Fat Tails, Serial Dependence, and Implied Volatility Index Connections

Michael Ellington*

September 26, 2021

Abstract

This paper accounts for fat tails and serial dependence for implied volatility index network connections among equity and commodity markets using Bayesian vector heterogeneous autoregressions. I analyse the information content of such connections over short-, medium- and long-horizons for predicting underlying asset returns and whether conventional asset pricing risk factors explain the variation of portfolios that sort on directional connections. Including network connections within the information set yields significant gains when forecasting underlying asset returns. Sorting underlying assets on directional connections shows that investors can hedge against temporary changes to investment opportunities at horizons of less than one month.

JEL: C11, C58, G10, G17

Keywords: Finance; Network Analysis; Forecasting; Volatility Spillovers; Serial Dependence; Fat Tails.

* University of Liverpool Management School, Chatham Building, Chatham Street, L69 7ZH, UK. Email addresses: Michael Ellington, m.ellington@liverpool.ac.uk

1 Introduction

Market participants monitor implied volatility indices to extract information around future price fluctuations of the underlying asset. These expectations around future financial conditions and asset market movements can dictate investor behaviour, influence trading decisions for speculation purposes, and also inform hedging decisions. Implied volatilities, like realized volatilities (see e.g. Herskovic et al., 2016), exhibit strong co-movement. Unlike realized volatility however, implied volatility is the ex-ante risk-neutral expectation of future volatility. The common component within volatilities indicates the possibility of an underlying network structure creating these connections (Herskovic et al., 2020)[1].

As agents base their decisions over different horizons (see e.g. Baruník et al., 2016), tracking implied volatility connections over different horizons may be of practical use for forecasting and portfolio formation. For instance, an active trader may wish to extract information relevant to daily or weekly horizons for speculation or hedging purposes. In contrast, an investor with longer-term goals will desire information at longer horizons, such as months or even years, to guide investment decisions.

Typically, work focusing on connections use linear measures of association like correlations or copulas, and/or Gaussian approximating models; something that is unsuitable for modelling financial data which typically exhibit features such as clustering and extreme values (e.g. Creal et al., 2011). Concerning the former, correlation based measures as in Engle and Kelly (2012) are pairwise linear. Regarding the latter, for example Acharya et al. (2012, 2017); Adrian and Brunnermeier (2016), are linear and Gaussian. A notable exception is Calabrese and Osmetti (2019), who propose a copula model to estimate systemic risk within the banking sector that captures non-linear tail dependence. Nevertheless, these studies track connections that aggregate over horizons, and also do not measure the connections among variable within the system as a whole.

In this paper, I show how one can track network connections over different horizons using Bayesian vector heterogeneous autoregressive (BVHAR) models that account for fat tails, serial dependence, and heteroskedasticity. I use this to assess changes in network structures among implied volatility indices of equity and commodity markets. The first objective of this paper is to understand how different error covariance structures influence implied volatility index connections. The second objective is to explore the practical uses of such network connections for decision makers. This study is the first to proxy network connections from variance decompositions using models that allow for non-Gaussian errors, serial dependence,

[1] Acemoglu et al. (2012) provides the foundation that network structures can influence aggregate volatility fluctuations. Note also that connections among financial systems promote stability, but also facilitate shock propagation which can lead to increases in frailty (Acemoglu et al., 2015).

and stochastic volatility to purely forward-looking networks.

I build on Diebold and Yılmaz (2014) and Baruník and Křehlík (2018) who provide a unifying framework to measure financial connections from vector autoregressions (VARs) by viewing the forecast error variance decomposition matrix as an adjacency matrix. I extend the heterogenous autoregression (HAR) model of Corsi (2009) to a Bayesian multivariate setting using the flexible error covariance structures in Chan (2020). This results in modifications to network connectedness measures from VAR models that allow one to properly account for the persistence inherent in financial time series such as implied volatilities.

My main results demonstrate the practical uses of implied volatility index connections in two empirical exercises. The first explores the out-of-sample forecasting performance of implied volatility index network connections for the underlying asset returns. Point and density forecasts reveal statistically significant and economically meaningful gains when including network connectedness. The second sorts the underlying assets on directional connections to create long-short portfolios and explores whether there is an appealing risk return profile for this strategy. These results indicate that investors can construct hedge portfolios trading on directional connections forming over horizons of less than one month because they act as intertemporal hedging devices against transient adverse changes to investment opportunities (Campbell et al., 2018). These findings are robust to a battery of alternatives.

This work relates to four main areas of literature. First it pertains to those providing innovations in the measurement network connections. Some focus on using VAR models (see e.g. Demirer et al., 2018; Baruník et al., 2020). Meanwhile others use network connections to describe volatility spillovers (see e.g. Engle et al., 2012; Diebold and Yilmaz, 2012, 2015). Typically, these studies use VARs with Gaussian error distributions (exceptions are e.g. Geraci and Gnabo, 2018; Barbaglia et al., 2020, the former tracks connections using pairwise VAR coefficients, the latter uses ex-post volatility). This paper build on this literature by showing how one can properly account for persistence in financial time-series along with flexible, non-Gaussian, error covariance structures to purely forward looking networks.

Next, it relates to the body of work on modelling network connections and volatility spillovers (see e.g. Engle et al., 2012; Diebold and Yilmaz, 2012, 2015). However, the majority of studies focus on ex-post volatility and thus provides descriptions of network connections using historical measures. Baruník et al. (2020) argue that it is more informative to look at purely forward looking measures of volatility. This is because they provide better measurement of expectations regarding future price movements of the underlying asset. There are almost no studies that depart from Gaussian error distributions within the Diebold and Yılmaz (2014) framework (an exception is e.g. Barbaglia et al., 2020, who nonetheless use ex-post volatility in their model). My results indicate models with an error covariance structure

that accounts for fat tails, serial dependence, stochastic volatility best fit the data relative to simpler alternatives.

Third, it relates to studies using frequency domain techniques to improve forecasting. Sévi (2014) and Baruník et al. (2016) use wavelets to decompose realized volatility to improve forecasts[2]. The former do so to forecast volatility of crude oil futures. The latter disentangle jump and integrated variation over different horizons for foreign exchange rates. My results first show how using these techniques enhances our understanding of how network structures change in the time and frequency domain. Next, they reveal substantial gains when predicting underlying asset returns. This implies that decision makers are able to exploit the information content within these network structures that form over different horizons.

Finally, this study relates to work on the measurement of systemic risk, financial contagion, and volatility spillovers (see, for example: Yang and Zhou, 2017; Calabrese and Osmetti, 2019; Barigozzi et al., 2020). Typically, this literature provides point estimates of systemic risk or volatility spillover measures over different time periods, which causes two issues. First, point estimates make inference difficult and require bootstrapping procedures. Second, these time-domain measures aggregate over frequencies and tell us nothing about frequency-specific connections. An important contribution of this paper is that it provides plausible resolutions to these issues by: i) using Bayesian VAR models which accounts for parameter uncertainty within the estimation process; and ii) decomposing connections among a network of implied volatility indices into horizons of interest for investors and market participants.

2 Data and Methodology

2.1 Data

The sample comprises six daily implied Chicago Board of Options Exchange (CBOE) volatility indices. The data is from Thomson Reuters Datastream and spans 16 March, 2011 to September 3, 2020. Specifically for stock markets, I use the CBOE EFA ETF volatility index (VXEFA) and the VIX index. The former proxies the expected volatility of the iShares MSCI EAFE Index Fund that includes large international firms listed in Europe, Asia, and Australia. The latter measures the expected volatility of the S&P500 over the next 30 days[3].

[2]Frequency domain techniques are a popular method for removing noise from data (see e.g. Haven et al., 2012; Sun and Meinl, 2012).

[3]Using the VXEFA index allows one to examine connections across international equity markets without explicitly accounting for their respective volatility indices. In the Online Appendix I repeat the analysis

I also use various other implied volatility indices. They are: The CBOE crude oil ETF volatility index (OVX), which measures the expected volatility of crude oil over the next 30 days. This tracks the United States Oil fund ETF (USO) that predominantly holds short-term (i.e. 1-month) NYMEX futures contracts on West Texas Intermediate (WTI) Crude oil. The CBOE Energy sector ETF volatility index (VXXLE), which measures the implied volatility of the XLE ETF. This ETF tracks an index using weights from market-capitalisation of US energy firms on the S&P500. The CBOE COMEX gold volatility index (GVX). GVX measures the market expectation of 30-day ahead volatility of Comex Gold futures. The CBOE silver ETF volatility index (VXSLV). SLV is the iShares Silver Trust ETF tracking the spot price of silver less expenses and liabilities. VXSLV measures the implied volatility of silver.

Table 1 reports the descriptive statistics and contemporaneous correlations for the implied volatility indices in Panels A and B, respectively. Panel A shows that all indices have similar means and medians. The standard deviations range from 5.30 to 20.16 with OVX and GVX displaying the respective highest and lowest standard deviations. All series exhibit positive skewness and are leptokurtic, with OVX and GVX possessing the respective highest and lowest skewness and kurtosis values. Turning to Panel B, it is clear that all series have strong positive contemporaneous correlations. The precious metal implied volatility indices exhibit a correlation of 0.87, and the correlation between the VXXLE and OVX is 0.84. The lowest correlations are OVX and GVX, and OVX and VXSLV. I plot all implied volatility indices in the Online Appendix.

Table 1 confirms that implied volatilities exhibit co-movement, a well-known feature of financial data (see e.g. Herskovic et al., 2016). Furthermore, we know that volatilities are persistent and possess high degrees of excess kurtosis (Diebold and Yılmaz, 2014). This implies that error distributions may require heavy tails and/or serial dependence when modelling connections among asset volatilities.

2.2 Bayesian Vector Heterogeneous Autoregressions

In order to account for the above features of the data, I use a multivariate version of the heterogeneous autoregressive (HAR) model of Corsi (2009) and combine this with the Bayesian estimation methodology in Chan (2020). The motivation for this is threefold. First, it allows one to better capture the persistence inherent in implied volatility indices whilst simultaneously reducing the computational burden by reducing the number of coefficients to estimate.

with leading volatility indices that proxy future volatility of the: EUROSTOXX50; the CAC40; the DAX; and emerging markets. I also conduct the same analysis as I report here using volatility indices underlying single stocks; these are available upon request.

Table 1: **Descriptive Statistics and Correlations of Implied Volatility Indices**
Notes: This table reports descriptive statistics of implied volatility indices from 16 March 2011 to 3 September 2020. VXEFA is the CBOE implied volatility index for the EFA ETF; VIX is the CBOE's VIX index; OVX is the CBOE's oil volatility index; VXXLE is the CBOE's energy ETF implied volatility index; GVX is the CBOE's gold implied volatility index; and VXSLV is the CBOE's silver ETF implied volatility index. Mean and Median denote the sample mean and median respectively. Std. is the sample standard deviation, and Skew. and Kurt. are the sample skewness and kurtosis respectively. Contemporaneous Correlations in Panel B with a * indicate statistical significance at the 1% level using the serial correlation adjustment in Cryer and Chan (2008)

Panel A: Descriptive Statistics						
	VXEFA	VIX	OVX	VXXLE	GVX	VXSLV
Mean	18.71	17.21	36.09	24.66	16.91	29.53
Median	16.63	15.07	32.28	21.29	16.05	27.37
Std.	8.32	7.42	20.16	11.52	5.30	11.02
Skew.	2.11	3.05	5.14	3.54	1.32	1.62
Kurt.	9.24	17.36	45.12	22.12	5.71	6.74
Panel B: Contemporaneous Correlations						
	VXEFA	VIX	OVX	VXXLE	GVX	VXSLV
VXEFA	–					
VIX	0.87*	–				
OVX	0.53*	0.71*	–			
VXXLE	0.77*	0.94*	0.84*	–		
GVX	0.76*	0.66*	0.45*	0.64*	–	
VXSLV	0.74*	0.63*	0.40*	0.61*	0.87*	–

Second, the HAR structure reflects differing reaction times of various market participants to the arrival of news as well as allowing one to relate volatility patterns over longer intervals to those over shorter intervals. This is particularly relevant in the case of short-term agents who may use longer-term information within implied volatility to influence trading behaviour (Corsi, 2009). Third in using Chan (2020), I can impose certain Kronecker structures in order to model non-Gaussian, heteroskedastic and serially dependent error terms that explicitly accounts for parameter and estimation uncertainty conditional on the data.

Formally, let y_τ be an $n \times 1$ vector of variables that we observe over $\tau = 1, \ldots, T$ periods with the following vector HAR (VHAR) specification

$$y_\tau = b_0 + \mathbf{B}_1 y_{\tau-1} + \mathbf{B}_5 y_{(\tau-1|t-5)} + \mathbf{B}_{22} y_{(\tau-1|t-22)} + u_\tau$$

where b_0 is a vector of constants, $\mathbf{B}$'s are $n \times n$ coefficient matrices that capture autoregressive interactions at daily, weekly and monthly intervals respectively. The lagged vectors of y_τ are $y_{(\tau-1|t-k)} = \frac{1}{k}\sum_{i=1}^{k} y_{\tau-j}$ that contains three volatility components pertaining to short-, medium-, and long-term investment horizons.

It is easy to see that we can write the VHAR model as a VAR(22) with zero restrictions on all autoregressive matrices apart from those at the first, fifth and twenty-second lags of y_τ. Therefore we have the following VAR(3) structure.

$$y_\tau = a_0 + \sum_{i=1}^{3} \mathbf{A}_i y_{\tau-i} + u_\tau$$

where a_0 is a vector of constants, $\mathbf{A}_i$, $i = \{1, 5, 22\}$ are $n \times n$ coefficient matrices that we associate to the first, fifth and twenty-second lags of $y_\tau^{'}$ respectively. Now let $x_t^{'} = (1, y_{\tau-1}^{'}, y_{\tau-5}^{'}, y_{\tau-22}^{'})$ be a $k \times 1$ vector of a constant and lags with $k = 1 + np$. Stacking the observations gives us

$$\mathbf{Y} = \mathbf{X}\mathbf{A} + \mathbf{U} \tag{1}$$

with $\mathbf{A} = (a_0, \mathbf{A}_1, \mathbf{A}_5, \mathbf{A}_{22})'$ is of dimension $k \times n$ and the matrices $\mathbf{Y}$, $\mathbf{X}$, and $\mathbf{U}$ are of dimensions $T \times n$, $T \times k$ and $T \times n$ respectively. The $n \times n$ covariance matrix of the VAR model is $\mathbf{\Sigma}$. Typically, $\mathbf{U}$ is $\text{vec}(\mathbf{U}) \backsim \text{N}(\mathbf{0}, \mathbf{\Sigma} \otimes \mathbf{I_T})$ where $\text{vec}(\mathbf{U})$ stacks the columns of $\mathbf{U}$ and $\mathbf{I_T}$ is a T-dimensional identity matrix, $\otimes$ denotes the Kronecker product.

However, replacing $\mathbf{I_T}$ with a $T \times T$ covariance matrix $\mathbf{\Omega}$ allows one to model the cross-sectional and serial covariance structures of $\mathbf{Y}$ separately. Specifically, $\mathbf{\Sigma}$ governs the cross-sectional covariance structure and $\mathbf{\Omega}$ governs the serial covariance structure. Formally,

$$\text{vec}(\mathbf{U}) \backsim \text{N}\left(\mathbf{0}, \mathbf{\Sigma} \otimes \mathbf{\Omega}\right) \tag{2}$$

The model in (1) and (2) nests a wide variety of specifications by choosing different covariance structures for $\mathbf{\Omega}$. This paper uses seven models with different covariance structures including: i) a Bayesian VHAR with innovations following a $t-$distribution (BVHAR-t) (e.g. Chiu et al., 2017; Clark and Ravazzolo, 2015); ii) a Bayesian VAR with a common stochastic volatility component (BVHAR-CSV) (e.g. Carriero et al., 2016; Mumtaz and Theodoridis, 2018); iii) a Bayesian VHAR with MA(1) errors (BVHAR-MA(1)) (e.g. Dimitrakopoulos and Kolossiatis, 2020); iv) a Bayesian VHAR with $t-$errors and a common stochastic volatility component (BVHAR-t-CSV); v) a Bayesian VHAR with MA(1) errors following a $t-$distribution (BVHAR-t-MA(1)); vi) a Bayesian VHAR with MA(1) errors and a common stochastic volatility component (BVHAR-CSV-MA(1)); and vii) a Bayesian VHAR with MA(1) errors following a $t-$distribution and a common stochastic volatility component (BVHAR-t-CSV-MA(1)). The Appendix provides an outline of different covariance structures and estimation algorithms.

The priors for parameters common to all models are set to be the same and follow Chan (2020) and Carriero et al. (2016). Further details regarding priors and posterior simulation algorithms are in the Online Appendix. I estimate each of the 7 models using a 252-day rolling window that moves forward through the sample 1 day at a time. This results in 2199 windows per model with a sample spanning 3 April, 2012 to September 3, 2020. At each window I generate 6,000 simulations from the posterior and discard the first 1,000 as burn-in. I obtain the posterior distribution of parameters and network measures from each model on a Desktop PC with a 6-core Intel i7-8700K CPU @3.70GHz with 64GB RAM. Computation time varies with specification of the error covariance structure, but take approximately 15.5 to 21 hours per model. All computations are in MATLAB version 2021a and utilise MATLAB's statistical package. I adapt computer code from both Joshua Chan's website for estimation purposes, and from my Github page to obtain network measures; replication code for this paper is also available from my Github page.

2.2.1 Tracking Network Connections in VHARs with a more General Error Covariance Structure

Diebold and Yılmaz (2014) and Baruník and Křehlík (2018) provide network measures using VAR models that stem from the forecast error variance decomposition matrix. Importantly, these measures stem from generalised forecast error variance decompositions (GFEVDs)

(Pesaran and Shin, 1998) that hold under the assumption of time series with no serial correlation (i.e. $\mathbf{\Omega} = \mathbf{I}_T$). Here I show how these measures modify for a more general structure on $\mathbf{\Omega}$. Prior to presenting the measures I outline notation. For a matrix $\mathbf{B}$, $(\mathbf{B})_{j,k}$ denotes the jth row and kth column. $(\mathbf{B})_j$ denotes the full jth row of $\mathbf{B}$, and $\sum \mathbf{B}$ denotes the sum of all elements in $\mathbf{B}$.

Note that we are able to write y_τ as $\mathbf{A}(L)y_\tau = u_\tau$ with $\mathbf{A}(L) = [\mathbf{I}_n - \mathbf{A}_1 L_1 - \cdots - \mathbf{A}_p L_p]$ being an $n \times n$ matrix lag-polynomial. Now assuming the roots of the VAR polynomial lie outside the unit circle we obtain the MA(∞) representation, $y_\tau = \mathbf{\Psi}(L)u_\tau$ where one can compute $\mathbf{\Psi}(L)$ from $\mathbf{A}(L) = [\mathbf{\Psi}(L)]^{-1}$ recursively. The specifications I outline in the previous section result in scaling the elements the VAR's covariance matrix $\mathbf{\Sigma}$ by some quantity $\mathbf{c}$ that depends on the structure of $\mathbf{\Omega}$ such that $u_\tau \backsim \mathrm{N}(0, \mathbf{c\Sigma})$[4]. For simplicity of notation let $\tilde{\mathbf{\Sigma}} = \mathbf{c\Sigma}$.

As an example, suppose that the errors follow an MA(1) process, with $u_\tau = \epsilon_t + \psi\epsilon_{t-1}$, $\epsilon_t \backsim \mathrm{N}(0, \mathbf{\Sigma})$[5]. It is easy to see that $\mathrm{var}(u_\tau) = (1+\psi^2)\mathbf{\Sigma}$. Now defining $\tilde{\mathbf{\Sigma}} = \mathbf{c\Sigma}$, with $\mathbf{c} = (1+\psi^2)$ in this case, the covariance matrix of the forecast error conditional on information at time $\tau - 1$ is

$$\Gamma_H = \sum_{h=0}^{H} \mathbf{\Psi}_h \tilde{\mathbf{\Sigma}} \mathbf{\Psi}_h^{'}.$$

Now defining the covariance matrix of the forecast error conditional on the knowledge of the current and future expected shocks to the jth equations. The conditional forecast error is

$$\varsigma_\tau^k(H) = \sum_{h=0}^{H} \mathbf{\Psi}_h \left[u_{\tau+H-h} - \mathbb{E}\left(u_{\tau+H-h} | u_{k,\tau+H-h} \right) \right].$$

Which under a Normal distribution is

$$\varsigma_\tau^k(H) = \sum_{h=0}^{H} \mathbf{\Psi}_h \left[u_{\tau+H-h} - \tilde{\sigma}_{kk}^{-1} \left(\tilde{\mathbf{\Sigma}} \right)_k u_{k,\tau+H-h} \right].$$

Then, the covariance matrix is

$$\Gamma_H^k = \sum_{h=0}^{H} \mathbf{\Psi}_h \tilde{\mathbf{\Sigma}} \mathbf{\Psi}_h^{'} - \tilde{\sigma}_{kk}^{-1} \sum_{h=0}^{H} \mathbf{\Psi}_h \left(\tilde{\mathbf{\Sigma}} \right)_k \left(\tilde{\mathbf{\Sigma}} \right)_k^{'} \mathbf{\Psi}_h^{'}.$$

[4]It is important to note here that that the above also holds when errors follow a multivariate-t distribution. Ding (2016) shows that this differs from the multivariate Normal distribution only by a random scaling factor, that I define in this paper as λ_t; more details are in the Appendix, as well as expressions for scaling factors using more elaborate error covariance structures.

[5]The Appendix provides details on the scaling factors for the seven different error covariance structures I outline above.

The unscaled H-step ahead forecast error variance of the jth variable with respect to the innovation in the kth variable is

$$\Delta_{(j)kH} = \left(\Gamma_H - \Gamma_H^k\right) = \tilde{\sigma}_{kk}^{-1} \sum_{h=0}^{H} \left(\left(\mathbf{\Psi}_h \tilde{\mathbf{\Sigma}}\right)_{j,k}\right)^2.$$

Scaling provides us with the expression for the Pesaran and Shin (1998) GFEVD:

$$\left(\mathbf{\Theta}_H\right)_{j,k} = \frac{\tilde{\sigma}_{kk}^{-1} \sum_{h=0}^{H} \left(\left(\mathbf{\Psi}_h \tilde{\mathbf{\Sigma}}\right)_{j,k}\right)^2}{\sum_{h=0}^{H} \left(\mathbf{\Psi}_h \tilde{\mathbf{\Sigma}} \mathbf{\Psi}_h^{'}\right)_{j,j}}$$

where $\tilde{\sigma}_{kk} = (\tilde{\mathbf{\Sigma}})_{k,k}$, $\mathbf{\Psi}_h$ are the $n \times n$ matrix of MA coefficients at lag h. $\left(\mathbf{\Theta}_H\right)_{j,k}$ denotes the contribution of variable k to the forecast error variance of variable j at horizon H[6]. It is clear the only difference here is replacing $\mathbf{\Sigma}$ with $\tilde{\mathbf{\Sigma}} = \mathbf{c}\mathbf{\Sigma}$ where c scales the VHAR's covariance matrix in a manner that accounts for the structure of the serial covariance structure $\mathbf{\Omega}$.

Therefore, the results in Baruník and Křehlík (2018) also hold here using $\tilde{\mathbf{\Sigma}}$. The power spectrum, $\mathbf{S}_y(\omega)$, allows us to examine networks that form over different horizons. Let ω denote the frequency component and $i = \sqrt{-1}$, then the spectral density is the Fourier transform of the infinite MA representation:

$$\mathbf{S}_y(\omega) = \sum_{h=-\infty}^{\infty} \mathbb{E}\left[y_\tau y_{\tau-h}\right] e^{-i\omega h} = \mathbf{\Psi}(e^{-i\omega}) \tilde{\mathbf{\Sigma}} \mathbf{\Psi}'(e^{+i\omega})$$

The above means one can define the generalised causation spectrum over frequency components $\omega \in [-\pi, \pi]$ and the weighting function to obtain a spectral representation of variance decompositions from variable j to k in y_τ. This proves the existence of horizon specific network measures. The generalised causation spectrum and weighting function are

$$\begin{aligned}
\left(\mathsf{f}(\omega)\right)_{j,k} &= \frac{\tilde{\sigma}_{kk}^{-1} |\left(\mathbf{\Psi}(e^{-i\omega})\right)_{j,k}|^2}{\left(\mathbf{\Psi}(e^{-i\omega}) \tilde{\mathbf{\Sigma}} \mathbf{\Psi}'(e^{-i\omega})\right)_{j,j}} \\
\mathbf{\Gamma}_j(\omega) &= \frac{\left(\mathbf{\Psi}(e^{-i\omega}) \tilde{\mathbf{\Sigma}} \mathbf{\Psi}'(e^{-i\omega})\right)_{j,j}}{\frac{1}{2\pi} \int_{-\pi}^{\pi} \left(\mathbf{\Psi}(e^{-i\lambda}) \tilde{\mathbf{\Sigma}} \mathbf{\Psi}'(e^{-i\lambda})\right)_{j,j} \mathrm{d}\lambda}
\end{aligned}$$

and allow us to define the spectral representation of the variance decomposition matrix:

[6]Note for estimation purposes, the number of lags to approximate the infinite MA representation is set to H=100 horizons; increasing this to $H = \{150, 200, 250\}$ does not change the results and conclusions I present in what follows.

$$(\mathbf{\Theta}_H)_{j,k} = \frac{1}{2\pi}\int_{-\pi}^{\pi} \mathbf{\Gamma}_j(\omega)\left(\mathsf{f}(\omega)\right)_{j,k}\mathrm{d}\omega, \quad H \to \infty$$

which reconstructs the H horizon variance decomposition matrix.

Instead of examining each component on the spectrum, ω, I partition the entire spectrum into bands. They constitute the short-, medium-, and long-horizons respectively. Therefore one can modify the above by allowing band $d = (a, b) \ : a, b \in [-\pi, \pi]\,,\ a < b$ such that

$$(\mathbf{\Theta}_d)_{j,k} = \frac{1}{2\pi}\int_{d} \mathbf{\Gamma}_j(\omega)\left(\mathsf{f}(\omega)\right)_{j,k}\mathrm{d}\omega$$

Note that summing over all horizons delivers $(\mathbf{\Theta}_H)_{j,k}$.

Then, scaling $\mathbf{\Theta}_d$ such that $(\tilde{\mathbf{\Theta}}_d)_{j,k} = (\mathbf{\Theta}_d)_{j,k}/\sum_k(\mathbf{\Theta}_H)_{j,k}$ delivers the adjacency matrix, $\tilde{\mathbf{\Theta}}_d$ over horizon band d. Manipulations allows one to define connectedness at horizon d as

$$\mathcal{C}_d = 100 \times \left(\frac{\sum \tilde{\mathbf{\Theta}}_d}{\sum \tilde{\mathbf{\Theta}}_H} - \frac{\mathrm{Tr}\{\tilde{\mathbf{\Theta}}_d\}}{\sum \tilde{\mathbf{\Theta}}_H}\right) \tag{3}$$

Equation 3 measures the contribution of forecast error variance attributable to all shocks within the system, minus own shocks. This measure infers system-wide connectedness at a horizon of interest.

Further manipulations of $\tilde{\mathbf{\Theta}}_d$ allow one to define directional connections at each horizon of interest. Specifically TO and FROM connections are given by

$$\mathcal{C}_{d,j\to\bullet} = 100 \times \frac{\sum_{k=1,k\neq j}^{n}(\tilde{\mathbf{\Theta}}_d)_{k,j}}{\sum_{j,k=1}^{n}(\tilde{\mathbf{\Theta}}_H)_{k,j}} \tag{4}$$

$$\mathcal{C}_{d,j\leftarrow\bullet} = 100 \times \frac{\sum_{k=1,k\neq j}^{n}(\tilde{\mathbf{\Theta}}_d)_{j,k}}{\sum_{j,k=1}^{n}(\tilde{\mathbf{\Theta}}_H)_{j,k}} \tag{5}$$

respectively. Equation (4) is precisely interpretable as in-degrees as the network literature names it. This indicates how much of asset j contributes to the variance of other variables within the system at horizon d. Equation (5) is precisely interpretable as out-degrees as the network literature names it. This indicates how much of asset j's variance is due to shocks from other variables within the system at horizon d. Taking the difference between (4) and (5) provides net directional connections of variable j at horizon d.

$$\mathcal{C}_{d,j}^{\text{NET}} = \mathcal{C}_{d,j\to\bullet} - \mathcal{C}_{d,j\leftarrow\bullet} \tag{6}$$

If the value is positive then the variable transmits its shocks to other variables after accounting for shock reception. Likewise, if the value is negative, the variable is a shock receptor, after accounting for its shock transmission capacity[7].

3 Results

3.1 Model Selection

Table 2: **Log Marginal Likelihoods**
Notes: This table reports the average log marginal likelihoods (Ave. ML) and median log marginal likelihoods (Med. ML) from each Bayesian Vector Heterogeneous Autoregressive (BVHAR) model using a rolling 252 day window throughout the sample 16 March, 2011 to 3 September, 2020. The suffixes: t- refer to errors following a $t-$distribution; CSV refer to a common stochastic volatility component; MA(1) refers to residuals following an MA(1) process. Combinations of these suffixes combine these residual covariance structures. Bold font indicates the highest Ave. ML and the highest Med. ML out of the seven models

	Ave. ML	Med. ML
BVHAR-t	-2261.47	-2196.18
BVHAR-CSV	-2217.93	-2151.13
BVHAR-MA(1)	-2430.39	-2342.96
BVHAR-t-CSV	-2193.86	-2133.04
BVHAR-t-MA(1)	-2249.72	-2188.78
BVHAR-CSV-MA(1)	-2211.78	-2149.42
BVHAR-t-CSV-MA(1)	**-2193.61**	**-2132.85**

Before reporting my main results, I conduct a model comparison exercise using log marginal likelihoods (log-MLs) following Chib (1995) and Chan (2020). These log-MLs have a built in penalty for complexity, which means it may not be the case that the most general covariance structure delivers the highest log-ML. Table 2 reports the average, and median, log-MLs for the seven BVHAR models. The higher the value, the better the model fit. It is clear that accounting for heavy tails, stochastic volatility, and serial correlation are empirically important. The model delivering the best log-ML is the BVHAR-t-CSV-MA(1).

[7]Note that further manipulations of $\tilde{\mathbf{\Theta}}_d$ permit one to define pairwise connections in an analogous manner to Equations (4), (5), and (6).

Therefore, I proceed by using network measures that stem from the BVHAR-t-CSV-MA(1) model[8].

3.2 Implied Volatility Network Connections

3.2.1 Horizon Specific Connectedness

I now examine horizon specific network connections using $\mathcal{C}_d$ in equation (3), with $d \in \{\text{Short, Medium, Long}\}$. Short proxies network connections from 1-day to 1-week, Medium proxies network connections from 1-week to 1-month, and Long proxies connections at horizons greater than 1-month. Figure 1 reports the posterior median and one-standard deviation percentiles of the posterior distribution of horizon specific network connectedness measures. The top panel shows short-horizon network connectedness, and the middle and bottom panels show medium- and long-horizon network connectedness respectively.

Three main points emerge from Figure 1. First, connectedness is strongest over the long-run with a gradual decline over the sample before rising in early 2020 amidst the global pandemic and ensuing recession. Next, error bands show that connectedness over the long-run is statistically different to short- and medium-horizon connectedness. Finally, surges in $\mathcal{C}_d$, $d \in \{\text{Short, Medium}\}$ occur at the time of transient stock market events throughout the sample. Specifically, during 2015 and 2018 short-run and medium-run network connectedness surge. Consistent with the sell off within Chinese stock markets, the fall in US stock markets that also links with drops in commodity prices, and the structural break in currencies that accompanies international stock market falls at the end of 2018.

Overall, these measures show that there are substantial differences in network connectedness at different horizons which also exhibit heterogeneous time dynamics. Transient events that occur in markets throughout the sample results in rising network connectedness over the short- and medium-run; for example, the downturn in financial and commodity markets in 2015. However, events that people expect to cause persistent uncertainty for a significant amount of time, such as the COVID-19 outbreak and the ensuing turmoil in energy markets (Baumeister et al., 2020) that causes economic recession in 2020, cause long-term network connections to intensify as evident from late 2019[9].

[8]The Online Appendix reports Bayes factors for log-MLs throughout the estimation sample. I benchmark the BVHAR-t-CSV-MA(1) against the other six models in Table 2. Overall, the plots reveal that the BVHAR-t-CSV-MA(1) is preferable relative to each of the simpler models. Note also that taking logs of the implied volatility indices prior to model estimation delivers the same results as I report in Table 2; these results are available on request.

[9]The Appendix plots the posterior median network measures from all 7 models I consider and shows differences are negligible.

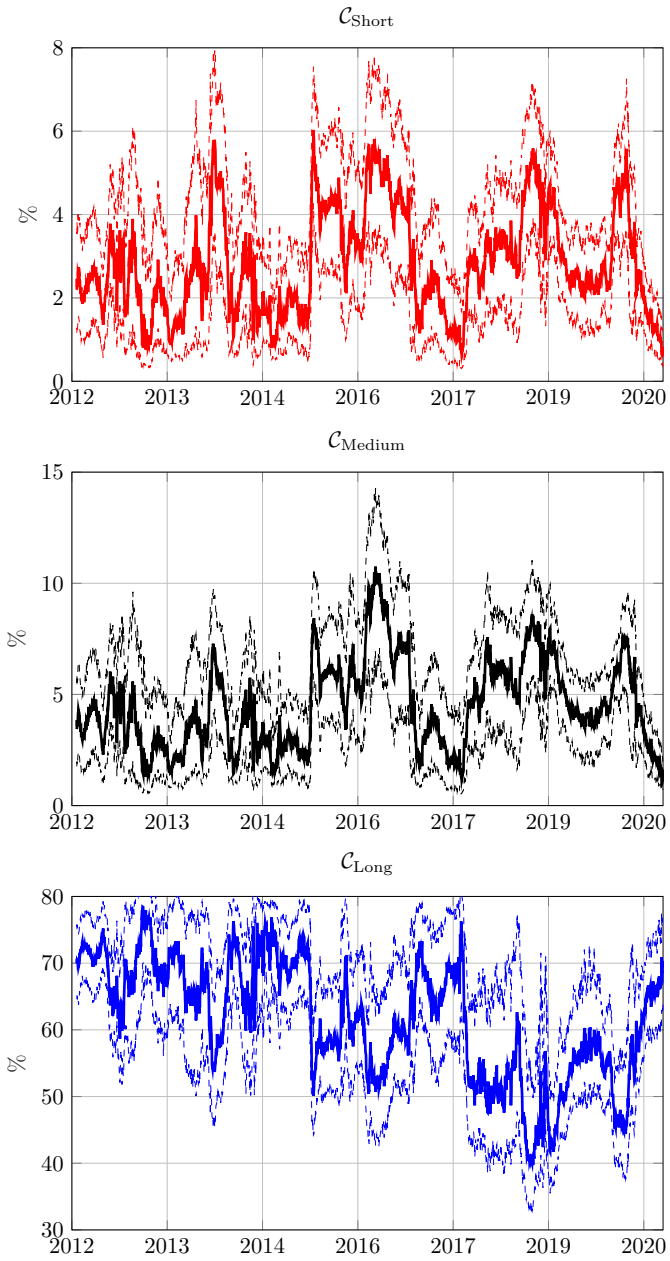

Figure 1: **Horizon Specific Network Connectedness Measures from 3 April, 2012 to 3 September, 2020.**

Notes: This figure plots the posterior median, and one-standard deviation percentiles of the posterior distribution of horizon specific network connectedness measures, $\mathcal{C}_d$. The top panel shows network connectedness at short horizons (i.e. 1-day to 1-week); the middle panel shows network connectedness at medium horizons (i.e. 1-week to 1-month); and the bottom panel shows network connectedness at long horizons (i.e. horizons greater than 1-month).

3.2.2 Net-Directional Connections

I now examine net-directional connections that form among stock and commodity markets at different horizons. I plot the posterior median and one-standard deviation percentiles of the posterior distributions for $\mathcal{C}^{\text{NET}}_{d,j}$, $d \in \{\text{Short, Medium, Long}\}$ in Figures 2, 3, 4 respectively. Note that variables $j = \{\text{VXEFA, VIX, OVX, VXXLE, GVX, VXSLV}\}$. Positive (Negative) values of $\mathcal{C}^{\text{NET}}_{d,j}$ indicate that variable j transmits (receives) shocks to (from) other variables within the system after accounting for shock reception (transmission) capacity.

In Figure 2 there are various noteworthy points. From 2013 to 2017, commodity markets transmit shocks throughout the system. The energy commodities–namely OVX and VXXLE–are prominent transmitters from 2014–2017, whereas precious metals transmit from 2013–2014. During this period, equity markets act as shock receptors. From mid-2017 to late 2019 VIX, OVX and VXXLE are dominant transmitters and VXEFA, GVX and VXSLV are shock receptors. These patterns are more pronounced in 3. It makes sense for this to occur as Figure 1 reveals one cannot distinguish statistically from $\mathcal{C}_{\text{Short}}$ and $\mathcal{C}_{\text{Medium}}$.

Figure 4 plots long horizon net-directional connections. Looking at equity markets, net-directional connections of VXEFA and VIX exhibit similar time dynamics up until 2019. Then, net-directional connections of the VIX remain positive (although error bands include zero), while the VXEFA becomes a (statistically credible) shock receptor. Considering energy markets, OVX and VXXLE net-directional connections oppose one another across various dates throughout the sample. In particular, we can see OVX (VXXLE) acts as a shock receptor (transmitter) during 2012–2013; the same phenomenon occurs in 2014–2016, and again at the end of the sample[10]. Turning to precious metals, over long-horizons net-directional connections appear to fluctuate around zero throughout the sample indicating that from a directional perspective shock transmission and reception capacity cancel one another out.

[10]This may occur for two reasons. First, the long-run cancels out any transient shocks that bear no permanent influence on price changes. Second, it could be down to what the volatility indices track. OVX tracks a fund holding largely US oil futures. VXXLE tracks an ETF following an index of US energy firms. Spot prices of commodities influence the performance of energy firms, and in turn help determine future prices. Indeed, the price of crude oil fell during 2014–2016 due to a decline in economic activity, along with falling demand in crude oil. At the end of the sample one can associate this observation with the COVID-19 outbreak and the Russia-Saudi oil price war (Baumeister et al., 2020).

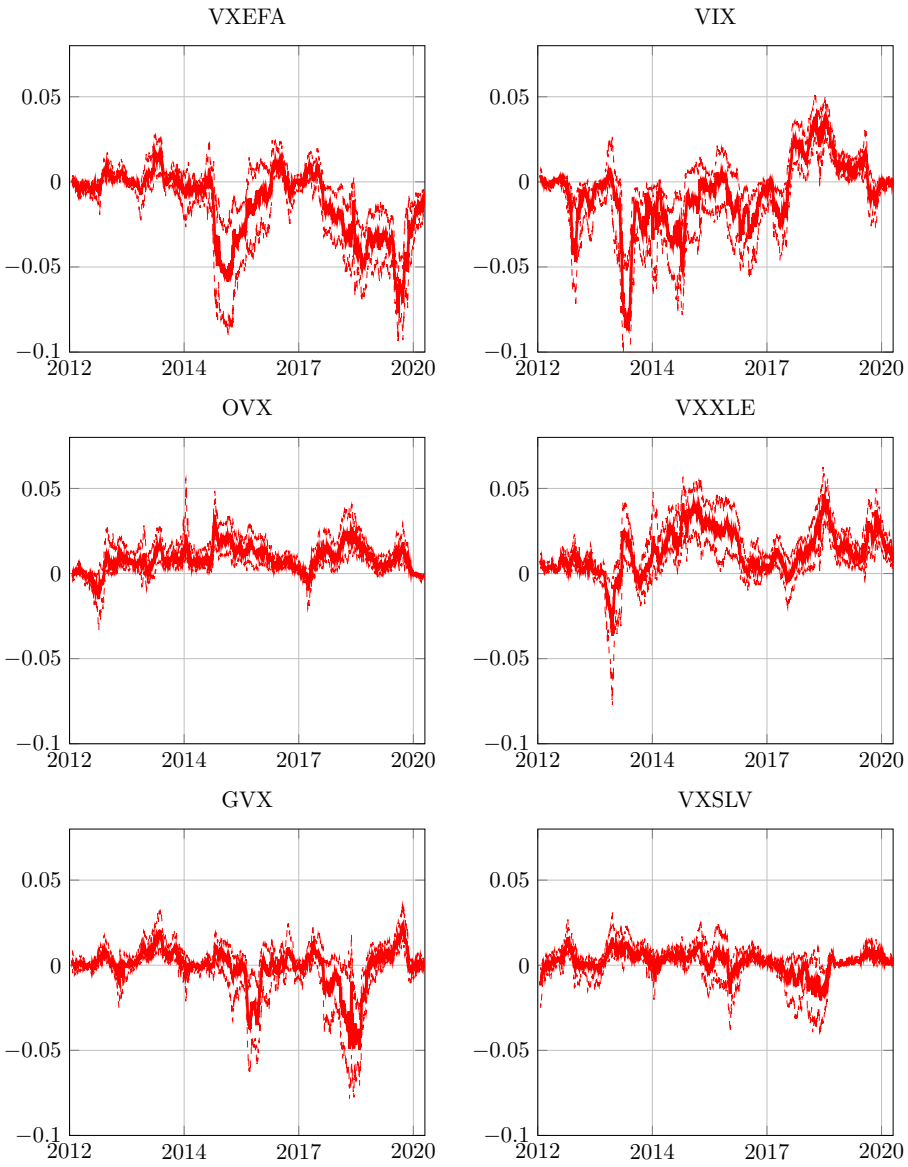

Figure 2: **Short-Horizon Net-directional Connections from 3 April, 2012 to 3 September, 2020.**

Notes: This figure plots the posterior median, and one-standard deviation percentiles of the posterior distribution of short-horizon (i.e. 1-day to 1-week) network connectedness measures, $\mathcal{C}_{d,j}^{\text{NET}}$. VXEFA is the EFA ETF volatility index; VIX is the VIX index; OVX is the crude oil ETF volatility index; VXXLE is the XLE ETF volatility index and GVX and VXSLV are the volatility indices of gold and silver respectively. Positive (Negative) values indicate that the variable acts as a shock transmitter (receptor) after accounting for its shock reception (transmission) capacity).

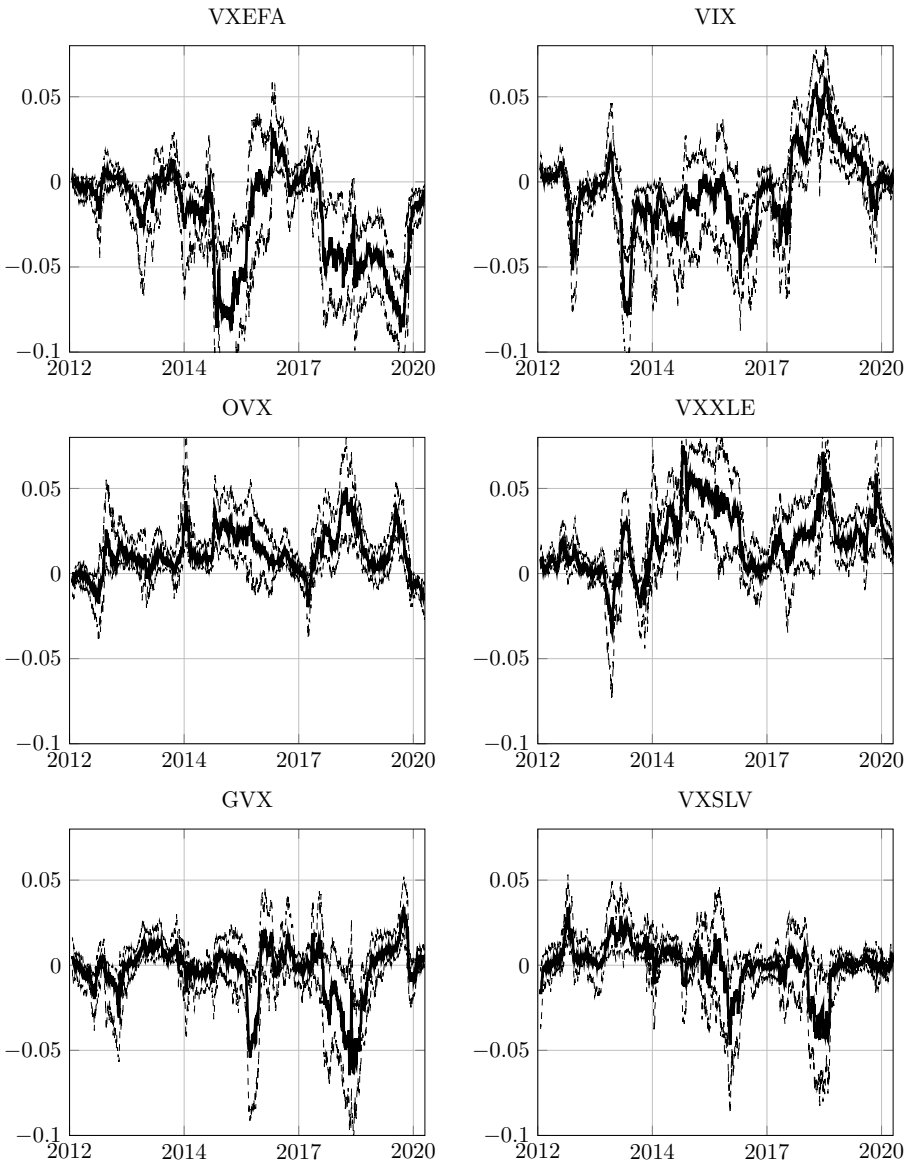

Figure 3: **Medium-Horizon Net-directional Connections from 3 April, 2012 to 3 September, 2020.**

Notes: This figure plots the posterior median, and one-standard deviation percentiles of the posterior distribution of medium-horizon (i.e. 1-week to 1-month) network connectedness measures, $\mathcal{C}_{d,j}^{\text{NET}}$. VXEFA is the EFA ETF volatility index; VIX is the VIX index; OVX is the crude oil ETF volatility index; VXXLE is the XLE ETF volatility index and GVX and VXSLV are the volatility indices of gold and silver respectively. Positive (Negative) values indicate that the variable acts as a shock transmitter (receptor) after accounting for its shock reception (transmission) capacity).

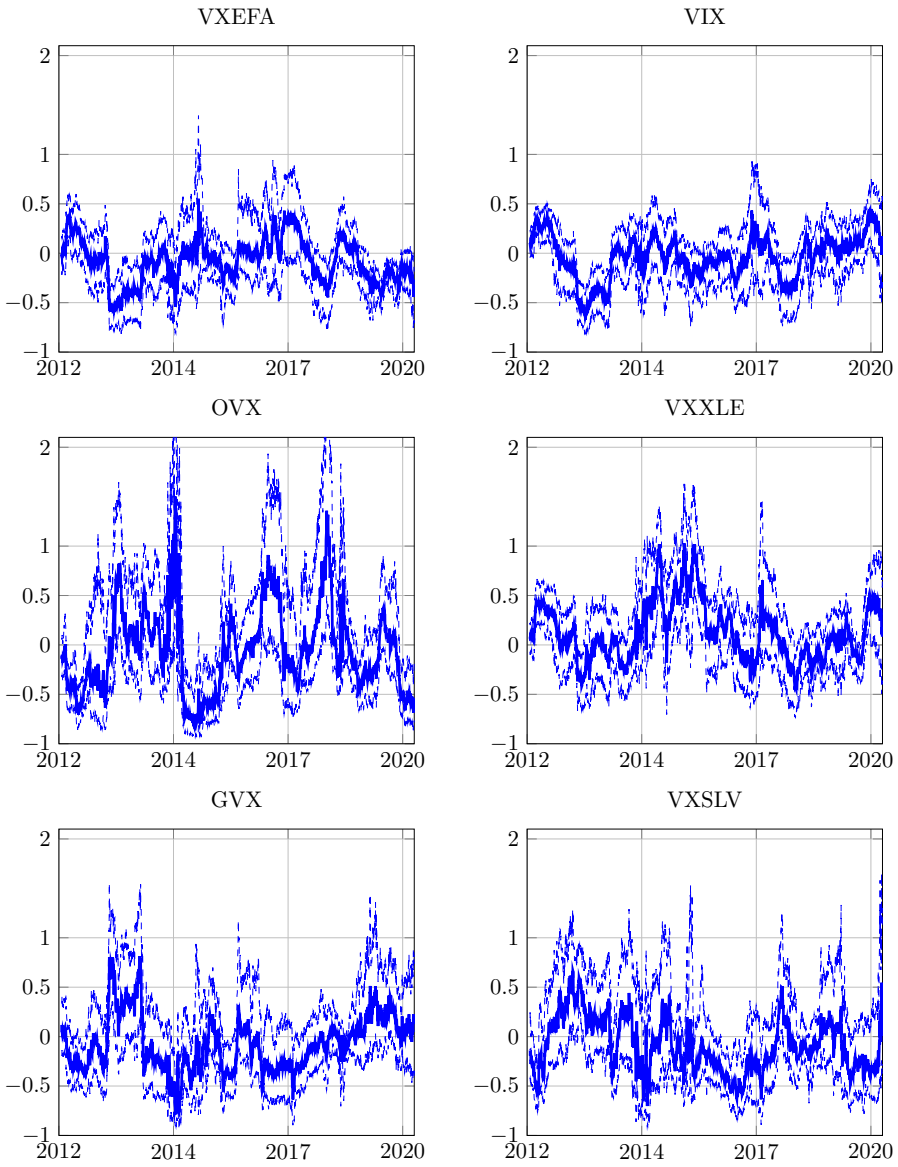

Figure 4: **Long-Horizon Net-directional Connections from 3 April, 2012 to 3 September, 2020.**

Notes: This figure plots the posterior median, and one-standard deviation percentiles of the posterior distribution of long-horizon (i.e. horizons > 1-month) network connectedness measures, $\mathcal{C}_{d,j}^{\text{NET}}$. VXEFA is the EFA ETF volatility index; VIX is the VIX index; OVX is the crude oil ETF volatility index; VXXLE is the XLE ETF volatility index and GVX and VXSLV are the volatility indices of gold and silver respectively. Positive (Negative) values indicate that the variable acts as a shock transmitter (receptor) after accounting for its shock reception (transmission) capacity).

3.3 Do Horizon Specific Network Connectedness Measures Help Predict Future Returns of the Underlying Assets?

I now investigate whether network connectedness measures on implied volatility indices have out-of-sample forecasting power for the underlying asset returns for each volatility index.

I take the daily prices of the iShares MSCI EAFE ETF fund (EFA), the SDPR S&P500 ETF trust (SPY), the United States Oil Fund ETF (USO), the Energy Select sector SPDR ETF (XLE), the SPDF Gold Shares ETF (GLD), and the iShares silver ETF (SLV) funds from Thomson Reuters Datastream and compute returns as 100 times the logarithmic ratio of prices on consecutive days. I also take rolling two-day averages for the first difference of each implied volatility index (as in, e.g. Yang and Zhou, 2017), to control for the effects of volatility. I also compute the daily changes in each of the horizon d connectedness measures. I create four datasets, one using the six underlying asset returns and changes in implied volatility indices, which I refer to as the baseline dataset. The remaining three datasets add one of the respective network connectedness measures $\mathcal{C}_d$, $d \in \{\text{Short, Medium, Long}\}$, to the underlying asset returns and changes in volatilities.

I consider three Bayesian VAR models, namely BVAR-CSV, BVAR-CSV-t, and a BVAR-CSV-t-MA(1), to generate forecasts of the underlying asset returns from the four datasets. The forecasting exercise is as follows: For each dataset and each model, I set the lag length p=5 and estimate models using a 252 day window that rolls throughout the sample adding one-day at time. At each recursion I generate 6,000 simulations of the model, discard the initial 1,000 as burn-in and use the remaining 5,000 draws to produce iterative forecasts up to q=10 periods (days) ahead. The recursive forecasting sample spans 14 March 2013, to August 21,2020.

I consider both point and density forecasts using mean absolute forecast errors (MAFEs) and log predictive scores (LS). The MAFE for variable j at horizon q and the LS for the jth variable at horizon q are

$$\text{MAFE}_{\tau,q,j} = \frac{1}{\tau - q - s + 1}\left|\sum_{i=s}^{\tau-q}\left(y^o_{\tau+q,j} - \mathbb{E}\left[y_{t+q,j}|y_1,\ldots,y_i\right]\right)\right| \tag{7}$$

$$\text{LS}_{\tau,q,j} = -\sum_{i=s}^{\tau-q}\log \text{p}\left(y_{i+q,j} = y^o_{\tau+q,j}|y_1,\ldots,y_i\right), \tag{8}$$

respectively. In (7) and (8), $y^o_{\tau+q,j}$ is the actual value of the jth variable at time $\tau+q$, s denotes the start of the forecasting period. In both cases, lower values indicate better forecast performance.

The forecasting methodology is as follows: First, I compute Hansen et al. (2011) 95% model confidence sets to arrive at the set of models containing the best forecasts. For every model, I gather the loss functions–either MAFEs or LSs–for variables of interest and forecast horizons in a $q \times n$ matrix, compute the Frobenius norm of these matrices and then run the Hansen et al. (2011) model confidence set procedure. Then given the model confidence set, I compare forecasts from models and datasets containing network connectedness measures against those with no network measures in the dataset. I assess the statistical significance of forecasts gains from models using network connectedness relative to those without using the Giacomini and White (2006) tests of equal conditional predictive ability.

Table 3: **Hansen et al. (2011) Model Confidence Sets from Frobenius Norms of Loss functions**

Notes: This table reports the Frobenius norms from mean absolute forecast errors and log scores for underlying asset returns within each dataset across forecast horizons $q = \{1, 5, 10\}$. Values highlighted in bold font indicate the dataset and models belonging to the 95% Hansen et al. (2011) model confidence sets.

Baseline Dataset					
Mean Absolute Forecast Error			**Log Scores**		
BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)	BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)
5.550	5.550	5.562	**8.021**	8.052	**7.981**
$\mathcal{C}_d$, d =**Short**					
BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)	BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)
5.544	**5.546**	5.557	**7.991**	8.162	8.090
$\mathcal{C}_d$, d =**Medium**					
BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)	BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)
5.544	**5.546**	5.556	**7.991**	8.161	8.091
$\mathcal{C}_d$, d =**Long**					
BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)	BVAR-CSV	BVAR-CSV-t	BVAR-CSV-t-MA(1)
5.544	**5.546**	5.557	**7.981**	8.154	8.083

Table 3 shows the Frobenius norms stemming from MAFEs and LS for each of competing forecasting models that I fit to the four datasets. Values in bold font belong to the 95% model confidence set (Hansen et al., 2011). Looking at point forecasts, BVAR-CSV and BVAR-CSV-t models fit to datasets containing a network connectedness measure all appear in the 95% model confidence set. It is difficult to distinguish between datasets as they all have very similar Frobenius norm values, but the BVAR-CSV looks to be the best performing model. Turning to density forecasts, the BVAR-CSV models fit to all datasets and the

BVAR-CSV-t-MA(1) model fit to the baseline dataset make up the model confidence set; the lowest Frobenius norm is the BVAR-CSV fit to the dataset containing long-horizon network connectedness with a value of 7.981.

As the model confidence set from point forecasts does not include the baseline dataset, I focus on the model confidence set from density forecasts. I benchmark point and density forecasts from the datasets containing a network connectedness measure, that all stem from BVAR-CSV models, against forecasts from: i) the baseline dataset using the BVAR-CSV model in Table 4; and ii) the baseline dataset using the BVAR-CSV-t-MA(1) model in Table 5. Within each Table, Panel A shows MAFEs and Panel B shows LSs. I express forecast metrics for the baseline dataset as in (7) and (8), and I express metrics from datasets containing network measures relative to the baseline. Negative values in both cases indicate that the forecast for the jth underlying asset return at horizon q is better when using a dataset containing a network connectedness measure. Numbers in parentheses are p-values from the Giacomini and White (2006) test of equal conditional predictive ability.

First considering Table 4, it is clear from Panel A that there are negligible differences in point forecasts of underlying asset returns when accounting a network connectedness at different horizons. We can see that gains are predominantly made at q=1 day ahead forecasts with statistically significant gains for XLE. There are also some statistically significant gains in forecasting SPY at q=5 day ahead forecasts. Turning to Panel B, we can see substantial benefits from density forecasts. There are considerable gains at 5 and 10 day ahead forecasts. At 5-day ahead horizons, forecasts of all underlying assets except GLD are significantly different from those stemming from the baseline dataset. Results for 10-day ahead forecasts largely correspond with the 5-day ahead forecasts, although the statistical significance of SLV disappears.

Next looking at Table 5, a similar story emerges. This shows that even allowing for a more flexible error covariance structure is not enough to outperform forecasts from a simpler model that include network connectedness within the information set. Overall, these results indicate that network connectedness provides: i) more accurate point forecasts 1-day ahead; and ii) more accurate density forecasts as the forecast horizon increases. These gains in density forecasts may occur as network connectedness permits a more precise estimation of the time-varying variances thus delivering lower LSs. This result highlights the importance of computing predictive probability distributions; a vital tool for market participants to make decisions (Geweke and Amisano, 2010). The implication here is that forecasters should track connections among implied volatility indices in order to predict the underlying asset returns. This is because asset managers whose performance relates to their decisions are willing to pay for even marginal improvements (Campbell and Thompson, 2008).

Table 4: **Point and Density Forecasts for Asset Returns I**

Notes: This table reports the $q = \{1, 5, 10\}$ day ahead mean absolute forecast errors and log scores of asset returns over the forecast sample 14 March, 2013 to 21 August, 2020 in Panels A and B respectively. Short, Medium, and Long add the daily change in $\mathcal{C}_d$ $d = \{\text{Short, Medium, Long}\}$ to the baseline dataset with forecast stemming from a BVAR-CSV model. MAFE and LS for these forecasts are relative to those from baseline dataset using the BVAR-CSV model. Negative values indicate better forecast performance from the dataset containing a network connectedness measure. Values in parentheses are p-values from the Giacomini and White (2006) test of equal conditional predictive ability. EFA is the the ishares MCFSI EFA ETF, SPY is the S&P500 ETF, USO is the US Oil Fund ETF, XLE is the SDPR energy selector ETF, GLD is the SPDR gold shares ETF, and SLV is the silver ETF.

	Panel A: MAFE							**Panel B: LS**						
Dataset, Model	Horizon	EFA	SPY	USO	XLE	GLD	SLV	Horizon	EFA	SPY	USO	XLE	GLD	SLV
Baseline, BVAR-CSV	q=1	0.705	0.654	1.605	1.101	0.680	1.083	q=1	2488.583	2329.955	4055.830	3288.391	2536.848	3453.075
	5	0.689	0.640	1.563	1.073	0.661	1.045	5	2502.321	2324.710	4063.232	3321.710	2464.410	3338.352
	10	0.688	0.640	1.566	1.073	0.659	1.046	10	2562.156	2373.192	4145.974	3399.004	2507.593	3375.774
$\mathcal{C}_d$ d=Short, BVAR-CSV	q=1	0.000	**-0.001**	**-0.005**	**-0.005**	**-0.002**	**-0.002**	q=1	5.792	2.068	**-5.571**	**-2.923**	3.218	5.351
		(0.86)	(0.33)	(0.10)	(0.02)	(0.50)	(0.30)		(0.48)	(0.86)	(0.52)	(0.74)	(0.62)	(0.52)
	5	0.000	**-0.001**	0.000	0.000	0.000	0.000	5	**-17.981**	**-20.378**	**-26.735**	**-21.039**	**-8.012**	**-11.735**
		(0.66)	(0.04)	(0.86)	(0.50)	(0.97)	(0.54)		(0.00)	(0.00)	(0.00)	(0.00)	(0.12)	(0.03)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	**-0.328**	**-11.537**	**-11.734**	**-1.137**	6.252	**-3.604**
		(0.68)	(0.46)	(0.58)	(0.16)	(0.73)	(0.54)		(0.01)	(0.02)	(0.00)	(0.01)	(0.17)	(0.22)
$\mathcal{C}_d$ d=Medium, BVAR-CSV	q=1	0.000	**-0.001**	**-0.007**	**-0.004**	**-0.002**	**-0.002**	q=1	9.167	3.312	**-7.461**	**-3.075**	6.580	7.182
		(0.84)	(0.76)	(0.06)	(0.06)	(0.37)	(0.59)		(0.14)	(0.72)	(0.37)	(0.58)	(0.34)	(0.30)
	5	0.000	**-0.001**	0.000	**-0.001**	0.000	0.000	5	**-17.558**	**-19.935**	**-27.592**	**-20.332**	**-6.040**	**-9.787**
		(0.78)	(0.05)	(0.78)	(0.62)	(0.94)	(0.94)		(0.00)	(0.00)	(0.00)	(0.00)	(0.19)	(0.08)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	1.130	**-11.418**	**-11.493**	**-0.893**	6.988	**-3.386**
		(0.89)	(0.32)	(0.96)	(0.15)	(0.97)	(0.28)		(0.01)	(0.04)	(0.00)	(0.01)	(0.16)	(0.20)
$\mathcal{C}_d$ d=Long, BVAR-CSV	q=1	**-0.001**	**-0.001**	**-0.006**	**-0.005**	**-0.001**	**-0.003**	q=1	**-3.681**	**-8.253**	**-11.816**	**-13.627**	1.013	4.455
		(0.48)	(0.35)	(0.12)	(0.01)	(0.47)	(0.38)		(0.76)	(0.21)	(0.12)	(0.08)	(0.91)	(0.55)
	5	0.000	**-0.001**	0.000	**-0.001**	0.000	0.000	5	**-20.533**	**-25.039**	**-26.851**	**-20.754**	**-7.271**	**-11.300**
		(0.84)	(0.25)	(1.00)	(0.34)	(0.87)	(0.94)		(0.00)	(0.00)	(0.00)	(0.00)	(0.14)	(0.04)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	0.209	**-12.617**	**-8.204**	2.603	8.713	**-1.684**
		(0.94)	(0.42)	(0.94)	(0.16)	(0.95)	(0.29)		(0.01)	(0.04)	(0.00)	(0.00)	(0.10)	(0.25)

Table 5: **Point and Density Forecasts for Asset Returns II**

Notes: This table reports the $q = \{1, 5, 10\}$ day ahead mean absolute forecast errors and log scores of asset returns over the forecast sample 14 March, 2013 to 21 August, 2020 in Panels A and B respectively. Short, Medium, and Long add the daily change in $\mathcal{C}_d$ $d = \{\text{Short, Medium, Long}\}$ to the baseline dataset with forecast stemming from a BVAR-CSV model. MAFE and LS for these forecasts are relative to those from baseline dataset using the BVAR-CSV-t-MA(1) model. Negative values indicate better forecast performance from the dataset containing a network connectedness measure. Values in parentheses are p-values from the Giacomini and White (2006) test of equal conditional predictive ability. EFA is the the ishares MCFSI EFA ETF, SPY is the S&P500 ETF, USO is the US Oil Fund ETF, XLE is the SDPR energy selector ETF, GLD is the SPDR gold shares ETF, and SLV is the silver ETF.

		Panel A: MAFE							**Panel B: LS**					
Dataset, Model	Horizon	EFA	SPY	USO	XLE	GLD	SLV	Horizon	EFA	SPY	USO	XLE	GLD	SLV
Baseline, BVAR-CSV-t-MA(1)	q=1	0.709	0.661	1.619	1.112	0.687	1.091	q=1	2475.492	2326.068	4026.929	3275.729	2509.330	3424.381
	5	0.688	0.640	1.562	1.073	0.661	1.047	5	2492.081	2336.934	4028.253	3305.414	2445.291	3344.474
	10	0.688	0.640	1.567	1.073	0.658	1.045	10	2584.014	2424.418	4134.514	3410.811	2495.843	3385.078
$\mathcal{C}_d$ d=Short, BVAR-CSV	q=1	**-0.005**	**-0.009**	**-0.020**	**-0.015**	**-0.008**	**-0.011**	q=1	18.884	5.955	23.329	9.738	30.736	34.045
		(0.10)	(0.00)	(0.01)	(0.01)	(0.00)	(0.05)		(0.10)	(0.71)	(0.04)	(0.72)	(0.00)	(0.00)
	5	0.001	0.000	0.001	**-0.001**	0.000	**-0.002**	5	**-7.741**	**-32.602**	8.244	**-4.743**	11.107	**-17.857**
		(0.28)	(0.04)	(0.78)	(0.80)	(0.51)	(0.03)		(0.03)	(0.01)	(0.03)	(0.06)	(0.07)	(0.08)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	**-22.186**	**-62.763**	**-0.274**	**-12.944**	18.002	**-12.908**
		(0.10)	(0.49)	(0.51)	(0.63)	(0.26)	(0.27)		(0.13)	(0.03)	(0.06)	(0.06)	(0.01)	(0.22)
$\mathcal{C}_d$ d=Medium, BVAR-CSV	q=1	**-0.004**	**-0.008**	**-0.021**	**-0.015**	**-0.009**	**-0.011**	q=1	22.259	7.199	21.439	9.587	34.098	35.876
		(0.13)	(0.01)	(0.01)	(0.01)	(0.00)	(0.04)		(0.06)	(0.69)	(0.06)	(0.76)	(0.00)	(0.00)
	5	0.001	0.000	0.001	**-0.001**	0.000	**-0.002**	5	**-7.318**	**-32.159**	7.387	**-4.037**	13.078	**-15.909**
		(0.16)	(0.04)	(0.87)	(0.72)	(0.54)	(0.04)		(0.03)	(0.00)	(0.04)	(0.07)	(0.04)	(0.12)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	**-20.728**	**-62.644**	**-0.033**	**-12.699**	18.738	**-12.690**
		(0.15)	(0.40)	(0.64)	(0.97)	(0.57)	(0.19)		(0.13)	(0.03)	(0.07)	(0.05)	(0.01)	(0.22)
$\mathcal{C}_d$ d=Long, BVAR-CSV	q=1	**-0.006**	**-0.009**	**-0.020**	**-0.015**	**-0.008**	**-0.011**	q=1	9.411	**-4.366**	17.084	**-0.965**	28.531	33.149
		(0.09)	(0.01)	(0.01)	(0.01)	(0.00)	(0.03)		(0.44)	(0.80)	(0.12)	(0.97)	(0.01)	(0.00)
	5	0.001	0.000	0.001	**-0.001**	0.000	**-0.002**	5	**-10.293**	**-37.263**	8.128	**-4.459**	11.847	**-17.422**
		(0.13)	(0.04)	(0.76)	(0.55)	(0.93)	(0.10)		(0.02)	(0.00)	(0.05)	(0.04)	(0.06)	(0.08)
	10	0.000	0.000	0.000	0.000	0.000	0.000	10	**-21.649**	**-63.843**	3.256	**-9.203**	20.464	**-10.988**
		(0.37)	(0.37)	(0.89)	(0.81)	(0.18)	(0.04)		(0.11)	(0.03)	(0.04)	(0.03)	(0.01)	(0.26)

3.4 Horizon Specific Directional Network Rotation Portfolios: Do Conventional Risk Factors Explain the Return?

Here I explore the economic benefits of network connections that form over different horizons within the implied volatility index networks. I do so by constructing long-short portfolios that sort on horizon d net-directional connections and name them directional rotation portfolios. Since these connections form among implied volatility indices, one expects to see negative risk adjusted returns on at least some of the horizon d directional rotation portfolios. This is because risk averse investors will accept lower returns in equilibrium for assets that positively covary with a variable forecasting future volatility because they act as intertemporal hedging devices (Campbell et al., 2018). Investors reduce current consumption to increase precautionary savings in light of uncertainty around future returns and adverse changes to the investment opportunity set (e.g. Ang et al., 2006b,a).

I construct directional rotation portfolios in the following manner. On day τ, I sort the 6 underlying asset returns for each implied volatility index on horizon d net-directional connections. Those above (below) the day τ median contain the most transmissive (receptive) of (to) volatility spillovers within the network. I take a long position is the transmitters and a short position in the receptors which gives the network rotation portfolio $R_p^{\text{NET}_d}$ with daily rebalancing and consider both value- and equally-weighted portfolios. The former weights, for the jth asset is is the day τ price of asset j divided by the sum of all day τ asset prices in the transmitter (receptor) portfolio.

Notably, all of the volatility indices have underlying assets that are ETFs; with the exception of the S&P500 underlying the VIX. To maintain consistency, I use the SDPR S&P500 ETF trust return here. When trading ETFs, investors pay management fees when long and repo costs when short[11]. In order to account for the likely trading costs the investor incurs here, I use daily bid and ask data from Thomson Reuters Datastream to construct effective half spreads (ES) and realized spreads (RS) (see e.g. Bessembinder and Venkataraman, 2010; Marshall et al., 2012). These are given by:

$$\text{ES}_{j,\tau} = 100 \times \delta_{j,\tau} \times \frac{(\text{P}_{j,\tau} - \text{V}_{j,\tau})}{\text{V}_{\tau,t}} \tag{9}$$

$$\text{RS}_{j,t} = 100 \times \delta_{j,\tau} \times \frac{(\text{P}_{j,\tau} - \text{V}_{j,\tau+s})}{\text{V}_{\tau,t}}, \tag{10}$$

[11]For these data the management fees range from 0.07% (iShares EAFE fund ETF) to 0.45% (United States Oil Fund ETF) per annum. These figures are taken from the most up-to-date fund information sheets. The Federal Reserve's overnight repo rate follows closely the Federal Funds rate, which for the sample I consider in this paper, is near to zero. Notably though, the Federal Reserve's overnight repo rate peaks at 2.25% per annum in 2019 before declining to 0.00 from February 2020.

respectively. In (9) and (10), $\delta_{j,\tau}$ is an indicator equal to 1 (-1) if position in the jth asset is long (short); that is, in the transmitter (receptor) portfolio. $P_{j,\tau}$ is the price of the jth asset on day τ. If long the asset the price is the day τ bid price, if short the asset the price is the day τ ask price. $V_{j,\tau}$, $V_{j,\tau+s}$ are observable proxies for the true value of the jth asset on day τ and day $\tau + s$ respectively. In this case I set this to the midpoint of the bid and ask for asset j. ES is an estimate of the execution cost the trader pays and the gross revenue earned by the liquidity provider. RS incorporates the price adjustment from $V_{j,\tau}$ to $V_{j,\tau+s}$ which reflects the market's assessment of private information the trade conveys[12].

Table 6 reports monthly expected returns and risk-adjusted returns for each horizon d directional network rotation portfolio. The risk-adjusted returns are the respective α's from the: Capital Asset Pricing Model (CAPM), α^{CAPM}; the Fama and French (1993) 3-factor model, α^{FF3}; and the Fama and French (2015) 5-factor model, α^{FF5}. Panels A and B account for transaction costs using ES and RS respectively.

Overall, we can see that the risk adjusted returns of short- and medium-horizon directional network rotation portfolios are negative, statistically significant at conventional levels and economically meaningful. Turning to long-horizon network rotation portfolios, risk adjusted returns are slightly positive or negative and statistical significance disappears. Network rotation portfolios sorted on directional connections that aggregate over all horizons, d=Total, exhibit the same pattern as those we form on long-horizon directional connections. Trading on horizon specific directional network connections indicates that long-horizon directional connections cancel out the statistical and economic significance of trading on short- and medium horizon directional connections.

The decomposition of network connections to short-, medium-, and long-horizons shows that trading on directional network connections that form over horizons of 1-week or less yield significant and negative returns. Combining the rationale above, with what we observe in Figures 2–4, these results suggest that short- and medium-horizon network rotation portfolios act as intertemporal hedging devices against transient adverse events affecting uncertainty around these markets. It is not the purpose of this research to uncover the underlying cause of this phenomenon. However, the main implication for decision makers is that these portfolios have the ability to hedge against adverse anticipated changes in investment opportunities that one expects to persist at horizons of less than 1-month[13]. What we can take away from this portfolio analysis is the importance of frequency dynamics for financial markets; something that Dew-Becker and Giglio (2016) and Bandi et al. (2019) highlight.

[12] In the presence of informed traders causes market prices to rise (fall) on average after customer buys (sells). These adverse price movements cause market makers to earn less than the effective spread for their services. RS captures the reversal from the trade price to the post-trade value.

[13] Conversely these portfolios could also be used for speculations purposes.

Table 6: **Risk Adjusted Returns for Directional Network Rotation Portfolios**

Notes: This table presents monthly portfolio expected returns, $\mathbb{E}\left[R_p^{\text{NET}_d}\right]$ for horizon d = {Short, Medium, Long} directional network rotation portfolios, as well as the risk adjusted returns: α^{CAPM}, α^{FF3}, and α^{FF5} from the CAPM, the Fama and French (1993) 3-factor model, and the Fama and French (2015) 5-factor model respectively. Panel A reports results after accounting for transaction costs using effective half spreads and Panel B shows analogous results accounting for transaction costs using realized spreads. t-stat denotes the t-statistic for each return.

Panel A: Transaction Costs = Effective Half Spread

	Value Weighted					Equal Weighted			
	$\mathbb{E}\left[R_p^{\text{NET}_d}\right]$	α^{CAPM}	α^{FF3}	α^{FF5}		$\mathbb{E}\left[R_p^{\text{NET}_d}\right]$	α^{CAPM}	α^{FF3}	α^{FF5}
d=Short	-0.85%	-0.86%	-0.56%	-0.47%	d=Short	-0.55%	-0.64%	-0.40%	-0.37%
t-stat	-1.40	-1.38	-0.89	-0.75	t-stat	-2.01	-2.31	-1.50	-1.39
d=Medium	-1.43%	-1.44%	-1.20%	-1.11%	d=Medium	-0.62%	-0.68%	-0.46%	-0.44%
t-stat	-2.67	-2.64	-2.14	-2.00	t-stat	-2.54	-2.76	-1.89	-1.78
d=Long	0.67%	0.97%	0.76%	0.72%	d=Long	0.20%	0.31%	0.25%	0.23%
t-stat	1.33	1.99	1.54	1.51	t-stat	0.84	1.27	1.01	0.94
d=Total	0.51%	0.77%	0.41%	0.36%	d=Total	0.12%	0.21%	0.09%	0.07%
t-stat	0.98	1.50	0.80	0.73	t-stat	0.49	0.83	0.38	0.28

Panel B: Transaction Costs = Realized Spread

	Value Weighted					Equal Weighted			
	$\mathbb{E}\left[R_p^{\text{NET}_d}\right]$	α^{CAPM}	α^{FF3}	α^{FF5}		$\mathbb{E}\left[R_p^{\text{NET}_d}\right]$	α^{CAPM}	α^{FF3}	α^{FF5}
d=Short	-1.51%	-2.97%	-2.29%	-2.16%	d=Short	-0.75%	-1.65%	-1.21%	-1.17%
t-stat	-1.14	-2.66	-2.06	-1.94	t-stat	-1.11	-3.30	-2.50	-2.40
d=Medium	-1.76%	-3.20%	-2.82%	-2.67%	d=Medium	-0.82%	-1.70%	-1.26%	-1.22%
t-stat	-1.41	-3.13	-2.67	-2.55	t-stat	-1.26	-3.53	-2.69	-2.59
d=Long	0.63%	-0.56%	-0.43%	-0.40%	d=Long	0.00%	-0.70%	-0.53%	-0.54%
t-stat	0.67	-0.78	-0.56	-0.54	t-stat	0.00	-1.86	-1.37	-1.40
d=Total	0.47%	-0.73%	-0.68%	-0.68%	d=Short	-0.08%	-0.79%	-0.67%	-0.69%
t-stat	0.50	-1.00	-0.89	-0.90	t-stat	-0.16	-2.13	-1.75	-1.79

4 Robustness Analysis

For the sake of brevity, I relegate all robustness analysis to the Online Appendix. I first assess the robustness of my main results by taking logs of the implied volatility indices and repeat all the analysis in Section 3. Next I explore changes in the information set by adding a variety of leading volatility indices to the six variables I focus on here. In doing so I construct an additional four alternative specifications. These results show that all network connectedness measures follow similar time profiles. Then, I assess the robustness of the forecasting and portfolio analysis under these alternative specifications. These results provide further substance that my main findings are robust to alternatives.

5 Conclusion

This paper shows how one can capture network connections among implied volatility indices over different horizons using Bayesian vector heterogeneous autoregressive (BVHAR) models with fat tails, serial dependence and heteroskedasticity. I show how the network measures of Diebold and Yılmaz (2014) and Baruník and Křehlík (2018) modify when one wishes to depart from conventional assumptions of innovation distributions.

My main results demonstrate the practical uses of implied volatility index connections for decision makers. Fist, point and density forecasts show statistically significant and meaningful gains when including network connectedness to predict the underlying asset returns. The latter show substantial gains as the forecast horizons increases. This result stresses the importance predictive probability distributions for market participants to make decisions (Geweke and Amisano, 2010) as well as the usefulness of network connectedness. The second sorts the underlying assets on directional connections that reveals an appealing risk return profile even after accounting for transaction costs. These results show that investors can construct hedge portfolios trading on directional connections forming over horizons of less than one month because they act as intertemporal hedging devices against transient adverse changes to investment opportunities (Campbell et al., 2018).

The implications are twofold. First, forecasters should track connections among implied volatility indices in order to predict the underlying asset returns. This is because asset managers whose performance relates to their decisions are willing to pay for even marginal improvements (Campbell and Thompson, 2008). Second, market participants should monitor directional connections either for short-term speculation purposes, or in order inform hedging decisions against adverse events they deem transient.

Acknowledgements

I would like to thank the editor Emanuele Borgonovo, and three anonymous referees for their invaluable comments. I am also grateful to Victoria Ivashina, Chris Florackis, Costas Milas, Chardin Wese Simen, Xi Fu, Marcin Michalski and Adnan Gazi for helpful comments and suggestions. I have nothing to declare and received no external funding support for this research.

References

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

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

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

Acharya, V.V., Pedersen, L.H., Philippon, T., Richardson, M., 2017. Measuring systemic risk. Review of Financial Studies 30, 2–47.

Adrian, T., Brunnermeier, M.K., 2016. CoVaR. American Economic Review 106, 1705–41.

Ang, A., Chen, J., Xing, Y., 2006a. Downside risk. Review of Financial Studies 19, 1191–1239.

Ang, A., Hodrick, R.J., Xing, Y., Zhang, X., 2006b. The cross-section of volatility and expected returns. Journal of Finance 61, 259–299.

Bandi, F.M., Chaudhuri, S., Lo, A.W., Tamoni, A., 2019. Spectral factor models. Johns Hopkins Carey Business School Research Paper .

Barbaglia, L., Croux, C., Wilms, I., 2020. Volatility spillovers in commodity markets: A large t-vector autoregressive approach. Energy Economics 85, 104555.

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 *In Press* .

Baruník, J., Bevilacqua, M., Tunaru, R., 2020. Asymmetric connectedness of fears in the US financial sector. Review of Economics and Statistics *(Forthcoming)* .

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

Baruník, J., Křehlík, T., Vácha, L., 2016. Modeling and forecasting exchange rate volatility in time-frequency domain. European Journal of Operational Research 251, 329–340.

Baumeister, C., Korobilis, D., Lee, T.K., 2020. Energy markets and global economic conditions. Review of Economics and Statistics, *Forthcoming* .

Bessembinder, H., Venkataraman, K., 2010. Bid-ask spreads: Measuring trade execution costs in financial markets. Encyclopedia of Quantitative Finance , 184–190.

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, 1053–1064.

Campbell, J.Y., Giglio, S., Polk, C., Turley, R., 2018. An intertemporal capm with stochastic volatility. Journal of Financial Economics 128, 207–233.

Campbell, J.Y., Thompson, S.B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? The Review of Financial Studies 21, 1509–1531.

Carriero, A., Clark, T.E., Marcellino, M., 2016. Common drifting volatility in large Bayesian VARs. Journal of Business & Economic Statistics 34, 375–390.

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

Chib, S., 1995. Marginal likelihood from the Gibbs output. Journal of the American Statistical Association 90, 1313–1321.

Chiu, C.W.J., Mumtaz, H., Pinter, G., 2017. Forecasting with VAR models: Fat tails and stochastic volatility. International Journal of Forecasting 33, 1124–1143.

Clark, T.E., Ravazzolo, F., 2015. Macroeconomic forecasting performance under alternative specifications of time-varying volatility. Journal of Applied Econometrics 30, 551–575.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics 7, 174–196.

Creal, D., Koopman, S.J., Lucas, A., 2011. A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. Journal of Business & Economic Statistics 29, 552–563.

Cryer, J.D., Chan, K.S., 2008. Time series analysis: with applications in R. Springer Science & Business Media.

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

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

Diebold, F.X., Yilmaz, K., 2012. Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of Forecasting 28, 57–66.

Diebold, F.X., Yılmaz, K., 2014. On the network topology of variance decompositions: Measuring the connectedness of financial firms. Journal of Econometrics 182, 119–134.

Diebold, F.X., Yilmaz, K., 2015. Trans-atlantic equity volatility connectedness: Us and european financial institutions, 2004–2014. Journal of Financial Econometrics 14, 81–127.

Dimitrakopoulos, S., Kolossiatis, M., 2020. Bayesian analysis of moving average stochastic volatility models: Modeling in-mean effects and leverage for financial time series. Econometric Reviews 39, 319–343.

Ding, P., 2016. On the conditional distribution of the multivariate t distribution. The American Statistician 70, 293–295.

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

Engle, R.F., Gallo, G.M., Velucchi, M., 2012. Volatility spillovers in East Asian financial markets: a MEM-based approach. Review of Economics and Statistics 94, 222–223.

Fama, E.F., French, K.R., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Fama, E.F., French, K.R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116, 1–22.

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, 1371–1390.

Geweke, J., Amisano, G., 2010. Comparing and evaluating Bayesian predictive distributions of asset returns. International Journal of Forecasting 26, 216–230.

Giacomini, R., White, H., 2006. Tests of conditional predictive ability. Econometrica 74, 1545–1578.

Hansen, P.R., Lunde, A., Nason, J.M., 2011. The model confidence set. Econometrica 79, 453–497.

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

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, 249–283.

Herskovic, B., Kelly, B.T., Lustig, H.N., Van Nieuwerburgh, S., 2020. Firm volatility in granular networks. Journal of Political Economy .

Marshall, B.R., Nguyen, N.H., Visaltanachoti, N., 2012. Commodity liquidity measurement and transaction costs. The Review of Financial Studies 25, 599–638.

Mumtaz, H., Theodoridis, K., 2018. The changing transmission of uncertainty shocks in the US. Journal of Business & Economic Statistics 36, 239–252.

Pesaran, H.H., Shin, Y., 1998. Generalized impulse response analysis in linear multivariate models. Economics letters 58, 17–29.

Sévi, B., 2014. Forecasting the volatility of crude oil futures using intraday data. European Journal of Operational Research 235, 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, 589–599.

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