European Journal of Operational Research 314 (2024) 393–407

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/eor

Interfaces with other disciplines

Persistence in financial connectedness and systemic risk ☆

Jozef Baruník [a], Michael Ellington [b,*]

[a] The Czech Academy of Sciences, IITA, Pod Vodárenskou Věží4, 182 08, Prague, Czech Republic
[b] University of Liverpool Management School, Chatham Building, Liverpool, L69 7ZH, UK

ARTICLE INFO

Keywords:
Finance
Network connections
Variance decompositions
Persistence
Spectral domain

ABSTRACT

This paper characterizes dynamic linkages arising from shocks with heterogeneous degrees of persistence. Using frequency domain techniques, we introduce measures that identify smoothly varying links of a transitory and persistent nature. Our approach allows us to test for statistical differences in such dynamic links. We document substantial differences in transitory and persistent linkages among US financial industry volatilities, argue that they track heterogeneously persistent sources of systemic risk, and thus may serve as a useful tool for market participants.

1. Introduction

Firms and economic units create connections through a variety of channels (see e.g. Garvey et al., 2015; Herskovic et al., 2020; Richmond, 2019).[1] These connections are dynamic and vary over time with changing states of an economy, as both stable and uncertain periods are associated with different shocks. At the same time, an increasing number of authors argue that economic variables are driven by shocks that influence their future value with heterogeneous levels of persistence (Bandi, Chaudhuri, Lo, & Tamoni, 2021; Dew-Becker & Giglio, 2016). Connectedness and the subsequent network structures that emerge from these relationships are central to risk measurement and management, as well as to understanding macroeconomic risks that emerge over business cycles.

The main objective of this paper is to introduce measures capable of identifying the smoothly varying persistence structure of such linkages. As an interesting and important example, we provide an analysis of US sectors, with a focus on the financial sector, and identify a heterogeneous persistence structure of systemic risk.

There are a variety of ways to measure connectedness.[2] In particular, Diebold and Yilmaz (2014) provides a unifying framework for measuring connectedness and other network properties using variance decompositions from an approximate model. Variance decompositions track how shocks affect the future variation of variables within a system, and are therefore a natural choice for inferring network connectedness from data.[3] However, the time evolution of such measures typically involves the estimation of static models that require covariance stationarity and roll through a time series of data. More importantly, such measures aggregate shocks and mask the persistence structures of a network.

In this paper, we provide the tools to identify dynamic connectedness and other key network measures from variance decompositions. We argue that different levels of connectedness can form around

☆ We thank Oliver Linton, Wolfgang Härdle, Melanie Schienle, Catherine Forbes, Lubos Hanus, Lukáš Vácha, Ryland Thomas, Chris Florackis, Costas Milas, Alex Kostakis and Charlie Cai for invaluable discussions and comments. We are grateful to Luboš Hanus for help in furnishing and converting estimation codes. We acknowledge insightful comments from numerous seminar presentations, such as: The 2019 and 2020 Society for Economic Measurement Conferences; the Danish National Bank; the 2019 STAT of ML conference; the 13th International Conference on Computational and Financial Econometrics; and many more. Jozef Baruník gratefully acknowledges support from the Czech Science Foundation under the 19-28231X (EXPRO) project. For estimation of dynamic horizon specific networks, we provide packages `DynamicNets.jl` in JULIA and `DynamicNets` in MATLAB. The packages are available at https://github.com/barunik/DynamicNets.jl and https://github.com/ellington/DynamicNets.

* Corresponding author.
E-mail addresses: barunik@fsv.cuni.cz (J. Baruník), m.ellington@liverpool.ac.uk (M. Ellington).

[1] Richmond (2019) measures connections through consumption growth. Garvey, Carnovale, and Yeniyurt (2015) tracks connections that describe the supply chain, and Herskovic, Kelly, Lustig, and Van Nieuwerburgh (2020) studies network connections between firm volatilities.

[2] Some propose measures derived from correlations or coefficient estimates (see e.g. Calabrese & Osmetti, 2019; Engle & Kelly, 2012; Geraci & Gnabo, 2018), while others track links between individual companies and broader market/economic movements (see e.g. Acharya, Engle, & Richardson, 2012; Acharya, Pedersen, Philippon, & Richardson, 2017; Adrian & Brunnermeier, 2016) .

[3] The subsequent literature has widely adopted the Diebold and Yilmaz (2014) approach to address a variety of issues in finance and economics (e.g. Baruník, Bevilacqua, & Tunaru, 2020; Yang & Zhou, 2017).

https://doi.org/10.1016/j.ejor.2023.11.023
Received 1 February 2023; Accepted 15 November 2023
Available online 17 November 2023
0377-2217/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

transitory and persistent components of shocks within a system of financial data. To capture the time-varying nature of such transitory and persistent components of connectedness, we further consider time-varying variance decomposition matrices from vector autoregressions (VARs) as dynamic adjacency matrices. To identify the persistence structure of the network, we propose to use localized spectral decompositions[4] of variance error forecasts.

The main contribution of our paper is to provide a novel framework for measuring dynamic relationships that relate to different horizons of interest in multivariate time series models. We use a locally stationary Bayesian time-varying parameter VAR model, which is readily available in high-dimensional settings. We also develop a test for differences in connectedness over different horizons and show how to infer differences over time. We provide Monte Carlo evidence that our measures are able to reliably track connections from different data generating processes (DGPs), including those that are non-Gaussian. Finally, we make computationally efficient packages `DynamicNets.jl` in JULIA and `DynamicNets` in MATLAB that allows one to obtain our measures on data the researcher desires.[5]

Our approach provides a solution to the problems of using rolling windows (see e.g. Demirer, Diebold, Liu, & Yilmaz, 2018) that does not suffer from dimensionality issues or inference problems. The Bayesian nature of our framework incorporates prior shrinkage and provides information about estimation uncertainty from the posterior distribution of the connectedness measures. This is in sharp contrast to conventional studies that only provide point estimates and rely on bootstrapping for confidence intervals. Our measures are also readily available for applications with large data systems. This extends the pairwise approach in Geraci and Gnabo (2018) to study linkages between firms.

The linkages that form over different horizons with heterogeneous persistence are important for a number of reasons. First, economic theory suggests that the marginal utility of agents' preferences depends on cyclical components of consumption (see e.g. Bandi & Tamoni, 2017; Giglio, Maggiori, & Stroebel, 2015) and also on investment horizons in their risk attitudes (Dew-Becker & Giglio, 2016). Such behaviour can be observed, for example, under myopic loss aversion, where an agent's decision depends on the valuation horizon.

Second, unanticipated shocks or news have the capacity to alter these preferences and can therefore generate transitory and persistent linkages of different strengths. For example, a shock that has an impact at longer horizons may reflect permanent changes in expectations of future price movements. Such a shock may lead to a permanent change in a firm's future dividend payments (Balke & Wohar, 2002). Conversely, a shock that affects shorter horizons may suggest temporary changes in future price movements. For example, suppose that the shock relates only to a change in an upcoming dividend payment. This would likely result in a very short term change, reflecting the transitory nature of the news.

Third, firms have different short-run and long-run objectives, and investors view short-run and long-run risks differently (see e.g. Drechsler & Yaron, 2011; Gerrard, Kyriakou, Nielsen, & Vodička, 2022). This behaviour motivates the long-run risk asset pricing literature pioneered by Bansal, Kiku, and Yaron (2010), Bansal and Yaron (2004). The implication here is that investment horizons may be a source of systematic risk that investors demand compensation for (contributions on this topic include e.g. Brennan & Zhang, 2018; Chaudhuri & Lo, 2019).

Identifying network structures that form due to idiosyncratic shocks is also relevant because they can determine aggregate fluctuations (Acemoglu, Carvalho, Ozdaglar, & Tahbaz-Salehi, 2012). These links between individual or firm-level entities create systemic risks for sectors and the economy as a whole (Acemoglu, Ozdaglar, & Tahbaz-Salehi, 2017; Billio, Getmansky, Lo, & Pelizzon, 2012). Such risks drive changes in uncertainty, which can be key factors in business cycles and financial distress (Bloom, Floetotto, Jaimovich, Saporta-Eksten, & Terry, 2018). Gabaix (2011) shows that sectoral co-movements are salient features of business cycles. Meanwhile, Atalay (2017) finds that the lion's share of variation in output growth is due to idiosyncratic industry-level shocks.

However, an understanding of the potential longevity of risks arising from these linkages and their importance in driving financial turmoil or business cycles is incomplete. We show how our approach provides measures of transitory and persistent linkages using the daily realized firm-level volatilities of S&P500 constituents. We classify the constituents into their eleven main sectors according to the Global Industry Classification Standard (GICS) and measure network connections from the transitory and persistent components of volatility shocks.

Our empirical results document substantial heterogeneities in transient and persistent measures of connectedness that reveal the nature of systemic risks arising from networks. Specifically, we document: (i) spikes in persistent network connectedness when long-lasting financial and economic events occur; and (ii) statistically significant differences between transitory and persistent network connectedness across sectors. Our network measures can serve as an online monitoring tool for sectoral uncertainty in markets of interest to macroprudential supervisors and investors alike.

The rest of the paper proceeds as follows. Section 2 derives our measures from locally stationary processes, discusses estimation, and proposes a test procedure for statistical differences in transitory and persistent connectedness. Section 3 provides Monte Carlo evidence that our measures are able to reliably track linkages and correctly identify statistical differences. In Section 4, we examine the links between firm-level volatilities of S&P500 sector constituents and assess the information content of sector connectedness measures beyond leading measures of uncertainty. Finally, Section 5 concludes.

2. Measuring transitory and persistent connections

Here we show how one can measure connectedness using the time-varying spectral decompositions. Our measures of connectedness are based on locally stationary processes. This assumes that the process is approximately stationary over a short time interval, which allows us to incorporate time variation into our analysis. This in turn allows us to construct our measures of frequency-dependent, time-varying network connectedness.

Formally, consider a doubly indexed N-variate time series $(\boldsymbol{X}_{t,T})_{1\le t\le T, T\in\mathbb{N}}$ with components $\boldsymbol{X}_{t,T} = (\boldsymbol{X}^1_{t,T}, \dots, \boldsymbol{X}^N_{t,T})^\top$ that describe all variables in an economy. Here t refers to a discrete time index and T is an additional index indicating the sharpness of the local approximation of the time series $(\boldsymbol{X}_{t,T})_{1\le t\le T, T\in\mathbb{N}}$ by a stationary one. Coarsely speaking, we can consider $(\boldsymbol{X}_{t,T})_{1\le t\le T, T\in\mathbb{N}}$ to be a weakly locally stationary process if, for a large T, given a set S_T of sample indices such that $t/T \approx u$ over $t \in S_T$, the sample $(\boldsymbol{X}_{t,T})_{t\in S_T}$ approximates the sample of a weakly stationary time series depending on the rescaled location u. Note that u is a continuous time parameter referred to as the rescaled time index, and T is interpreted as the number of available observation, hence $1 \le t \le T$ and $u \in [0, 1]$, see Dahlhaus (1996) for further details.

We assume that the economy follows a locally stationary TVP-VAR of lag order p as

$$\boldsymbol{X}_{t,T} = \boldsymbol{\Phi}_1(t/T)\boldsymbol{X}_{t-1,T} + \cdots + \boldsymbol{\Phi}_p(t/T)\boldsymbol{X}_{t-p,T} + \boldsymbol{\epsilon}_{t,T}, \tag{1}$$

[4] Note that frequency domain techniques are useful tools for denoising (Haven, Liu, & Shen, 2012; Sun & Meinl, 2012) and forecasting (Barunik, Krehlik, & Vacha, 2016; Sévi, 2014) financial time series.

[5] The packages are available at https://github.com/barunik/DynamicNets.jl and https://github.com/ellington/DynamicNets

where $\epsilon_{t,T} = \Sigma^{-1/2}(t/T)\eta_{t,T}$ with $\eta_{t,T} \sim NID(0, I_M)$ and $\Phi(t/T) = (\Phi_1(t/T), \dots, \Phi_p(t/T))^\top$ are the time varying autoregressive coefficients. In a neighbourhood of a fixed time point $u_0 = t_0/T$, we approximate the process $X_{t,T}$ by a stationary process $\widetilde{X}_t(u_0)$ as

$$\widetilde{X}_t(u_0) = \Phi_1(u_0)\widetilde{X}_{t-1}(u_0) + \dots + \Phi_p(u_0)\widetilde{X}_{t-p}(u_0) + \epsilon_t, \tag{2}$$

with $t \in \mathbb{Z}$ and under suitable regularity conditions $|X_{t,T} - \widetilde{X}_t(u_0)| = O_p\big(|t/T - u_0| + 1/T\big)$ which justifies the notation "locally stationary process". Crucially, the process has time varying VMA(∞) representation (Dahlhaus et al., 2009; Roueff & Sanchez-Perez, 2016)

$$X_{t,T} = \sum_{h=-\infty}^{\infty} \Psi_{t,T}(h)\epsilon_{t-h} \tag{3}$$

where $\Psi_{t,T}(h) \approx \Psi(t/T, h)$ is a stochastic process satisfying $\sup_\ell \|\Psi_t - \Psi_\ell\|^2 = O_p(h/t)$ for $1 \le h \le t$ as $t \to \infty$. Specifically, $\Psi_{t,T}(h) = \big[\Phi_{t,T}(h)\big]^{-1}$, which is key to understanding dynamics. Since $\Psi_{t,T}(h)$ contains an infinite number of lags, we approximate the moving average coefficients at $h = 1, \dots, H$ horizons (see the detailed discussion below). The network characteristics rely on variance decompositions, which are transformations of the impulse response functions, $\Psi_{t,T}(h)$, and permit the measurement of the contribution of shocks to the system.

Since a shock to a variable in the model does not necessarily appear alone, i.e. orthogonally to shocks to other variables, an identification scheme is crucial in calculating variance decompositions. We adapt the generalized identification scheme in Pesaran and Shin (1998) to locally stationary processes. A natural way to disentangle connections that form over transitory and persistent components of shocks is to consider a spectral representation of the approximating model.[6] Hence instead of impulse responses, we propose to use the (local) frequency response of a shock. The building block of our measures consider a time-varying frequency response function $\Psi_{t,T}(e^{-i\omega}) = \sum_h e^{-i\omega h}\Psi_{t,T}(h)$ which we obtain from a Fourier transform of the coefficients with $i = \sqrt{-1}$.

Before introducing our network measures, we define the time varying spectral density of $X_{t,T}$ at frequency ω which is locally the same as the spectral density of $\widetilde{X}_t(u)$ at $u = t/T$ as a Fourier transform of VMA(∞) filtered series as

$$S_X(u,\omega) = \sum_{h=-\infty}^{\infty} \mathbb{E}\Big[\widetilde{X}_{t+h}(u)\widetilde{X}_t^\top(u)\Big]e^{-i\omega h} = \Big\{\Psi(u, e^{-i\omega})\Big\}\Sigma(u)\Big\{\Psi(u, e^{+i\omega})\Big\}^\top. \tag{4}$$

The time-varying spectral density is a key quantity for understanding frequency dynamics. It describes the distribution of the time varying covariance of $X_{t,T}$ over frequency components ω. The local spectral density captures the influence of the time-varying parameters through the impulse transfer functions $\Psi(u, e^{-i\omega})$, $\Psi(u, e^{+i\omega})$ above. Using the spectral representation for the local covariance that is associated with the local spectral density,

$$\mathbb{E}\left[\widetilde{X}_{t+h}(u)\widetilde{X}_t^\top(u)\right] = \int_{-\pi}^{\pi} S_X(u,\omega)e^{i\omega h}d\omega \tag{5}$$

we can naturally introduce time-varying frequency domain counterparts of variance decompositions. This is important since, as Diebold and Yilmaz (2014) note, we can view the variance decomposition matrix as an adjacency matrix forming asymmetric connections among a system of variables. In our case, this allows us to define dynamic adjacency matrices with different degrees of persistence.

2.1. A route towards transitory and persistent connectedness

Variance decompositions are transformations of impulse responses $\Psi_{t,T}(h)$ that allow us to measure the contribution of shocks to the system and thus to characterize the networks that form in response to shocks. While it may seem natural to choose a forecast horizon $h = 1, \dots, H$ of interest, this choice is costly in terms of information aggregation and hence loss. In contrast to the cumulative information with increasing h, the spectral representation of the impulse responses $\Psi_{t,T}(e^{-i\omega})$ contains much richer and more precise information. Switching to the frequency domain allows one to trace network connections arising from transitory and persistent components of shocks.

To illustrate, consider a simple bivariate system in which the links between two variables are of interest. Specifically, we are interested in how variable b responds to two different shocks to variable a.[7] The first is a *transitory shock* that causes the variable b to rise by one unit in period 1 and fall by one unit in period 2, before returning to zero from period onwards. Here we expect the proportion of error variation to be large at short horizons and small at long horizons due to the purely transitory effect.

The second is a *persistent shock* that results in a unit increase in the variable b and a gradual decrease to zero over the impulse horizon. In this case, the fraction of error variation is expected to be large at longer horizons and smaller at shorter horizons. How these quantities in the time and frequency domain capture the responses to these two shocks is the main motivation for our measures.

Fig. 1 shows the impulse response functions of variable b to shocks to variable a in the left panel, the corresponding fraction of forecast error variances in the middle panel, and their corresponding spectral decompositions of forecast error variance fractions in the right panel. Note that one can hardly identify persistence of the shock from the almost indistinguishable forecast error variance shares (variance decompositions) in the time domain depicted by the middle panel of Fig. 1. A persistent shock results in a slightly larger value of forecast error variation relative to the variation due to a transitory shock. At the same time, if we were to estimate these quantities from the data and take into account the uncertainty of the estimates (the figure plots the theoretical values of a simple example), they would become statistically indistinguishable.

In contrast, the spectral representation of the forecast error variance shares in the right-hand panel of Fig. 1 accurately captures the heterogeneous impact of the two shocks across frequencies. The transitory shock in the variable a has a negligible impact at low frequencies (i.e. close to $\omega = 0$), indicating that this shock has no importance for the long-run variation of the variable b, and larger weights at higher frequencies, showing the transitory nature of the link established by this shock. Conversely, the persistent shock affects low frequencies and correctly identifies a persistent link between the variables.

The main implication is that we can construct network measures that take into account the nature of the shocks that form such links. Thus, using spectral decompositions, we are able to identify transitory and persistent network links that are not apparent in the time domain.

2.2. Measuring connectedness

The following proposition establishes the time-varying spectral representation of the variance decomposition of shocks from variable j to variable k. This is central to the existence of network measures in the time–frequency domain.[8]

[6] Baruník and Křehlík (2018) disentangle long-run and short-run unconditional network connections using standard VAR models.

[7] In this example, we assume that the impact of own shocks on variable b increases by 1 unit and is persistent. This means that the impact of a 1 unit shock to variable b affects the value of b up to 20 horizons after we observe the shock. Note that our analysis in Fig. 1 remains the same if the own shocks are transitory.

[8] Note to notation: $[\boldsymbol{A}]_{j,k}$ denotes the jth row and kth column of matrix $\boldsymbol{A}$ denoted in bold. $[\boldsymbol{A}]_{j,\cdot}$ denotes the full jth row; $[\boldsymbol{A}]_{\cdot,j}$ denotes the full jth column. A $\sum A$, where A is a matrix that denotes the sum of all elements of the matrix A.

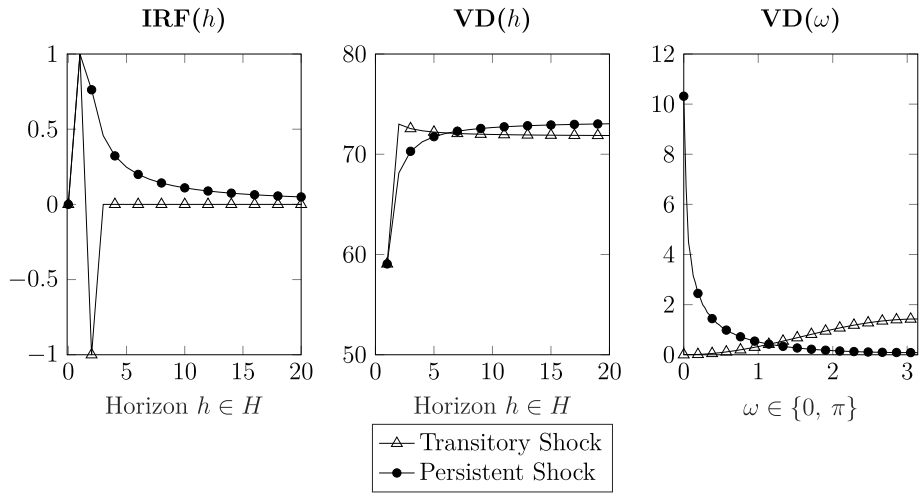

Fig. 1. Impulse Response Functions and Variance Decompositions in time and frequency domain.
Notes: This figure reports impulse response functions (left), variance decompositions in time domain (middle) and the corresponding spectral representation of the forecast error variance decompositions (right), of a variable b within a bivariate system with respect to a shock in the variable a. We consider two types of shocks, a transitory shock (triangles) and a persistent shock (dots) impacting horizon $h = 1, \dots, H$ and frequency $\omega \in \{0, \dots, \omega\}$.

Proposition 1 (*Dynamic Adjacency Matrix*). *Suppose $\boldsymbol{X}_{\mathbf{t,T}}$ is a weakly locally stationary process with $\sigma_{kk}^{-1} \sum_{h=0}^{\infty} \left| \left[\boldsymbol{\Psi}(u,h)\boldsymbol{\Sigma}(u) \right]_{j,k} \right| < +\infty, \forall j,k$. Then the **time–frequency variance decompositions** of the jth variable at a rescaled time $u = t_0/T$ due to shocks in the kth variable on the frequency band $d = (a,b) : a,b \in (-\pi,\pi), a < b$ form a **dynamic adjacency matrix** defined as*

$$\left[\theta(u,d)\right]_{j,k} = \frac{\sigma_{kk}^{-1} \int_a^b \left| \left[\boldsymbol{\Psi}(u,e^{-i\omega})\boldsymbol{\Sigma}(u) \right]_{j,k} \right|^2 d\omega}{\int_{-\pi}^{\pi} \left[\left\{ \boldsymbol{\Psi}(u,e^{-i\omega}) \right\} \boldsymbol{\Sigma}(u) \left\{ \boldsymbol{\Psi}(u,e^{+i\omega}) \right\}^{\top} \right]_{j,j} d\omega} \tag{6}$$

where $\boldsymbol{\Psi}(u,e^{-i\omega}) = \sum_h e^{-i\omega h}\boldsymbol{\Psi}(u,h)$ is local impulse transfer function or frequency response function computed as the Fourier transform of the local impulse response $\boldsymbol{\Psi}(u,h)$

Proof. See Appendix A. □

It is important to note that $\left[\theta(u,d)\right]_{j,k}$ is a natural disaggregation of traditional variance decompositions to time-varying frequency bands. This is because a portion of the local error variance of the jth variable at a given frequency band due to shocks in the kth variable is scaled by the variance of the jth variable. Note that while the Fourier transform of the impulse response generally takes on complex values, the quantity in proposition 1 is the squared modulus of weighted complex numbers, thus producing a real quantity.

This relationship is an identity which means the integral is a linear operator, summing over disjoint intervals covering the entire range $(-\pi,\pi)$ recovers the time domain counterpart of the local variance decomposition for $h \to \infty$. The following remark formalizes this fact.

Remark 1 (*Aggregation of Adjacency Matrix*). Denote by d_s an interval on the real line from the set of intervals D that form a partition of the interval $(-\pi,\pi)$, such that $\cap_{d_s \in D} d_s = \emptyset$, and $\cup_{d_s \in D} d_s = (-\pi,\pi)$. Due to the linearity of integral and the construction of d_s, we have

$$\left[\theta(u)\right]_{j,k} = \sum_{d_s \in D} \left[\theta(u,d_s)\right]_{j,k}.$$

Remark 1 is important as it establishes the aggregation of network connectedness measures across different frequency bands to its time domain, *total* counterpart. Hence one can easily obtain time varying network measures across any horizon of interest using frequency bands that will always sum up to an aggregate time domain counterpart.

As the rows of the time–frequency network connectedness do not necessarily sum to one, we normalize the element in each by the corresponding row sum

$$\left[\widetilde{\theta}(u,d)\right]_{j,k} = \left[\theta(u,d)\right]_{j,k} \Big/ \sum_{k=1}^{N} \left[\theta(u)\right]_{j,k} \tag{7}$$

Our notion that we can approximate well the process $\boldsymbol{X}_{t,T}$, by a stationary process $\widetilde{\boldsymbol{X}}_t(u)$ in a neighbourhood of a fixed time point $u = t/T$, means that all associated local quantities approximate well their time varying counterparts. Following the arguments in Dahlhaus (1996), and using mild assumptions, one can easily see that local variance decompositions at a frequency band $\widetilde{\theta}(u,d)$ approximate well the time-varying variance decompositions of the process $\boldsymbol{X}_{t,T}$.

Note that the local generalized variance decompositions form a dynamic adjacency matrix that defines a time-varying network at a given frequency band. Thus, we can directly use our measures as time-varying network characteristics that contain richer information in comparison to typical network analysis. In our notion, variance decompositions can be viewed as weighted links showing the strengths of connections. In addition, the links are directional, meaning that the j to k link is not necessarily the same as the k to j link, and hence the adjacency matrix is asymmetric. Even more important, the adjacency matrix is time-varying and frequency specific that allows the study of time-varying network characteristics at various frequency bands of the user's choice. The simplest is to measure transitory network connections over the short-run and persistent ones over the long-run.

Now we can define network connectedness measures that characterize a time-varying and frequency specific network. We define local network connectedness measures at a given frequency band as the ratio of the off-diagonal elements to the sum of the entire matrix

$$C(u,d) = 100 \times \sum_{\substack{j,k=1 \\ j \neq k}}^{N} \left[\widetilde{\theta}(u,d)\right]_{j,k} \Big/ \sum_{j,k=1}^{N} \left[\widetilde{\theta}(u)\right]_{j,k} \tag{8}$$

This measures the contribution of forecast error variance attributable to all shocks in the system, minus the contribution of own shocks over frequency band d and infers system-wide connectedness over such frequency band. We can also define measures that reveal when an individual variable in the economy is a transmitter or receiver of shocks. Local directional connectedness measures how much of each variables's j variance is due to shocks in other variables $k \neq j$ in the

economy over frequency band d is given by

$$C_{j\leftarrow\bullet}(u,d) = 100 \times \sum_{\substack{k=1\\k\neq j}}^{N} \left[\widetilde{\theta}(u,d)\right]_{j,k} \Big/ \sum_{j,k=1}^{N} \left[\widetilde{\theta}(u)\right]_{j,k}, \tag{9}$$

defining the so-called FROM connectedness. One can precisely interpret this quantity as from-degrees (often called out-degrees in the network literature) associated with the nodes of the weighted directed network represented by the variance decompositions matrix generalized to a time-varying frequency specific quantity. Likewise, the contribution of variable j to variances in other variables is computed as

$$C_{j\rightarrow\bullet}(u,d) = 100 \times \sum_{\substack{k=1\\k\neq j}}^{N} \left[\widetilde{\theta}(u,d)\right]_{k,j} \Big/ \sum_{k,j=1}^{N} \left[\widetilde{\theta}(u)\right]_{k,j} \tag{10}$$

and is the so-called TO connectedness. Again, one can precisely interpret this as to-degrees (often called in-degrees in the network literature) associated with the nodes of the weighted directed network represented by the variance decompositions matrix. These two measures show how other variables contribute to the variation of variable j, and how variable j contributes to the variation of others, respectively, in a time varying fashion at a chosen frequency band. We note here that taking the difference between TO connectedness and FROM connectedness summarizes information regarding directional connections in net-terms. Further, one can track pairwise connections over frequency bands in an analogous manner to the above as differences between the $j-k$th element and $k-j$th elements.

Importantly, the following proposition shows one can always reconstruct time domain network connectedness measures from our frequency-dependent networks.

Proposition 2 (*Reconstruction of Dynamic Network Connectedness*). *Denote by d_s an interval on the real line from the set of intervals D that form a partition of the interval $(-\pi,\pi)$, such that $\cap_{d_s\in D} d_s = \emptyset$, and $\cup_{d_s\in D} d_s = (-\pi,\pi)$. We then have that*

$$\begin{aligned} C(u) &= \sum_{d_s\in D} C(u,d_s) \\ C_{j\leftarrow\bullet}(u) &= \sum_{d_s\in D} C_{j\leftarrow\bullet}(u,d_s) \\ C_{j\rightarrow\bullet}(u) &= \sum_{d_s\in D} C_{j\rightarrow\bullet}(u,d_s) \end{aligned} \tag{11}$$

where $C(u)$ are local network connectedness measures aggregated over frequencies.

Proof. See Appendix A. □

In light of the above, all local frequency connectedness measures $C(u,d)$ for $u = t/T$ approximate well the time-varying frequency connectedness of the process $\boldsymbol{X}_{t,T}$.

2.3. Obtaining transitory and persistent network measures

In light of the assumptions that underpin our measures, we conjecture that the economy follows a stable time-varying parameter heteroskedastic VAR (TVP-VAR) model as in (1). We follow Petrova (2019) who establishes a Quasi Bayesian Local-Likelihood approach for inference in the presence of time-varying parameters.

For consistent estimation under the QBLL approach, let $\boldsymbol{X}_{t,T}$ be a time-series we observe with log probability density $l_t\left(\boldsymbol{X}_{t,T}|\boldsymbol{X}_{t-1,T}, \tilde{\boldsymbol{\Phi}}(t/T)\right)$, $\tilde{\boldsymbol{\Phi}}(t/T)$ stacks the time-varying autoregressive coefficient matrices into a finite-dimensional vector that satisfies one of the following conditions.

(i) $\tilde{\boldsymbol{\Phi}}_t = \tilde{\boldsymbol{\Phi}}(t/T)$ is a deterministic process where $\boldsymbol{\Phi}(.)$ is a piecewise differentiable function.

(ii) $\tilde{\boldsymbol{\Phi}}(t/T)$ is a stochastic process satisfying: $\sup_{j:|j-t|\leq h} \|\tilde{\boldsymbol{\Phi}}_t - \tilde{\boldsymbol{\Phi}}_j\|^2 = O_p(h/t)$ for $1 \leq h \leq t$, $t \to \infty$.

Both of the above indicate that the parameter sequence drafts gradually over time. The first condition is standard of Dahlhaus (2000) for locally stationary processes which requires the parameter process is a piecewise smooth deterministic function; thus allowing for breaks in parameters. The second condition is a generalization of the first to include stochastic parameter processes exhibiting degrees of persistence necessary for consistent estimation of stochastic driven time-variation. Such condition includes bounded random walk processes and some fractionally integrated processes. The parameters may feature any combination of deterministic trends and/or breaks satisfying conditions (i) and (ii) above. Our data generating processes (DGPs) in Section 3 are examples of such DGPs.[9]

To obtain the time-varying coefficient estimates at a fixed time point $u = t_0/T$, $\hat{\boldsymbol{\Phi}}_1(u),\ldots,\hat{\boldsymbol{\Phi}}_p(u)$, and the time-varying covariance matrices, $\hat{\boldsymbol{\Sigma}}(u)$, we follow the QBLL approach of Petrova (2019). Specifically, this approach uses a kernel weighting function that provides larger weights to observations that surround the period whose coefficient and covariance matrices are of interest. Using conjugate priors, the (quasi) posterior distribution of the parameters of the model are analytical. This alleviates the need to use a Markov Chain Monte Carlo (MCMC) simulation algorithm and permits the use of parallel computing. Note also that in using (quasi) Bayesian estimation methods, we obtain a distribution of parameters that we use to construct network measures that provide confidence bands. Details of the model and estimation algorithm are in Appendix B.[10] We provide a computationally efficient package `DynamicNets.jl` in JULIA and `DynamicNets` in MATLAB that allows one to obtain our measures on data the researcher desires.[11]

To estimate the elements of dynamic adjacency matrix, we first need to truncate the infinite VMA(∞) representation of the approximating model with a choice of finite horizon H. Here we note that in the frequency domain quantities, H serves only as an approximation factor, and it has no interpretation as in the time domain. Hence in the applications we advise setting the H sufficiently high to obtain a better approximation, particularly when lower frequencies are of interest. We obtain horizon specific measures using Fourier transforms and set our truncation horizon $H = 100$. Note, we run all results in this paper for $H \in \{50, 100, 200\}$, they are qualitatively similar and available upon request.

Next, estimating dynamic network measures requires the user to choose a kernel and its bandwidth. Typically the larger the bandwidth, the smoother time-evolution of our frequency specific network measures. Therefore, prior to tracking dynamic network connections, it is important the user considers the time-series properties of their data. For example if common peaks (troughs) in the time-series occur frequently and are transient, then a shorter bandwidth may be necessary. Conversely, if tracking network connections among data that evolves gradually over time, like interest rates, a larger bandwidth may be more appropriate. In the context of our study, we use a Normal kernel and explore the implication of bandwidth choice for a variety of data generating processes (DGPs) in Section 3.

[9] Figure 1 of Petrova (2019) provides figures of examples and provides further discussion around conditions the parameter sequences must satisfy for consistent estimation.

[10] Unlike traditional TVP VARs time-variation evolves in a non-parametric manner thus making no assumption on the laws of motion within the model. Typically, the model of Primiceri (2005), and many extensions, assume parameters evolve as random walks or autoregressive processes.

[11] The packages are available at https://github.com/barunik/DynamicNets.jl and https://github.com/ellington/DynamicNets.

It is noteworthy to mention that the choice of a two-sided kernel can come at a cost; especially if one wishes to use network measures for forecasting purposes. In these cases one may wish to: (i) estimate the dynamic network recursively throughout time such that the Normal kernel truncates to use only past values at the time T estimate; or (ii) choose a one-sided kernel such as those in Barigozzi, Hallin, Soccorsi, and von Sachs (2020), Hahn, Todd, and Van der Klaauw (2001). In practice, we encourage researchers to experiment with a variety of bandwidths to ensure results are not driven by its selection. We also encourage authors to use reasonable bandwidths given the data for their application. For example, if one was using these measures for forecasting daily stock return volatilities, a researcher might consider combination forecasts using Bayesian Model averaging to trade on. Alternatively one might look to use variance minimizing kernels in an attempt to reduce uncertainty around the forecast.[12]

We estimate the j,k element of our dynamic adjacency matrix at time $u = t_0/T$ and horizon $d = (a,b) : a,b \in (-\pi,\pi)$ and $a < b$ such that it corresponds to the transitory (high frequency band) and persistent (low frequency band) element of the adjacency matrix respectively as:

$$\left[\widehat{\boldsymbol{\theta}}(u,d)\right]_{j,k} = \frac{\widehat{\sigma}_{kk}^{-1} \sum_{\omega\in d} \left(\left[\widehat{\boldsymbol{\Psi}}(u,\omega)\widehat{\boldsymbol{\Sigma}}(u)\right]_{j,k} \right)^2}{\sum_{\omega\in(-\pi,\pi)} \left[\widehat{\boldsymbol{\Psi}}(u,\omega)\widehat{\boldsymbol{\Sigma}}(u)\widehat{\boldsymbol{\Psi}}^{\top}(u,\omega)\right]_{j,j}}, \tag{12}$$

where $\widehat{\boldsymbol{\Psi}}(u,\omega) = \sum_{h=0}^{H-1}\sum_h \widehat{\boldsymbol{\Psi}}(u,h)e^{-i\omega h}$ is an estimate of the impulse transfer function from Fourier frequencies $\omega \in \{aH/2\pi,\ldots,bH/2\pi\}$ of impulse response functions that cover a specific horizon.[13] From this, estimates of Eqs. (8)–(10) directly follow. For example if the application uses daily data, one may define transitory (short-term) as horizons corresponding to 1–5 days and persistent (long-term) as horizons corresponding to horizons greater than 5 days. This would require defining the band as $(a,b) = (2\pi/5, 2\pi)$ for the transitory and $(a,b) = (0, 2\pi/5)$ for the persistent networks.

2.4. Testing for statistical differences in connectedness

We now consider how one can determine, from a statistical perspective, differences between connectedness. We discuss in detail here how one can test for differences in connectedness one computes over different frequency bands.

In a Bayesian setting there are three alternatives for hypothesis testing. The first is the Bayes factor, the second uses posterior credible intervals, and the third follows statistical decision theory. We follow the latter and utilize the work of Li, Liu, and Yu (2015), Li, Zeng, and Yu (2014) and Liu, Li, Yu, and Zeng (2022).[14] These studies focus on developing test statistics of a point null hypothesis using the posterior distribution of parameters from a Bayesian model. The approach requires only the posterior distribution of parameters and has various advantages. First, they overcome the problem of the Jeffreys-Lindley paradox. Second, are not sensitive to the prior and are pivotal quantities. Third, they are easy to compute. Crucially, these statistics directly come from quadratic loss functions, as with classical test statistics, and therefore possess the same distributions as their frequentist counterparts.

Li et al. (2015) develop a Bayesian Lagrange-Multiplier (LM) type test that is asymptotically equivalent to a classical LM test. Using similar assumptions, Liu et al. (2022) develop a Bayesian Wald type test that is asymptotically equivalent to a classical Wald test and requires only the posterior mean and posterior variance of parameters under the null hypothesis.[15] Noting that the VMA(∞) representation of the VAR are nothing more than a transformation of the VAR parameters, we follow the assumptions in Liu et al. (2022) and therefore establish a Bayesian Wald type test for differences between network connectedness across frequency bands. We emphasize our network connectedness measures are manipulations of the VAR parameters themselves and possess a posterior distribution.[16]

We test the null hypothesis, $\mathrm{H}_0 : C(u,d_a) = C(u,d_b)$ that network connectedness across frequency band d_a and d_b are equivalent against the alternative, $\mathrm{H}_1 : C(u,d_a) \neq C(u,d_b)$. This is equivalent to testing $\mathrm{H}_0 : C(u,d_a) - C(u,d_b) = 0$ against $\mathrm{H}_1 : C(u,d_a) - C(u,d_b) \neq 0$. The following proposition establishes the test statistic and its asymptotic distribution under the null. We outline the regularity conditions in Appendix A along with a brief discussion on the importance of such conditions.

Proposition 3 (*Testing for Heterogeneities in Network Connectedness*). *Let $\bar{\mathcal{D}}(u)$ and $\mathbf{V}_D\left(\bar{\mathcal{D}}(u)\right)$ denote the time u posterior mean and variance of the difference between network connectedness across frequency band d_a and d_b. Then the test statistic under H_0*

$$\mathbf{W}\left(\boldsymbol{X}, \mathcal{D}_0(u)\right) = q_D + \left(\bar{\mathcal{D}}(u) - \mathcal{D}_0(u)\right)^{\top}\left[\mathbf{V}_{DD}\left(\bar{\mathcal{D}}(u)\right)\right]^{-1}\left(\bar{\mathcal{D}}(u) - \mathcal{D}_0(u)\right) \tag{13}$$

$$= q_D + \mathbf{Wald} \tag{14}$$

where $\mathbf{Wald} = \left(\bar{\mathcal{D}}(u) - \mathcal{D}_0(u)\right)^{\top}\left[\mathbf{V}_{DD}\left(\bar{\mathcal{D}}(u)\right)\right]^{-1}\left(\bar{\mathcal{D}}(u) - \mathcal{D}_0(u)\right)$, *$q_D$ is the number of restrictions, and $\mathcal{D}_0(u) = C(u,d_a) - C(u,d_b) = 0$.*

$$\mathbf{W}\left(\boldsymbol{X}, \mathcal{D}_0(u)\right) - q_D = \mathbf{Wald} + O_p(1) \to^d \chi^2\left(q_D\right)$$

Proof. See Appendix A. □

It is important to note that it is straightforward to generalize Eq. (13) to include multiple restrictions. This may be applicable if one requires testing equivalence among more than two frequency bands. We also note that one may also wish to utilize the above to test for differences between directional network connections over frequency bands. In the context of the above one would take the difference between net-directional connections, or pairwise directional connections, over frequency bands and compute the test in an analogous manner to below.

For estimation purposes, the test statistic only requires the posterior mean and the posterior variance of $\mathcal{D}(u)$, $\bar{\mathcal{D}}(u)$ and $\mathbf{V}\left(\bar{\mathcal{D}}\right)$. Let $\left\{\mathcal{D}^{[r]}\right\}_{r=1}^{R}$ denote the posterior draws such that $\mathcal{D}^{[r]} = \widehat{\mathcal{C}}^{[r]}(u,d_a) - \widehat{\mathcal{C}}^{[r]}(u,d_b)$ is the rth posterior draw of the difference between estimates of network connectedness over frequency band d_a and d_b. Then, the estimate of our test statistic for heterogeneities between network connectedness is

[12] In the context of our empirical application below where we look at daily realized volatilities of stock returns, we use a bandwidth equal to 8. We also estimate the models using bandwidths of 12, 18, and $\sqrt{T} = \sqrt{3278} \approx 57$. Increasing the bandwidth smooths our network connectedness measures because it assigns larger weights to more distant observations.

[13] Note that $i = \sqrt{-1}$.

[14] Bayes factors involve comparing the marginal likelihoods of two competing models and extensively appear in the literature (see e.g. Chan, 2020; Koop, Leon-Gonzalez, & Strachan, 2010). This is not appropriate in our setting because the network connectedness measures come from a manipulation of a sequence of posterior parameters from the same model; we have no alternative model to specify the marginal likelihood. Using posterior credible intervals is possible and something we consider in the spirit of Cogley, Primiceri, and Sargent (2010). In particular we use the joint posterior distribution of network connectedness measures to compute the probability that connectedness across one frequency band is larger than an analogous measure across another frequency band. These results are in Appendix C.

[15] They also show asymptotic equivalence between their Wald-type test the LM type test in Li et al. (2015).

[16] Lütkepohl (1990) provide the asymptotic distribution for impulse response functions from conventional VAR models one estimates using OLS Petrova (2019).

given by

$$\widehat{\mathbf{W}}(\boldsymbol{X}, \mathcal{D}(u)=0) = \frac{\frac{1}{R}\sum_{r=1}^{R}\left(\mathcal{D}^{[r]}(u)\right)^2}{\frac{1}{R}\sum_{r=1}^{R}\left(\mathcal{D}^{[r]}(u)-\bar{\bar{\mathcal{D}}}(u)\right)^2}, \text{ with } \bar{\bar{\mathcal{D}}}(u) = \frac{1}{R}\sum_{r=1}^{R}\mathcal{D}^{[r]}(u) \tag{15}$$

Under the null, we have $\mathbf{W}\left(\boldsymbol{X}, \mathcal{D}_0(u)\right) - q_D \to^d \chi^2\left(q_D\right)$ with $q_D = 1$ in this particular case.[17] Thus, we only need to compare $\widehat{\mathbf{W}}(\boldsymbol{X}, \mathcal{D}(u)=0)$ to the critical values of the $\chi^2(1)$ distribution. Rejecting the null implies that the time u network connectedness over frequency band d_a is statistically different to the corresponding time u network connectedness over frequency band d_b.

It may also be pertinent to test for differences in connectedness over time. Here we fix the frequency band, and now consider differences over time. In this case we test the null hypothesis, $\mathrm{H}_0 : C(u_1, d) = C(u_2, d)$ that network connectedness at time u_1 and u_2 across frequency band d are equivalent against the alternative, $\mathrm{H}_1 : C(u_1, d) \neq C(u_2, d)$. This is equivalent to testing $\mathrm{H}_0 : C(u_1, d) - C(u_2, d) = 0$ against $\mathrm{H}_1 : C(u_1, d) - C(u_2, d) \neq 0$. Now letting $\bar{\mathcal{D}}(s)$ and $\mathbf{V}_{DD}\left(\bar{\mathcal{D}}(s)\right)$ denote the posterior mean and variance of the difference between network connectedness at times u_1 and u_2 and replacing $\bar{\mathcal{D}}(u)$ and $\mathbf{V}_{DD}\left(\bar{\mathcal{D}}(u)\right)$ in Proposition 3 with these quantities delivers a Wald-type test statistic following a $\chi^2(1)$ distribution. Again, all we need to do is compare the test statistic with critical values of the $\chi^2(1)$ distribution.

3. Monte Carlo study

In this section, we conduct a Monte Carlo exercise to understand the finite sample properties of our connectedness measures. In order to motivate the need to focus on connections forming conditional on the persistence of shocks, we generate data with different levels of persistence throughout time and also changes in the covariance structure. This will induce differences in network connectedness measures we compute over different frequency bands. Here we concentrate on low and high frequency bands and consider four different data generating processes (DGP) to highlight their uses.

For simplicity, we focus on bivariate VAR(2) models with time-varying parameters and time-varying covariance matrices:

$$\begin{aligned}\boldsymbol{X}_{t,T} &= \boldsymbol{\Phi}_0(u) + \boldsymbol{\Phi}_1(u)\boldsymbol{X}_{t-1,T} + \boldsymbol{\Phi}_2(u)\boldsymbol{X}_{t-2,T} + \boldsymbol{\epsilon}_{t,T},\\ \boldsymbol{\epsilon}_{t,T} &= \boldsymbol{\Sigma}^{-1/2}(u)\boldsymbol{\eta}_{t,T},\ \boldsymbol{\eta}_{t,T} \backsim \left(0, \mathbf{I}_2\right)\end{aligned}$$

where $\boldsymbol{\Phi}_0(u)$ contains the time-varying intercepts and $\boldsymbol{\Phi}_1(u)$ and $\boldsymbol{\Phi}_2(u)$ contain the time-varying autoregressive parameters. The time-varying covariance matrix $\boldsymbol{\Sigma}(u) = \mathbf{A}^{-1}(u)\mathbf{H}(u)\left(\mathbf{A}^{-1}(u)\right)^{\top}$ with $\mathbf{A}^{-1}(u)$ being a lower triangular matrix with a unit diagonal and $\mathbf{H}(u)$ is a 2×2 diagonal matrix.

DGPI: Our first DGP has residuals such that, $\boldsymbol{\eta}_{t,T} \backsim \mathrm{NID}\left(0, \mathbf{I}_2\right)$. The time-varying intercepts follow the process:

$$\left[\boldsymbol{\Phi}_0(u)\right]_j = 0.0025\sin(0.004\pi t) + 0.15\sum_{i=1}^{t}\frac{v_i}{\sqrt{t}},\ v_i \backsim \mathrm{NID}\left(0, 0.001^2\right),\ j = 1, 2$$

For the time-varying autoregressive parameters, we have

$$\left[\boldsymbol{\Phi}_g(u)\right]_{j,k} = \begin{cases} 0.05\sin(0.002\pi t) + 0.75\sum_{i=1}^{t}\frac{\kappa_i}{\sqrt{t}}, & t \in \{1, \dots, 500\}\ \forall g, j, k = 1, 2\\ 0.45\sin(0.002\pi t) + 0.75\sum_{i=1}^{t}\frac{\kappa_i}{\sqrt{t}}, & t \in \{501, \dots, 1000\}\ j = k = 1,\ j = k = 2\\ 0.05\sin(0.002\pi t) + 0.75\sum_{i=1}^{t}\frac{\kappa_i}{\sqrt{t}}, & t \in \{501, \dots, 1000\}\ j = 1, k = 2\ \&\ j = 2, k = 1\end{cases}$$

with $\kappa_i \backsim \mathrm{NID}\left(0, 0.0001^2\right)$. The (2,1) element of $\mathbf{A}(u)$ have the following dynamics:

$$[\mathbf{A}(u)]_{2,1} = \begin{cases} 0.03\sin(0.002\pi t) + 0.7\sum_{i=1}^{t}\frac{v_i}{\sqrt{t}}, & t \in \{1, \dots, 500\}\\ 1.5\sin(0.002\pi t) + 0.7\sum_{i=1}^{t}\frac{v_i}{\sqrt{t}}, & t \in \{501, \dots, 1000\}\end{cases}$$

with $v_i \backsim \mathrm{NID}\left(0, 0.3^2\right)$. The diagonal elements of $\mathbf{H}(u)$ follow

$$\log\left[\mathbf{H}(u)\right]_{j,j} = \mu_j + \lambda_j\left(\log\left[\mathbf{H}(u-1)\right]_{j,j} - \mu_j\right) + \xi_{j,t}$$

where $\xi_{j,t} \backsim \mathrm{NID}\left(\mu_j, 0.1^2/(1-\lambda_j)\right)$, $\mu_j = 0.01$, $\lambda_j = 0.95$.

This DGP has little to no dependence for the first 500 observations which means connectedness at both high and low frequency bands will be low and close to zero. The latter half of the sample sees the AR coefficients in each equation become persistent as the sin wave becomes negative. Note also that the contemporaneous relationship intensifies. This induces connections at high frequency bands while connections at low frequency bands should be low and close to zero. Note we also allow for non-Gaussian residuals in this DGP such that we draw $\boldsymbol{\eta}_{t,T}$ from a multivariate student-t distribution with 5 degrees of freedom. These results are in Appendix C.

DGPII: For our second DGP, the time-varying intercepts follow the process:

$$\left[\boldsymbol{\Phi}_0(u)\right]_j = 0.25\sin(0.004\pi t) + 0.15\sum_{i=1}^{t}\frac{v_i}{\sqrt{t}},\ v_i \backsim \mathrm{NID}\left(0, 0.1^2\right),\ j = 1, 2$$

and the time-varying autoregressive parameters follow:

$$\left[\boldsymbol{\Phi}_g(u)\right]_{j,k} = 0.25\sin(0.004\pi t) + 0.75\sum_{i=1}^{t}\frac{\kappa_i}{\sqrt{t}}, \forall g, j, k = 1, 2$$

with $\kappa_i \backsim \mathrm{NID}\left(0, 0.3^2\right)$. The (2,1) element of $\mathbf{A}(u)$ and the diagonal elements of $\mathbf{H}(u)$ follow the processes:

$$\begin{aligned}[\mathbf{A}(u)]_{2,1} &= 0.3\sin(0.008\pi t) + 0.7\sum_{i=1}^{t}\frac{v_i}{\sqrt{t}}\\ \log\left[\mathbf{H}(u)\right]_{j,j} &= \mu_j + \lambda_j\left(\log\left[\mathbf{H}(u-1)\right]_{j,j} - \mu_j\right) + \xi_{j,t},\end{aligned}$$

where $v_i \backsim \mathrm{NID}\left(0, 0.3^2\right)$ and $\xi_{j,t} \backsim \mathrm{NID}\left(\mu_j, 0.1^2/(1-\lambda_j)\right)$, $\mu_j = 0.01$, $\lambda_j = 0.95$.

This DGP induces two distinct periods of persistence during observations 100-200 and 600-700 which amplifies connectedness at the low frequency band.

DGPIII: Our third DGP is the same as DGPII but relaxes the assumption that $\boldsymbol{\eta}_{t,T}$ are Gaussian. Instead we assume that the residuals follow a multivariate student-t distribution with 5 degrees of freedom.

DGPIV: Our fourth DGP is the same as DGPII, but increases the periodicity of the sin functions in the time-varying autoregressive matrices from $\sin(0.004\pi t)$ to $\sin(0.006\pi t)$. This generates three distinct periods of persistence during observations 50-150, 300-450, and 700-850.

For each of the four DGPs, we generate 100 simulations of length $T = 1000$ and compute the network connectedness measures. We use the median over these simulations as the true network connectedness. Then, for each of the 100 simulations of DGPI–DGPIV, we fit the TVP-VAR model we outline in Section 2.3. In fitting this model we take 1000 draws from the posterior distribution, calculate our network connectedness measures and then save the posterior median. For this exercise, we compute network connectedness on two frequency bands

[17] In the classical setting one uses Wald tests to check a variety of restrictions, such as a parameter of interest being equal to zero, or equivalence between two parameters of interest under the null hypothesis. When one tests the latter, the resulting Wald test is $\backsim \chi^2(1)$.

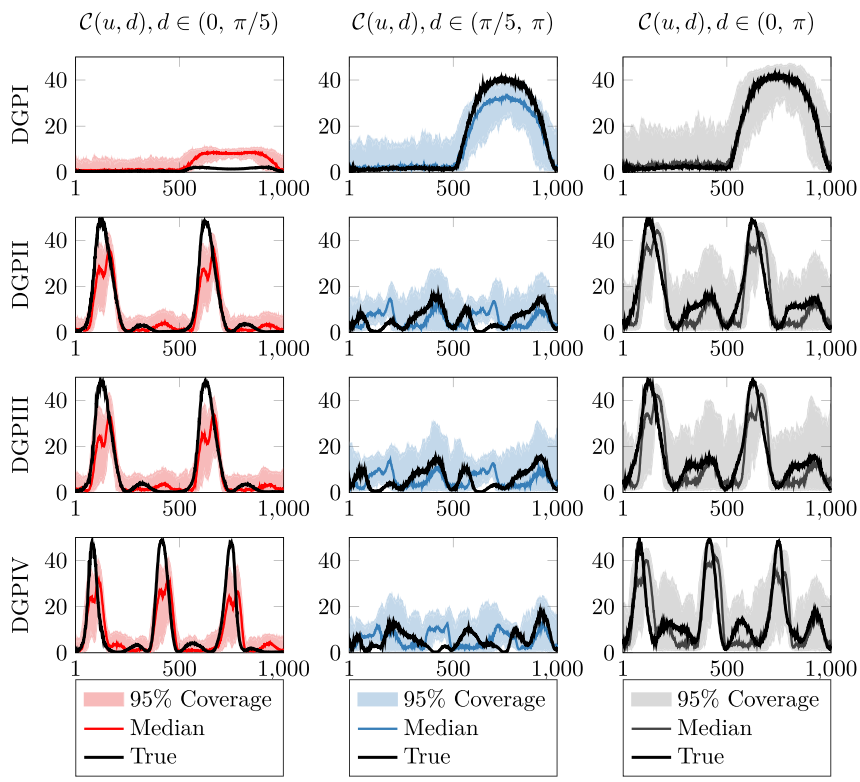

Fig. 2. Dynamic network connectedness measures: True and fitted values.
Notes: This figure plots the true network connectedness measures for three data generating processes following bi-variate TVP VAR(2) models along with the median and 95% quantiles of estimated network connectedness measures with bandwidth $W = 8$. The left columns report network connectedness on the low-frequency band, $d \in (0, \pi/5)$, the middle columns show network connectedness on the high-frequency band, $d \in (\pi/5, \pi)$, and the right columns show the aggregate network connectedness such that $d \in (0, \pi)$. DGPI (top row) is a TVP VAR(2) model with Gaussian errors, we introduce a break in the time-varying autoregressive matrices and contemporaneous relations from observation 500 that induces large connections across the high frequency band. DGPII (second row) is a TVP VAR(2) where time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIII (third row) is a TVP VAR(2) model with student-t errors, time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIV (bottom row) is the same as DGPII, but with an increase in the periodicity of the respective sin wave functions the time-varying intercepts and autoregressive matrices follow.

that cover the spectrum. The low-frequency band, which empirically pertains to persistent network connections, is $d \in (0, \pi/5)$, and the high-frequency band, pertaining to transitory network connections, is $d \in (\pi/5, \pi)$. For completeness, we compute the aggregate connectedness measures that considers the entire spectrum such that $d \in (0, \pi)$; this corresponds to a dynamic version of the Diebold and Yilmaz (2014) connectedness measure.

A final noteworthy point is the choice of bandwidth, W, for the kernel weights that induce time-variation into the TVP-VAR model. In this exercise, we consider $W = \{8, 12, 18\}$. However, for ease of exposition, we only report plots of the network connectedness measure estimates using $W = 8$ in the main text, results of fitted values from $W = \{12, 18\}$ are in Appendix C. The larger the bandwidth, the smoother the network measures become. This is because larger bandwidths assign weights to a higher number of observations around the one of interest. In general, we find that larger W results in poorer fit. This highlights the importance of selecting an appropriate bandwidth for the kernel weights relative to the data application.

As we discuss in 2.3, from a practical perspective, we encourage researchers to explore the robustness of their results to different bandwidths. For example, consider a low frequency forecaster looking to predict returns one-month ahead today, using our network connectedness measures, would likely place zero weight on data from the burst of the dot-com bubble and 2008 recession. Likewise, a high frequency investor would likely place little to no weight on data from one-year prior. As we show in Appendix C, lengthening the bandwidth in our Monte Carlo experiment causes the surges in estimates of connectedness to be more gradual.

If the data is low frequency data such as monthly yields, then one could argue to use a wider bandwidth as changes in these data are far smoother than returns or return volatility. If one is looking to describe the nature of connections that form on persistent and transitory components of shocks then we encourage researchers to use multiple bandwidths to check how these dynamics are influenced by such changes. We do not suspect these changes would drastically change the conclusions or results in most applications.

In Appendix C, we conduct further robustness checks for our simulation analysis. In Appendix C.3 we provide analysis on the performance of rolling VAR models for our DGPs. These results show that such connectedness estimates less accurate relative to our approach. In particular, estimates are highly sensitive to the window size and fail to accurately capture the peaks and troughs in connectedness; the latter is prominent as the complexity of the DGP increases (e.g. DGPII-DGPIV). Meanwhile in Appendix C.4 we conduct a Monte Carlo study on larger scale VAR models under DGPI containing $N = 10$ and $N = 25$ variables respectively. These results show that our approach is robust to increasing the number of variables and tracks connections well within larger-scale models.

Fig. 2 reports the true network connectedness measures and the median and 95% quantiles of corresponding estimates from the TVP VAR model using a kernel bandwidth of $W = 8$. We report network connectedness over the low-frequency-band, the high-frequency-band, and aggregate, in the left, middle, and right columns respectively.

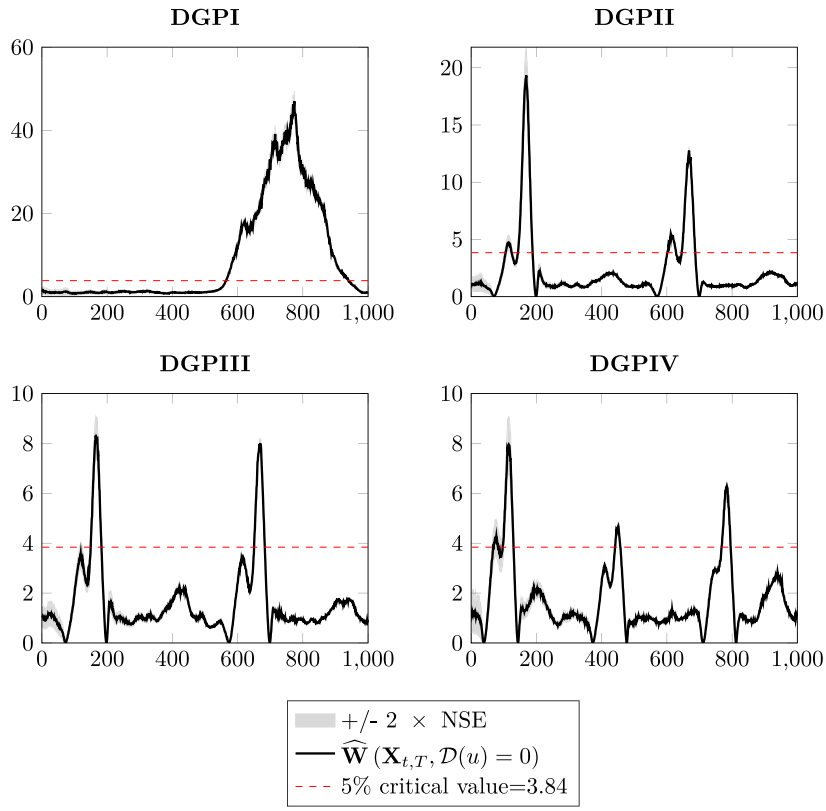

Fig. 3. Tests for differences between network connectedness across different frequency bands.
Notes: This figure plots the test statistics and +/− 2 × the numerical standard errors for heterogeneities between network connectedness measures across different frequency bands from four DGPs. The test statistic is $\widehat{\mathbf{W}}\left(\boldsymbol{X}_{t,T}, \mathcal{D}(u) = 0\right) - 1$ where $\mathcal{D}(u) = \widehat{C}(u,d) - \widehat{C}(u,c)$ with $d \in (0, \pi/5)$ corresponding to the low frequency band and $c \in (\pi/5, \pi)$ corresponding to the high frequency band. The dashed black line is the 5% critical value from the $\chi^2(1)$ distribution = 3.84. Values greater than 3.84 reject the null hypothesis of equivalent network connections across frequency band d and c in favour of differences. DGPI (top left panel) is a TVP VAR(2) model with Gaussian errors, we introduce a break in the time-varying autoregressive matrices and contemporaneous relations from observation 500 that induces large connections across the high frequency band. DGPII (top right panel) is a TVP VAR(2) where time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIII (bottom left panel) is a TVP VAR(2) model with student-t errors, time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIV (bottom right panel) is the same as DGPII, but with an increase in the periodicity of the respective sin wave functions the time-varying intercepts and autoregressive matrices follow.

The top row corresponds to DGPI, and the second, third and fourth rows results from DGPII, DGPIII, and DGPIV respectively. As we can see, the distribution of estimates for each DGP track the true values remarkably well. In almost all cases, the true value lies within the 95% quantiles of the distribution from model estimates. This plot shows that our method provides an accurate representation of horizon specific network connectedness, even when the underlying process has complex dynamics and the true error distribution is non-Gaussian.

Fig. 3 further reports the estimates, and numerical standard error bounds of our test for heterogeneities between high frequency band and low frequency band network connectedness measures (Eq. (15)) for each of the four DGPs we use in the Monte Carlo study. Specifically, for each observation we test the null hypothesis $\mathcal{D}(u) = 0$ where $\mathcal{D}(u) = \widehat{C}(u,d) - \widehat{C}(u,c)$ with $d \in (0, \pi/5)$ corresponding to the low frequency band and $c \in (\pi/5, \pi)$ corresponding to the high frequency band. In each plot, we also report the 5% critical value from the $\chi^2(1)$ distribution of 3.84. Test statistics exceeding this value reject the null in favour of heterogeneities between network connectedness measures across frequency bands.[18] As we can see, our test statistic identifies significant differences between low and high frequency band network connectedness measures that correspond with the peaks we observe in Fig. 2 for each DGP. We can see with the non-Gaussian DGP, DGPIII the estimates of the test statistics are smaller relative to the analogous Gaussian DGP, DGPII. However, there are still clear rejections.

We now test for differences in connectedness measures across the same frequency band over time. To do so, we test the first time period against all remaining 999 observations from our DGPs, $u_1 = 1$, $u_2 = \{2, 3, \ldots, 1,000\}$. We do this for both the low-frequency band, $d \in (0, \pi/5)$, and the high-frequency band, $d \in (\pi/5, \pi)$. Fig. 4 reports the estimates of the test statistics, their numerical standard error and the corresponding 95% critical value. First, considering DGPI, it is clear that connectedness across the high frequency band exhibits significant differences between the first observation when we increase connections at observation 500. It is noteworthy to mention that we also see this for the low frequency band. This is expected as we can see that the estimates from our simulations exhibit slight bias here. However, for DGPII–DGPIV we have rejections relative to the first observation at the corresponding periods where we create connectedness across the low frequency band, and no rejections across the high frequency band.

Overall, this shows that our testing procedure indicates rejections of equality in connectedness forming over different frequency bands, and over time for connectedness across the same frequency band, where we should expect to see such differences.

[18] We obtain numerical standard errors in a similar manner to Li et al. (2015).

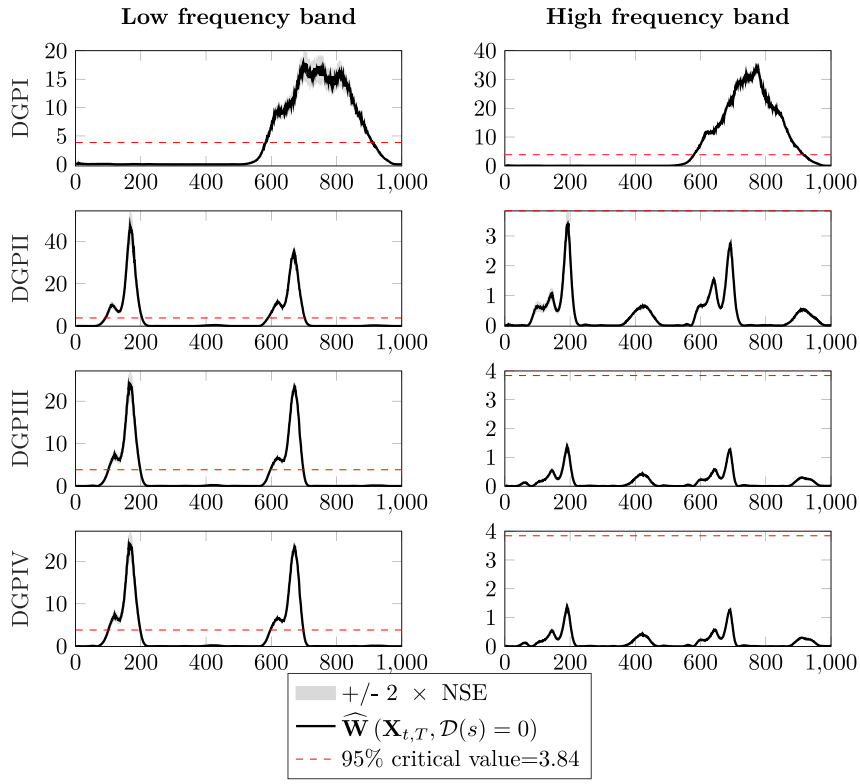

Fig. 4. Tests for differences between network connectedness over time.
Notes: This figure plots the test statistics and +/− 2 × the numerical standard errors for heterogeneities between network connectedness measures across different frequency bands from four DGPs. The test statistic is $\widehat{\mathbf{W}}\left(\boldsymbol{X}_{t,T}, \mathcal{D}(s)=0\right)-1$ where $\mathcal{D}(u)=\widehat{C}(u_1,d)-\widehat{C}(u_2,d)$ with $u_1=1$, $u_2=\{2,3,\ldots,1{,}000\}$. Here $d \in (0, \pi/5)$ corresponds to the low frequency band (LHS plots) and $d \in (\pi/5, \pi)$ corresponds to the high frequency band (RHS plots). The dashed black line is the 5% critical value from the $\chi^2(1)$ distribution = 3.84. Values greater than 3.84 reject the null hypothesis of equivalent network connections at time u_1, u_2 across frequency band d. DGPI (top row) is a TVP VAR(2) model with Gaussian errors, we introduce a break in the time-varying autoregressive matrices and contemporaneous relations from observation 500 that induces large connections across the high frequency band. DGPII (second row) is a TVP VAR(2) where time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIII (third row) is a TVP VAR(2) model with student-t errors, time-varying intercepts and autoregressive matrices following sin wave functions with a stochastic error, time-varying covariance matrix where the off-diagonals follow sin wave functions with a stochastic error, and the diagonal elements follow a stationary AR(1) processes. DGPIV (fourth row) is the same as DGPII, but with an increase in the periodicity of the respective sin wave functions the time-varying intercepts and autoregressive matrices follow.

4. Monitoring persistence in uncertainty networks using S&P500 stocks

Changes in uncertainty can play a key role in driving business cycles and financial turmoil (Bloom et al., 2018). Identifying the sources of such risks is a focus for researchers and practitioners. Some quantify systemic risks emanating from financial markets and sectors (e.g. Acemoglu, Ozdaglar, & Tahbaz-Salehi, 2015; Billio et al., 2012), while others examine how sectoral shocks affect aggregate fluctuations (e.g. Acemoglu et al., 2017; Atalay, 2017; Gabaix, 2011). Related to the above, and in response to financial crises, many countries are implementing policies to monitor systemic risk and financial stability. Therefore, we use our framework to identify new measures of transitory and persistent linkages for S&P500 constituents. In Section 4.1, we examine sectoral connectedness as well as network structures at a granular level for financial firms.

We use high-frequency tick data for all stocks listed on the S&P500 from 5 July 2005 to 31 August 2018 and compute realized volatility (RV) for all stocks in the sample. To obtain firm-level RVs, we restrict our analysis to five-minute returns during New York Stock Exchange (NYSE) trading hours (i.e. 09:30-16:00). The data are time-synchronized using the same timestamps, eliminating transactions executed on Saturdays and Sundays, US holidays, 24-26 December and 31 December to 2 January due to low activity on these days. This leaves us with 3278 trading days. After cleaning the data, we are left with a cross section of 496 stocks.

To obtain our network connectivity measures, we estimate a TVP VAR model on $N = 496$ stocks with $p = 2$ lags on our $T = 3278$ days of data. We estimate our horizon-specific dynamic network measures on a 48-core server. For each $t \in \{1, 2, \ldots, T\}$, we generate 500 simulations of the (quasi) posterior distribution, resulting in a total estimation time of 10 days. We define transitory (short-term) network links as those that form over a 1-5 day horizon, and persistent (long-term) network links at horizons greater than 5 days (i.e. 5 days to the ∞ horizon). Our choice of these horizons stems from the existing literature on volatility modelling using high-frequency data, which shows that daily and weekly fluctuations contain salient information for future volatility (e.g. Corsi & Renò, 2012).[19]

Fig. 5 plots measures of transitory and persistent network connectedness from 5 July 2005 to 31 August 2018. Overall, there are significant differences in the level of horizon-specific connectedness across our estimation sample. In general, long-term linkages are muted during periods of economic/financial calm. However, it is clear that long-term connectedness spikes during periods of economic recession or major stock market events. For example, long-term connectedness starts to rise in 2006 and continues to rise during the 2007–2009 recession. In addition, we can see that long-term connectedness increases during 2010–2012. This may be due to fears of contagion from the European

[19] In this case, due to the size of the system, we diagonalize the covariance matrix of the VAR as an additional precaution to avoid overfitting. For the sectoral network measures in Section 4.1, we use a full covariance matrix.

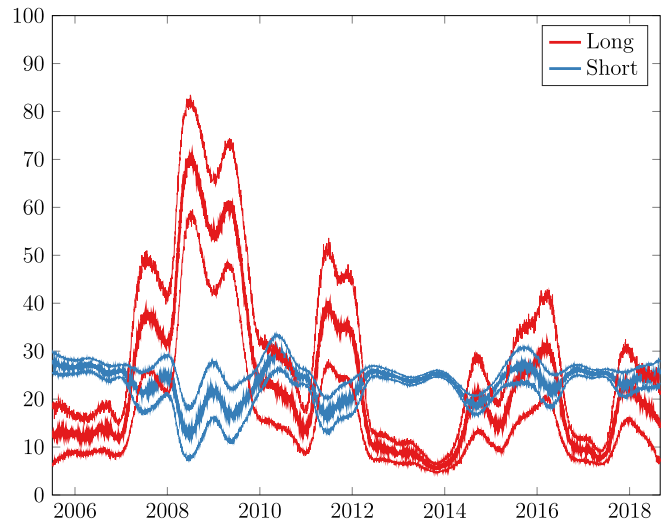

Fig. 5. Horizon specific dynamic total network connectedness for S&P500 constituents. This figure plots the posterior median and 1-standard deviation percentiles of horizon specific dynamic total network connectedness, $C(u,d)$, $d \in \{S, L\}$ from July 5, 2005 to August 31, 2018. S refers to the short-term, transitory connectedness, which we define as 1 day to 1 week; and L refers to long-term, persistent connectedness, which we define as horizons > 1 week. The spectrum with which the horizons stem from link to the frequency with which we observe the data.

sovereign debt crisis, the 2010 flash crash and when the S&P500 entered a bear market in 2011, albeit a short-lived one. We can also see an increase in short- and long-term connectedness in mid-to-late 2015, which is consistent with the stock market sell-off starting in August 2015; this may also be related to fears of contagion from the Chinese stock market crash in late 2015. Overall, we see that long-term network connectedness increases during periods of high systemic risk across our sample.

4.1. Transitory and persistent network connectedness of S&P500 sectors

Here we focus on the overall network connectedness driven by transitory and persistent shocks to companies in a given sector. We classify stocks into eleven main sectors according to the Global Industry Classification Standard (GICS).[20] These are: Consumer Discretionary (COND) with 73 stocks; Consumer Staples (CONS) with 34 stocks; Health Care (HLTH) with 53 stocks; Industrials (INDU) with 73 stocks; Information Technology (INFT) with 67 stocks; Materials (MATR) with 33 stocks; Real Estate (REAS) with 29 stocks; Financials (SPF) with 66 stocks; Energy (SPN) with 36 stocks; Communication Services (TELS) with 6 stocks; and Utilities (UTIL) with 26 stocks. Further details, including descriptive statistics (Table D.1 of Appendix D) for annualized daily RVs, which we compute as $100 \times \sqrt{252 \times RV_t}$ pooling information across companies within each sector over the sample period from 5 July 2005 to 31 August 2018, are reported in Appendix D. The energy sector has the highest mean, while the financial sector has the highest standard deviation, skewness and kurtosis. Overall, we can see that these sectoral RVs show significant differences in terms of the first four moments as well as the minimum and maximum values.

For each of the 11 sectors, we obtain dynamic network measures by estimating a TVP VAR model on all stocks with two lags on our 3278 days of data. On each day, we take 500 draws from the (quasi) posterior distribution.[21] We define transient and persistent network connectivity in the same way as above.

In Fig. 6 we plot the posterior median and 95% confidence bands for transitory and persistent network connectivity for each sector as in Eq. (8). These follow directly from the manipulations of the estimates of the dynamic adjacency matrices (see Eq. (12)). Grey bars in these figures represent periods where there are statistically significant differences at the 5% level between transient and persistent network connectivity.[22] Overall, we observe significant differences between these horizon-specific networks for each sector. In general, network connectedness due to the persistent component of shocks exceeds that due to the transitory component of shocks during periods of market turbulence. Then, during periods of calm, the transitory part of the networks becomes more pronounced.

Comparing these measures across sectors, we can see from the real estate and financial sectors that surges in network connectedness driven by persistent shocks drive uncertainty in the sector for much longer periods of time during the Great Recession relative to other sectors, e.g. CONS and HLTH. In addition, the magnitude of persistent network connectedness from the real estate and financial sectors is much greater. Note also that throughout this period we observe much more frequent evidence in favour of statistical differences between networks driven by transitory and persistent shocks for these sectors; particularly relative to COND, CONS, HLTH and MATR. Although the other sectors show spikes in persistent network connectedness, these do not occur until around February 2008. This highlights how long-term systemic risks within the real estate and financial sectors intensify during this period and are stronger relative to other sectors.

We also see clear spikes in persistent sectoral network connectedness in May–October 2011 and again in 2015–2016. The former coincides with the S&P500 entering a bear market and the latter with declines in major stock markets around the world. In 2015–2016, we see much lower long-term connectedness of consumer staples, utilities, real estate and telecoms relative to other sectors. We expect the long-term linkages of industrials, materials, energy, information technology and financials to be strong during this period, as the decline in global equity markets is linked to falling commodity prices and the depreciation of Asian currencies against the US dollar.

In terms of sectoral network linkages due to transitory shocks, the main differences we observe are in the size of the linkages. Note, however, that there are subtle differences in the time profiles. For each sector, we see that transitory linkages are much stronger during periods of uncertainty at the firm level than persistent linkages during periods of calm.

Digging deeper into our investigation of whether there are significant differences between transitory and persistent network connectedness for each sector, two main points emerge. First, we document significant heterogeneity in persistence over substantial periods of time. Fig. 6 shows that during periods of tranquillity, transitory network linkages between sectors are stronger relative to linkages from persistent shocks. During periods of turbulence, however, our results show that persistent linkages intensify in all sectors, but their magnitude differs considerably.[23] Second, statistical differences between transitory and persistent connectedness in one sector do not necessarily imply differences in other sectors. For example, over the 2014–2017 period, we see clear differences in the rejections that transitory and persistent connectedness are equivalent in REAS, SPF, SPN and UTIL. In general, the temporal nature of our results aligns well with Bianchi, Billio,

[20] GICS is an industry taxonomy developed by MSCI and Standard & Poor's for use by the global financial community.

[21] We estimate our horizon-specific dynamic network measures on a 64-core server, resulting in a total estimation time of about 4–5 h to obtain network estimates from the 11 sectors.

[22] We plot the values of these test statistics in Figure D.1 in Appendix D.

[23] From Figure D.1 in Appendix D we also document that the magnitude of the differences evolves substantially over time.

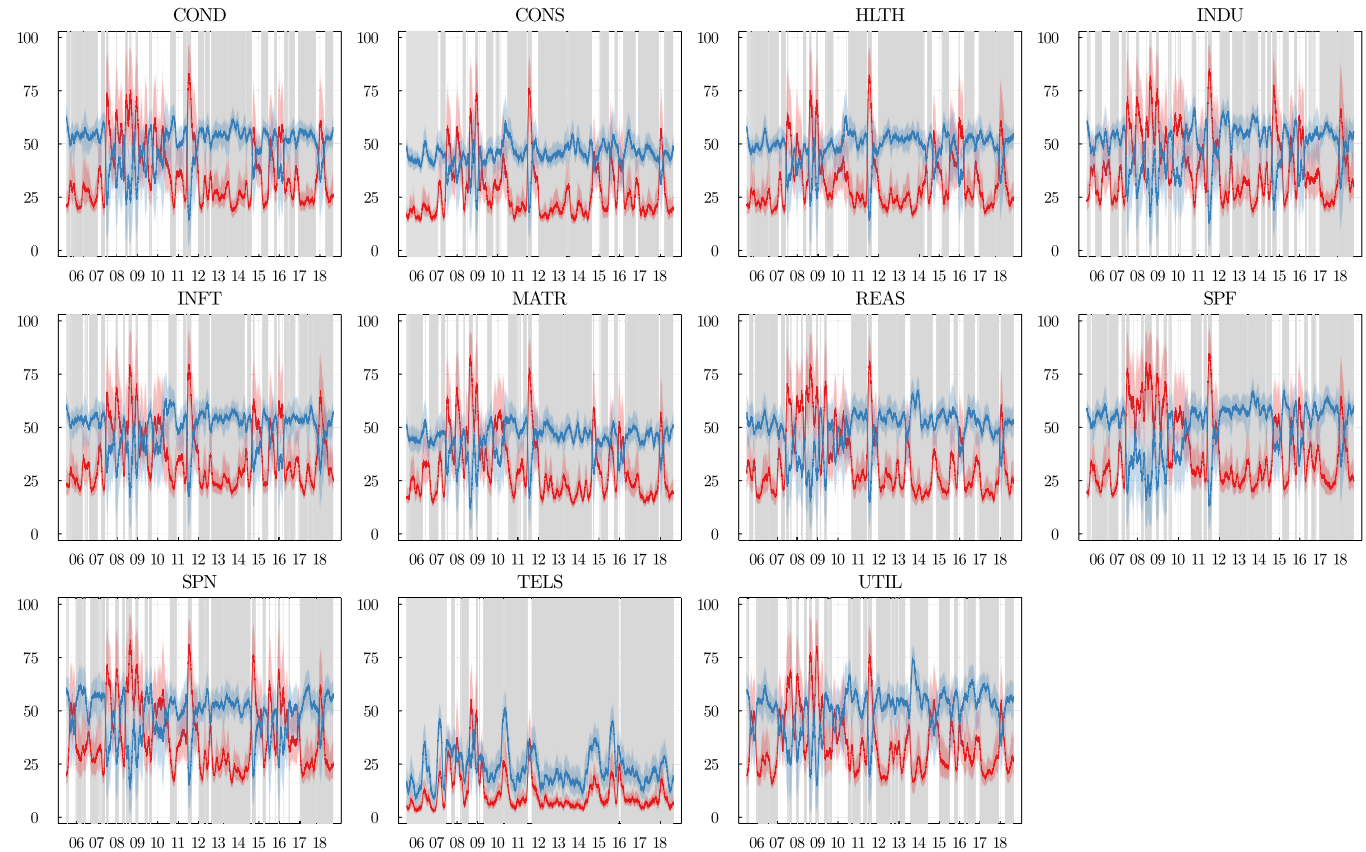

Fig. 6. Dynamics and Persistence in the U.S. Uncertainty Networks.
This figure plots the (quasi) posterior median and 95% confidence bands of network connectedness specific to the transitory, or short-term (in blue) and persistent, or long-term (in red) shocks to realized volatility of the S&P500 sectors: Consumers Discretionary (COND), Consumer Staples (CONS), Health Care (HLTH), Industrials (INDU), Information Technology (INFT), Materials (MATR), Real Estate (REAS), Financials (SPF), Energy (SPN), Communication Services (TELS), and Utilities (UTIL) from July 8, 2005 to August 31, 2018. We define transitory (short-term), as connections made over the 1 day to 1 week horizon; we characterize connections greater than 1 week as persistent (long-term). Grey bars indicate periods with significant differences in transitory and persistent connectedness. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Casarin, and Guidolin (2019), which documents a regime dependent impact of systemic risk on financial markets. They also provide empirical support for Gabaix (2011) and Acemoglu et al. (2012), and uncover new measures of sectoral uncertainties (or sector-wide risks). Testing for statistical differences between transitory and persistent sectoral network connectedness adds further substance to our suggestion that one should consider dynamic network structures that form across frequency bands.

For researchers, our measures of network connectedness may contain useful information for real economy or forecasting purposes; there is already evidence that network connectedness contains predictive content for the real economy (e.g. Baruník et al., 2020). For practitioners, tracking persistence in sectoral networks can be useful for informing macroprudential policy. This is because one can use these measures as online monitoring tools to study the evolution and persistence of sectoral network connectedness.

Although we provide evidence of substantial heterogeneity in the persistence of network structures, there are commonalities in the time profiles. We attribute this to the high degree of correlation between the data used to proxy uncertainty in our investigation. Herskovic, Kelly, Lustig, and Van Nieuwerburgh (2016) exploits the correlation structure of idiosyncratic return volatilities and shows that a common factor among the drivers of firm-level volatilities has pricing implications. Our network connectedness measures, by definition, refer to this correlation structure and provide an aggregate description of the network at each point in time. However, our network measures are more informative. We are able to obtain measures that contain information about the overall network structure. These may relate to directional connections (in and out degrees) or concentration (i.e. a high influence of a small number of firms/nodes on the overall network); both of which have been shown to contain information with economic implications (see e.g. Herskovic, 2018; Herskovic et al., 2020).

4.2. Network connections at a granular level

By focusing on shocks to a single financial institution that affect the wider system, our research contributes to the large literature on measuring systemic risk. Many understand systemic risk as many market participants realizing severe losses as a result of propagation throughout the (financial) system.[24] During periods of financial turbulence, uncertainty shocks, the drying up of liquidity, and insolvencies have the ability to spread rapidly affecting many institutions across the market. Lessons from recent financial crises spur the demand for financial regulations in order to mitigate firm behaviours consistent with increasing systemic risk.

From a prudential perspective, safeguarding against this type of risk requires quantifying systemic risk. The existing literature offers many measures of such types of risk, these include: the expected shortfall measure of Acharya et al. (2017); Co-Value-at-Risk (CoVaR) (Adrian & Brunnermeier, 2016); and network connectedness measures (see e.g.

[24] For a comprehensive review of the literature on systemic risk, see Benoit, Colliard, Hurlin, and Pérignon (2017).

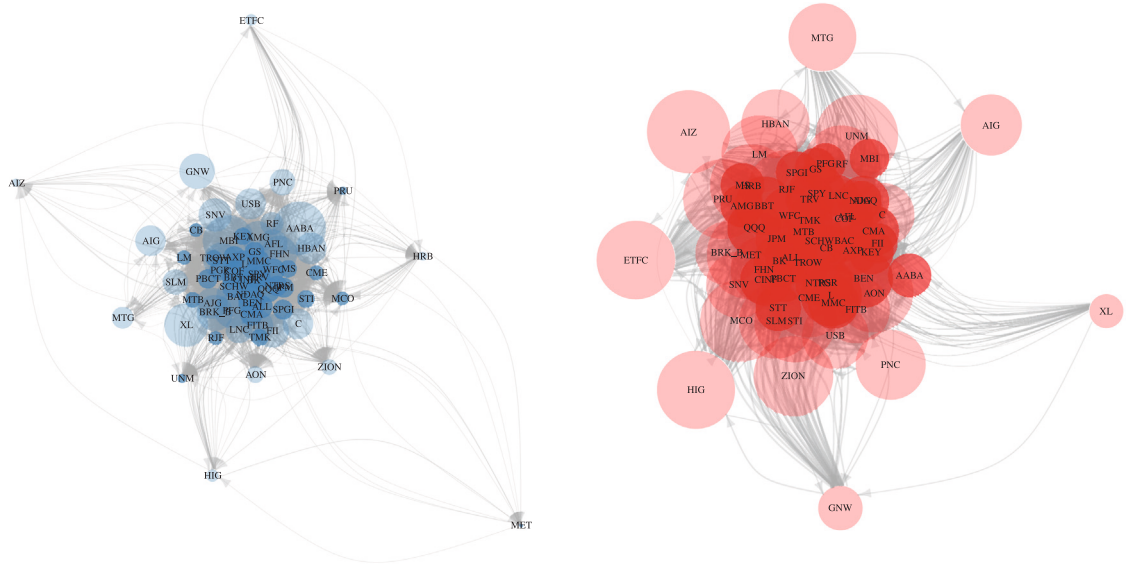

Fig. 7. Transitory and persistent networks of finance: 24 October 2008.
The left (right) figure shows the network connections between the assets comprising the SPF sector driven by transitory (persistent) shocks on 24 October 2008, corresponding to the day when the VIX peaked. Arrows indicate the direction of the connections and the strength of the lines indicates the strength of the connections. Grey (black) vertices indicate firms that receive (send) more shocks than they send (receive). The size of the vertices indicates the net amount of shocks. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Demirer et al., 2018). The former measures relate to specific risk channels and as such can aid in calibrating regulatory tools. The latter approach permits one to quantify the overall influence of individual institutions to overall systemic risk, and hence identify systemically important financial institutions (SIFIs).

Our approach relates closely with the above. However, under our framework one can characterize SIFIs, or important variables for different applications, throughout time as well as understanding whether the influence is persistent or transitory in nature. The benefit of this is twofold. First, enhancing prudential authorities' understanding of whether SIFIs influence are transient or long lasting can help refine policies in order to mitigate adverse firm behaviour. For example, one could tailor policies by increasing the capital requirements of SIFIs for those who transmit persistent shocks that contribute significantly to systemic risk. Second, since systemic risk threatens the stability of the entire financial sector, knowing the frequency-specific sources of instability facilitates monitoring of changes to the such risks.

To illustrate how policy makers might use our approach, we examine the network structures of the financial sector at the granular level of 65 firms. Since our application uses daily data we have transitory and persistent network structures at every observation in our sample. We therefore focus on two different dates. Fig. 7 shows the network structures driven by transitory and persistent shocks for the SPF sector on October 24, 2008. This date corresponds to the start of the global financial crisis. Fig. 8 reports the corresponding network structures one year later on October 24, 2009. For each plot, arrows indicate the direction and strength of the connections, while a transparent (full colour) vertex indicates a stock that sends (receives) more shocks than it receives (sends). The size of the vertices indicates the net direction of the connections.

We can see that on 24 October 2008, the persistent links are larger relative to the transitory ones, suggesting that shocks within the financial sector create links that relate to the long-term. This suggests that systemic risk within the systems stems from persistent network structures. Now looking one year later, it is clear that connections are far weaker across both transitory and persistent network structures and systemic risk is relatively lower. The main takeaway from these plots is the strong differences in the overall structure of the horizon-specific networks.

In Appendix D, we plot heatmaps showing the strength of financial institutions connections across persistent and transitory network structures for these same two dates in Figures D.2–D.3. Persistent shocks tend to drive the links with greater strength, and so we focus our discussion here. Zooming in to examine the contribution of specific firms, we can see that Truist Financial Corp (BBT), Franklin Resources (BEN), Loews Corporation (L), SPY, or Wells Fargo & Co (WFC) transmitted persistent shocks to the financial sector and thus are identified as SIFIs that affected the system with persistent shocks. We can follow the contributions from the columns of the heatmap in Figure D.2. The impact of a bank increases with the number of rows containing a stronger and warmer red colour. As we can see, those banks we name above affect many other financial institutions at the start of the financial crisis, and such impacts are long-lasting.

Conversely, Metlife (MET), Moody's (MCO), Unum (UNM), H&R Block (HRB) and Assurant (AIZ) receive the most shocks on 24 October 2008. One year later, we can see that the structure changes dramatically. While WFC, SPY and L seem to be strong SIFIs, BEN is a nearly non-contributing bank. This highlights how our approach tracks dynamics of key financial institutions within the system and their influence across persistent and transitory network structures. For completeness, we rank all institutions in the financial sector according to the strength of transitory and persistent shocks they transmit/receive during the same two dates in Table D.2 in Appendix D.

Overall, these network structures show how the role of a firm can change not only over time, but also in terms of persistence. Researchers may wish to use these network structures to assess the pricing implications of such risks. For example, users could assess the role of directional linkages in an empirical asset pricing application that builds on the theoretical work of Branger, Konermann, Meinerding, and Schlag (2020). Our framework also allows one to track dynamic network structures that would complement studies such as (Herskovic, 2018) and Gofman, Segal, and Wu (2020). The advantage of our approach is that one does not have to rely on monthly or annual data to capture such networks. Finally, looking at dynamic adjacency matrices can help economists understand how shocks dynamically determine network structures in models of monetary policy (e.g. Pasten, Schoenle, & Weber, 2020) or the Phillips curve (e.g. Rubbo, 2020).

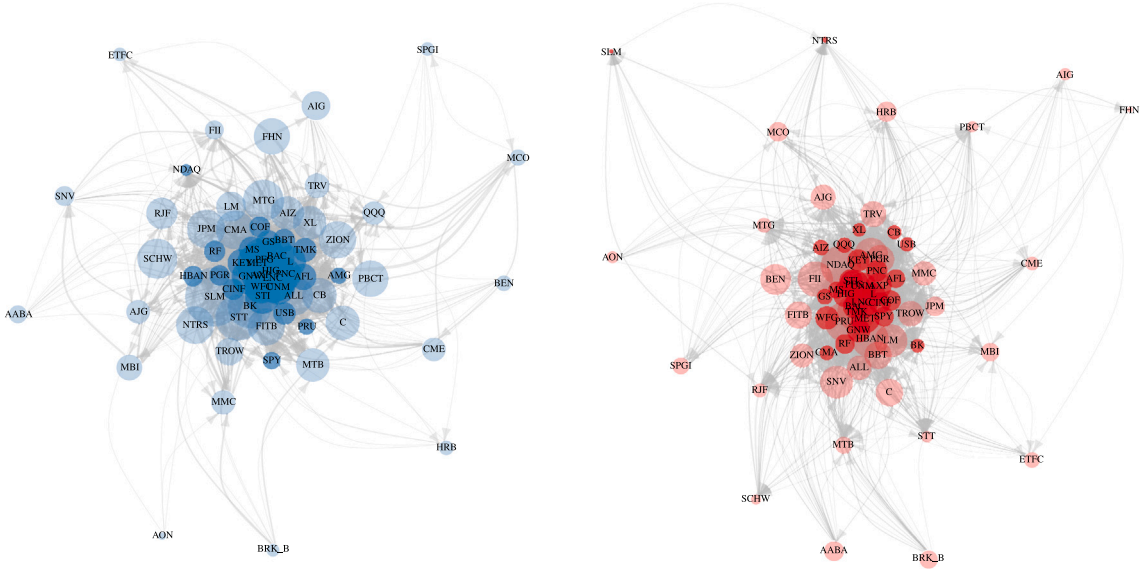

Fig. 8. Transitory and Persistent Network of Financials: 24 October 2009.
The left (right) figure shows the network connections between the assets that make up the SPF sector by transitory (persistent) connectedness. Arrows indicate the direction of the connections and the strength of the lines indicates the strength of the connections. Grey (black) vertices denote firms that receive (send) more shocks than they send (receive). The size of the vertices indicates the net amount of shocks. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

5. Conclusion

This paper proposes a novel framework for measuring dynamic network connections in a multivariate data system. We track dynamic connections driven by different degrees of persistence using a spectral decomposition of time-varying variance decomposition matrices. Our approach properly accounts for the characteristics of the shocks that create such links. We outline a procedure that allows one to test for statistical differences in connectedness over time and frequency. We provide Monte Carlo evidence that our measures are able to reliably track connectedness and correctly identify statistical differences from different data generating processes.

Empirically, we show that transitory and persistent measures of network connectedness improve our understanding of systemic risks arising from uncertainty networks. This is because our approach allows one to track connectedness across the transitory and persistent components of shocks. This is particularly useful during periods of heightened uncertainty, as our measures indicate whether systemic risks from network connections are transitory or persistent in nature. Ultimately, this could lead to better decision-making by macroprudential supervisors and investment decisions by market participants.

Disclosure statement

Jozef Baruník and Michael Ellington have nothing to disclose.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ejor.2023.11.023.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., & Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, *80*(5), 1977–2016.

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, *105*(2), 564–608.

Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2017). Microeconomic origins of macroeconomic tail risks. *American Economic Review*, *107*(1), 54–108.

Acharya, V., Engle, R., & Richardson, M. (2012). Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review*, *102*(3), 59–64.

Acharya, V. V., Pedersen, L. H., Philippon, T., & Richardson, M. (2017). Measuring systemic risk. *The Review of Financial Studies*, *30*(1), 2–47.

Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, *106*(7), 1705–1741.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, *9*(4), 254–280.

Balke, N. S., & Wohar, M. E. (2002). Low-frequency movements in stock prices: A state-space decomposition. *The Review of Economics and Statistics*, *84*(4), 649–667.

Bandi, F. M., Chaudhuri, S. E., Lo, A. W., & Tamoni, A. (2021). *Journal of Financial Economics*, *142*(1), 214–238.

Bandi, F., & Tamoni, A. (2017). Business-cycle consumption risk and asset prices. Available at SSRN 2337973.

Bansal, R., Kiku, D., & Yaron, A. (2010). Long run risks, the macroeconomy, and asset prices. *American Economic Review*, *100*(2), 542–546.

Bansal, R., & Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, *59*(4), 1481–1509.

Barigozzi, M., Hallin, M., Soccorsi, S., & von Sachs, R. (2020). Time-varying general dynamic factor models and the measurement of financial connectedness. *Journal of Econometrics*.

Baruník, J., Bevilacqua, M., & Tunaru, R. (2020). Asymmetric network connectedness of fears. *The Review of Economics and Statistics*, 1–41.

Baruník, J., & Křehlík, T. (2018). Measuring the frequency dynamics of financial connectedness and systemic risk. *Journal of Financial Econometrics*, *16*(2), 271–296.

Barunik, J., Krehlik, T., & Vacha, L. (2016). Modeling and forecasting exchange rate volatility in time-frequency domain. *European Journal of Operational Research*, *251*(1), 329–340.

Benoit, S., Colliard, J.-E., Hurlin, C., & Pérignon, C. (2017). Where the risks lie: A survey on systemic risk. *Review of Finance*, *21*(1), 109–152.

Bianchi, D., Billio, M., Casarin, R., & Guidolin, M. (2019). Modeling systemic risk with Markov switching graphical SUR models. *Journal of Econometrics*, *210*(1), 58–74.

Billio, M., Getmansky, M., Lo, A. W., & Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, *104*(3), 535–559.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., & Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, *86*(3), 1031–1065.

Branger, N., Konermann, P., Meinerding, C., & Schlag, C. (2020). Equilibrium asset pricing in directed networks. *Review of Finance*, *25*(3), 777–818.

Brennan, M. J., & Zhang, Y. (2018). Capital asset pricing with a stochastic horizon. *Journal of Financial and Quantitative Analysis*, 1–45.

Calabrese, R., & Osmetti, S. A. (2019). A new approach to measure systemic risk: A bivariate copula model for dependent censored data. *European Journal of Operational Research*, *279*(3), 1053–1064.

Chan, J. C. (2020). Large Bayesian VARs: A flexible Kronecker error covariance structure. *Journal of Business & Economic Statistics*, *38*(1), 68–79.

Chaudhuri, S. E., & Lo, A. W. (2019). Dynamic alpha: A spectral decomposition of investment performance across time horizons. *Management Science*, *65*(9), 4440–4450.

Cogley, T., Primiceri, G. E., & Sargent, T. J. (2010). Inflation-gap persistence in the US. *American Economic Journal: Macroeconomics*, *2*(1), 43–69.

Corsi, F., & Renò, R. (2012). Discrete-time volatility forecasting with persistent leverage effect and the link with continuous-time volatility modeling. *Journal of Business & Economic Statistics*, *30*(3), 368–380.

Dahlhaus, R. (1996). On the Kullback-Leibler information divergence of locally stationary processes. *Stochastic Processes and Their Applications*, *62*(1), 139–168.

Dahlhaus, R. (2000). A likelihood approximation for locally stationary processes. *The Annals of Statistics*, *28*(6), 1762–1794.

Dahlhaus, R., Polonik, W., et al. (2009). Empirical spectral processes for locally stationary time series. *Bernoulli*, *15*(1), 1–39.

Demirer, M., Diebold, F. X., Liu, L., & Yilmaz, K. (2018). Estimating global bank network connectedness. *Journal of Applied Econometrics*, *33*(1), 1–15.

Dew-Becker, I., & Giglio, S. (2016). Asset pricing in the frequency domain: theory and empirics. *The Review of Financial Studies*, *29*(8), 2029–2068.

Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, *182*(1), 119–134.

Drechsler, I., & Yaron, A. (2011). What's vol got to do with it. *The Review of Financial Studies*, *24*(1), 1–45.

Engle, R., & Kelly, B. (2012). Dynamic equicorrelation. *Journal of Business & Economic Statistics*, *30*(2), 212–228.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, *79*(3), 733–772.

Garvey, M. D., Carnovale, S., & Yeniyurt, S. (2015). An analytical framework for supply network risk propagation: A Bayesian network approach. *European Journal of Operational Research*, *243*(2), 618–627.

Geraci, M. V., & Gnabo, J.-Y. (2018). Measuring interconnectedness between financial institutions with Bayesian time-varying vector autoregressions. *Journal of Financial and Quantitative Analysis*, *53*(3), 1371–1390.

Gerrard, R., Kyriakou, I., Nielsen, J. P., & Vodička, P. (2022). On optimal constrained investment strategies for long-term savers in stochastic environments and probability hedging. *European Journal of Operational Research*.

Giglio, S., Maggiori, M., & Stroebel, J. (2015). Very long-run discount rates. *Quarterly Journal of Economics*, *130*(1), 1–53.

Gofman, M., Segal, G., & Wu, Y. (2020). Production networks and stock returns: The role of vertical creative destruction. *The Review of Financial Studies*, *33*(12), 5856–5905.

Hahn, J., Todd, P., & Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, *69*(1), 201–209.

Haven, E., Liu, X., & Shen, L. (2012). De-noising option prices with the wavelet method. *European Journal of Operational Research*, *222*(1), 104–112.

Herskovic, B. (2018). Networks in production: Asset pricing implications. *The Journal of Finance*, *73*(4), 1785–1818.

Herskovic, B., Kelly, B., Lustig, H., & Van Nieuwerburgh, S. (2016). The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics*, *119*(2), 249–283.

Herskovic, B., Kelly, B., Lustig, H., & Van Nieuwerburgh, S. (2020). Firm volatility in granular networks. *Journal of Political Economy*, *128*(11), 4097–4162.

Koop, G., Leon-Gonzalez, R., & Strachan, R. W. (2010). Dynamic probabilities of restrictions in state space models: an application to the phillips curve. *Journal of Business & Economic Statistics*, *28*(3), 370–379.

Li, Y., Liu, X.-B., & Yu, J. (2015). A Bayesian chi-squared test for hypothesis testing. *Journal of Econometrics*, *189*(1), 54–69.

Li, Y., Zeng, T., & Yu, J. (2014). A new approach to Bayesian hypothesis testing. *Journal of Econometrics*, *178*, 602–612.

Liu, X., Li, Y., Yu, J., & Zeng, T. (2022). Posterior-based wald-type statistics for hypothesis testing. *Journal of Econometrics*, *230*(1), 83–113.

Lütkepohl, H. (1990). Asymptotic distributions of impulse response functions and forecast error variance decompositions of vector autoregressive models. *The Review of Economics and Statistics*, 116–125.

Pasten, E., Schoenle, R., & Weber, M. (2020). The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics*, *116*, 1–22.

Pesaran, H. H., & Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, *58*(1), 17–29.

Petrova, K. (2019). A quasi-Bayesian local likelihood approach to time varying parameter VAR models. *Journal of Econometrics*.

Primiceri, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, *72*(3), 821–852.

Richmond, R. J. (2019). Trade network centrality and currency risk premia. *The Journal of Finance*, *74*(3), 1315–1361.

Roueff, F., & Sanchez-Perez, A. (2016). Prediction of weakly locally stationary processes by auto-regression. arXiv preprint arXiv:1602.01942.

Rubbo, E. (2020). Networks, phillips curves and monetary policy. Available https://economics.yale.edu/sites/default/files/rubbo_jmp.pdf.

Sévi, B. (2014). Forecasting the volatility of crude oil futures using intraday data. *European Journal of Operational Research*, *235*(3), 643–659.

Sun, E. W., & Meinl, T. (2012). A new wavelet-based denoising algorithm for high-frequency financial data mining. *European Journal of Operational Research*, *217*(3), 589–599.

Yang, Z., & Zhou, Y. (2017). Quantitative easing and volatility spillovers across countries and asset classes. *Management Science*, *63*(2), 333–354.