen countries in three regions (North America, Europe, and East Asia). There are several key results. First, we identified major players (large net senders of future uncertainty) on both global and local scales. Under clustered identification, the two North American countries, USA and CAN, are the largest global net senders by far, and FRA is a key European net sender.

Second, we documented important time variation in connectedness. Under clustered identification, system-wide connectedness varies importantly, and its cross-cluster (as opposed to within-cluster) component is responsible for most of the variation. Regional net directional

connectedness also varies importantly for the key regional net sender, North America, but much less so for the Europe and East Asia.

Finally, we found important differences in connectedness patterns for clustered vs generalized identifications, and we provided an explanation. Generalized identification is unable to uncover causal connections, instead attributing all node connections to simple correlation.

Now, in closing, let us mention some additional literature links and some promising directions for future research. Regarding additional literature links, it is interesting to note that the clustered network connectedness measurement methods introduced in this paper are part of a wave of recent econometric contributions addressing measurement in other clustered contexts. One prominent example is the now-large literature on estimation of panel data models with clustered fixed effects, beginning with Bonhomme and Manresa (2015). Another prominent example is estimation with clustered covariance matrices in both cross sections (MacKinnon, Nielsen, and Webb, 2023) and time series (Tong, Hansen, and Archakov, 2024).

As for future research, one direction is empirical cluster classification, whether from a frequentist perspective in the tradition of Bonhomme, Lamadon, and Manresa (2022) and Chiang, Sasaki, and Wang (2025), or from a Bayesian perspective as in Zhang (2024). Of particular interest may be improved methods for detecting sender/receiver clusters, as in Gudmundsson and Brownlees (2021) and Brownlees, Gudmundsson, and Lugosi (2022). An elegant approach, incorporating not only clustering but also parameter variation, is provided by Li et al. (2025), who build on Zhu et al. (2017).

An important related issue is cluster misclassification, regardless of whether the clusters are determined endogenously or specified exogenously. In this paper we have assumed that the DGP features clustering, and we have emphasized the network estimation distortions produced in that environment when non-clustered identification (most notably, generalized identification) is used. Conversely, however, clustered identification may produce distortions when clustering is imposed on a non-clustered DGP, or more generally when an adopted clustering pattern differs from the actual pattern.

Appendices

A Cluster Orthogonalization by Linear Projection

Here we derive the matrix $\mathbf{Q}_C^{-1}$ such that transforming the vector of VAR residuals $\mathbf{u}_t$ by $\mathbf{Q}_C^{-1}$ orthogonalizes them across C clusters. For clarity we display results for $C = 3$; extension to $C > 3$ is immediate but more tedious. That is, we seek $\mathbf{Q}_3^{-1}$ such that

$$\mathbb{V}[\mathbf{Q}_3^{-1}\mathbf{u}_t] = \mathbb{V}[\boldsymbol{\epsilon}_t] = \boldsymbol{\Omega} = \begin{pmatrix} \boldsymbol{\Omega}_{11} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Omega}_{22} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \boldsymbol{\Omega}_{33} \end{pmatrix},$$

where

$$\mathbb{V}[\mathbf{u}_t] = \boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} & \boldsymbol{\Sigma}_{13} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} & \boldsymbol{\Sigma}_{23} \\ \boldsymbol{\Sigma}_{31} & \boldsymbol{\Sigma}_{32} & \boldsymbol{\Sigma}_{33} \end{pmatrix},$$

and $\mathbf{0}$ denotes a matrix of zeros.

We begin with $\boldsymbol{\epsilon}_{1,t}$ and proceed sequentially, orthogonalizing residuals across clusters by linear projection, precisely as in the well-known Gram-Schmidt procedure:

$$\boldsymbol{\epsilon}_{1,t} = \mathbf{u}_{1,t} \tag{A.1}$$

$$\boldsymbol{\epsilon}_{2,t} = \mathbf{u}_{2,t} - \boldsymbol{\Sigma}_{21}\boldsymbol{\Sigma}_{11}^{-1}\mathbf{u}_{1,t} \tag{A.2}$$

$$\begin{aligned} \boldsymbol{\epsilon}_{3,t} &= \mathbf{u}_{3,t} - \begin{pmatrix} \boldsymbol{\Sigma}_{31} & \boldsymbol{\Sigma}_{32} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \end{pmatrix} \\ &= \mathbf{u}_{3,t} - \begin{pmatrix} \boldsymbol{\Sigma}_{31} & \boldsymbol{\Sigma}_{32} \end{pmatrix} \begin{pmatrix} \boldsymbol{\Sigma}^{11} & \boldsymbol{\Sigma}^{12} \\ \boldsymbol{\Sigma}^{21} & \boldsymbol{\Sigma}^{22} \end{pmatrix} \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \end{pmatrix} \\ &= \mathbf{u}_{3,t} - \begin{pmatrix} \boldsymbol{\Sigma}_{31}\boldsymbol{\Sigma}^{11} + \boldsymbol{\Sigma}_{32}\boldsymbol{\Sigma}^{21} & \boldsymbol{\Sigma}_{31}\boldsymbol{\Sigma}^{12} + \boldsymbol{\Sigma}_{32}\boldsymbol{\Sigma}^{22} \end{pmatrix} \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \end{pmatrix}, \end{aligned} \tag{A.3}$$

where $\boldsymbol{\epsilon}_{1,t}$, $\boldsymbol{\epsilon}_{2,t}$, and $\boldsymbol{\epsilon}_{3,t}$ are the orthogonalized counterparts of $\mathbf{u}_{1,t}$, $\mathbf{u}_{2,t}$, and $\mathbf{u}_{3,t}$, respectively.

Moreover, the block inverse matrix inside equation (A.3) is

$$\begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}^{-1} \equiv \begin{pmatrix} \Sigma^{11} & \Sigma^{12} \\ \Sigma^{21} & \Sigma^{22} \end{pmatrix} = \begin{pmatrix} \Sigma_{11}^{-1} + \Sigma_{11}^{-1}\Sigma_{12}K\Sigma_{21}\Sigma_{11}^{-1} & -\Sigma_{11}^{-1}\Sigma_{12}K \\ K\Sigma_{21}\Sigma_{11}^{-1} & K \end{pmatrix},$$

where

$$K = (\Sigma_{22} - \Sigma_{21}\Sigma_{11}^{-1}\Sigma_{12})^{-1}$$

is the Schur complement of Σ . Stacking equations (A.1)-(A.3) then yields

$$\begin{aligned} \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \epsilon_{3,t} \end{pmatrix} &= \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \\ \mathbf{u}_{3,t} \end{pmatrix} - \begin{pmatrix} 0 & 0 & 0 \\ \Sigma_{21}\Sigma_{11}^{-1} & 0 & 0 \\ \Sigma_{31}\Sigma^{11} + \Sigma_{32}\Sigma^{21} & \Sigma_{31}\Sigma^{12} + \Sigma_{32}\Sigma^{22} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \\ \mathbf{u}_{3,t} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{I} & 0 & 0 \\ -\Sigma_{21}\Sigma_{11}^{-1} & \mathbf{I} & 0 \\ -\Sigma_{31}\Sigma^{11} - \Sigma_{32}\Sigma^{21} & -\Sigma_{31}\Sigma^{12} - \Sigma_{32}\Sigma^{22} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{u}_{1,t} \\ \mathbf{u}_{2,t} \\ \mathbf{u}_{3,t} \end{pmatrix}, \end{aligned}$$

or

$$\epsilon_t = \mathbf{Q}_3^{-1}\mathbf{u}_t.$$

References

- Bonhomme, S., T. Lamadon, and E. Manresa (2022), “Discretizing Unobserved Heterogeneity,” *Econometrica*, 90, 625–643.
- Bonhomme, S. and E. Manresa (2015), “Grouped Patterns of Heterogeneity in Panel Data,” *Econometrica*, 83, 625–643.
- Brownlees, C., G.S. Gudmundsson, and G. Lugosi (2022), “Community Detection in Partial Correlation Network Models,” *Journal of Business and Economic Statistics*, 40, 216–226.
- Chiang, H.D., Y. Sasaki, and Y. Wang (2025), “Genuinely Robust Inference for Clustered Data,” Working Paper, <https://arxiv.org/abs/2308.10138>.
- Demirer, M., F.X. Diebold, L. Liu, and K. Yilmaz (2018), “Estimating Global Bank Network Connectedness,” *Journal of Applied Econometrics*, 33, 1–15.

- 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 (2012), “Better to Give Than to Receive: Predictive Directional Measurement of Volatility Spillovers,” *International Journal of Forecasting*, 28, 57–66.
- Diebold, F.X. and K. Yilmaz (2014), “On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms,” *Journal of Econometrics*, 182, 119–134.
- Diebold, F.X. and K. Yilmaz (2023), “On the Past, Present, and Future of the Diebold-Yilmaz Approach to Dynamic Network Connectedness,” *Journal of Econometrics*, (50th Anniversary Special Issue), 234, 115–120.
- Forbes, K.J. and R. Rigobon (2002), “No Contagion, Only Interdependence: Measuring Stock Market Comovements,” *Journal of Finance*, 57, 2223–2261.
- Goyal, A. and I. Welch (2008), “A Comprehensive Look at The Empirical Performance of Equity Premium Prediction,” *Review of Financial Studies*, 21, 1455–1508.
- Gudmundsson, G.S. and C. Brownlees (2021), “Detecting Groups in Large Vector Autoregressions,” *Journal of Econometrics*, 225, 2–26.
- Jacomy, M., S. Heymann, T. Venturini, and M. Bastian (2014), “ForceAtlas2, A Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software,” *PLoS ONE*, 9, www.plosone.org.
- Jorda, O. (2005), “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- Jorda, O. and A.M. Taylor (2024), “Local Projections,” Working Paper, <https://www.nber.org/papers/w32822>.
- Koop, G., M.H. Pesaran, and S.M. Potter (1996), “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74, 119–147.
- Li, D., B. Peng, S. Tang, and W. Wu (2025), “Estimation of Grouped Time-Varying Network Vector Autoregression Models,” Working paper, <https://fba.um.edu.mo/wp-content/uploads/RePEc/doc/202526.pdf>.

- MacKinnon, J.G., M.O. Nielsen, and M.D. Webb (2023), “Cluster-Robust Inference: A Guide to Empirical Practice,” *Journal of Econometrics*, 232, 272–299.
- Montiel Olea, J.L. and M. Plagborg-Moller (2021), “Local Projection Inference is Simpler and More Robust Than You Think,” *Econometrica*, 89, 1789–1823.
- Pesaran, M.H. and Y. Shin (1998), “Generalized Impulse Response Analysis in Linear Multivariate Models,” *Economics Letters*, 58, 17–29.
- Rambachan, A. and N. Shephard (2025), “When do Common Time Series Estimands Have Nonparametric Causal Meaning?” ArXiv Working Paper, <https://arxiv.org/abs/1903.01637>.
- Sims, C.A. (1980), “Macroeconomics and Reality,” *Econometrica*, 48, 1–48.
- Stock, J.H. and M.W. Watson (2005), “Implications of Dynamic Factor Models for VAR Analysis,” Working Paper, Harvard University and Princeton University, <https://www.princeton.edu/~mwatson/papers/favar.pdf>.
- Tibshirani, R. (1996), “Regression Shrinkage and Selection via the Lasso,” *Journal of the Royal Statistical Society, Series B (Methodological)*, 58, 267–288.
- Tong, C., P.R. Hansen, and I. Archakov (2024), “Cluster GARCH,” Working Paper, <https://arxiv.org/abs/2406.06860>.
- Zhang, B. (2024), “Incorporating Prior Knowledge of Latent Group Structure in Panel Data Models,” Working Paper, https://boyuan-zhang.github.io/assets/pdf/CBG_latest.pdf.
- Zhu, X., R. Pan, G. Li, Y. Liu, and H. Wang (2017), “Network Vector Autoregression,” *Annals of Statistics*, 45, 1096–1123.
- Zou, H. and H. Zhang (2009), “On the Adaptive Elastic Net with a Diverging Number of Parameters,” *Annals of Statistics*, 37, 1733.