



# Factor-Driven Network Informed Restricted Vector Autoregression

Brendan Martin  
Imperial College London  
London, United Kingdom  
b.martin22@imperial.ac.uk

Mihai Cucuringu\*  
University of California, Los Angeles  
Los Angeles, USA  
mihai@math.ucla.edu

Alessandra Luati†  
Imperial College London  
London, United Kingdom  
a.luati@imperial.ac.uk

Francesco Sanna Passino  
Imperial College London  
London, United Kingdom  
f.sannapassino@imperial.ac.uk

## Abstract

High-dimensional financial time series often exhibit complex dependence relations driven by both common market structures and latent connections among assets. To capture these characteristics, this paper proposes Factor-Driven Network Informed Restricted Vector Autoregression (FNIRVAR), a model for the common and idiosyncratic components of high-dimensional time series with an underlying unobserved network structure. The common component is modelled by a static factor model, which allows for strong cross-sectional dependence, whilst a network vector autoregressive process captures the residual co-movements due to the idiosyncratic component. An assortative stochastic block model underlies the network VAR, leading to groups of highly co-moving variables in the idiosyncratic component. For estimation, a two-step procedure is proposed, whereby the static factors are estimated via principal component analysis, followed by estimation of the network VAR parameters. The method is demonstrated in financial applications to daily returns, intraday returns, and FRED-MD macroeconomic variables. In all cases, the proposed method outperforms a static factor model, as well as a static factor plus LASSO-estimated sparse VAR model, in terms of forecasting and financial performance metrics.

## Keywords

Factor models, financial returns prediction, network time series, stochastic block model.

### ACM Reference Format:

Brendan Martin, Mihai Cucuringu, Alessandra Luati, and Francesco Sanna Passino. 2025. Factor-Driven Network Informed Restricted Vector Autoregression. In *6th ACM International Conference on AI in Finance (ICAIF '25)*, November 15–18, 2025, Singapore, Singapore. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3768292.3770412>

\*Also with University of Oxford

†Also with University of Bologna



This work is licensed under a Creative Commons Attribution 4.0 International License. *ICAIF '25, Singapore, Singapore*  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2220-2/25/11  
<https://doi.org/10.1145/3768292.3770412>

## 1 Introduction

Datasets in which the number of variables is comparable to or greater than the number of observations are ubiquitous in finance and economics. In a high-dimensional time series setting, correctly characterizing the dependence among variables is crucial for interpreting the relationships between series as well as for prediction tasks. Two main techniques for imposing low-dimensional structure on high-dimensional time series have emerged in the literature [15, 19, 25]. First, factor models are used to capture co-movements due to shocks common across all series [see 20, and references therein]. Second, sparse models, whereby each series depends only on a small subset of the other series, have proven effective tools for parameter estimation and forecasting [see, for example, 18]. In fact, any second-order stationary process  $\{X_{i,t} : i \in \mathbb{N}, t \in \mathbb{Z}\}$  can be decomposed into a so-called common or systematic component, driven by a finite number of (possibly dynamic) factors, and an orthogonal idiosyncratic component [7]. As such, combining dense factor models with sparse methods has been widely used to capture both sources of co-movement [6, 13, 19, 32, 41]. Dense, here, refers to any model in which all explanatory variables play a role in prediction, even though the impact of some of them may be small.

In this work, we propose a novel factor-based framework for modelling multivariate time series with an underlying network structure, called Factor-Driven Network Informed Restricted Vector Autoregression (FNIRVAR). The proposed approach combines a static factor model [3, 45] with a network-based vector autoregressive (VAR) model for the idiosyncratic component. Network VAR models, introduced by Zhu et al. [52] and Knight et al. [31], treat each univariate time series as being observed on a vertex of a graph, with the VAR coefficient matrix encoding the edge structure of the graph. The network structure can be thought of as a variable selection mechanism, with only those vertices which are connected having non-zero VAR coefficients. This regularisation allows for parameter estimation in high dimensions. The majority of existing works on network VAR methods assume the network to be observable [see, for example, 5, 30, 50]. In financial and economic applications, however, the assumption of an observed underlying network is often unrealistic. A recent framework proposed by Martin et al. [39], called Network Informed Restricted Vector Autoregression (NIRVAR), allows for an unobserved network, instead using a data-driven method to estimate the network. We therefore adopt the

NIRVAR methodology to model the idiosyncratic component. Since we augment the NIRVAR model by including unobserved factors, we denote this factor driven model as FNIRVAR.

For the estimation of FNIRVAR, we propose a two-step procedure consisting in factor estimation based on principal component analysis (PCA), followed by NIRVAR estimation on the remaining component. The NIRVAR estimation framework entails finding groups of highly co-moving time series via spectral embedding and clustering, and subsequently estimating the VAR coefficients for series in the same group only [39]. All VAR parameters for variables belonging to different groups are restricted to zero. In a financial prediction context, this grouping method could be construed as a generalisation of a pairs trading strategy [see 49, for a discussion of pairs trading]: instead of building a strategy based upon two highly correlated assets, we compare groups of many highly correlated assets. In particular, the future returns of a given asset are forecasted as a weighted sum of the assets within its group (after controlling for common factors). In our proposed estimation procedure, we follow the approach of Fan et al. [19] by assuming a large eigengap between the covariance matrix of the common component and that of the idiosyncratic component. An alternative approach is the joint optimisation method of Agarwal et al. [2]. This, however, requires a sparsity assumption on the idiosyncratic component which we do not impose. In Section 4, we conduct a simulation study showing that assuming  $r$  strong factors in the FNIRVAR model leads to the largest (in magnitude)  $r$  eigenvalues exploding with dimension  $N$ , whereas the idiosyncratic eigenvalues are bounded. This justifies our proposed two-step estimation procedure.

### 1.1 Main contributions

The primary contribution of this paper is the development of FNIRVAR, a factor driven NIRVAR model, and a related estimation framework. To the best of our knowledge, this is the first study that combines a dense factor model with a sparse network time series model without assuming the network to be observed. The benefits of this approach are three-fold: (i) accounting for unobserved common factors mitigates endogeneity issues that arise if the common component is not modelled; (ii) the NIRVAR estimation method captures residual co-movements between groups of series which are not picked up by the estimated factors, leading to enhanced forecasting performance; (iii) the estimated factors can lead to meaningful interpretations about the forces driving the common component whilst estimated NIRVAR groups can shed light on the network driving the idiosyncratic component. The usefulness of the proposed methodology is demonstrated in an extensive application to financial returns prediction, as well as in a prediction task on the macroeconomic FRED-MD dataset [40].

### 1.2 Outline

The rest of this paper is organised as follows: the remainder of Section 1 discusses related literature and establishes the notation used in the paper. Next, Section 2 defines the FNIRVAR model and gives the conditions under which the model is stationary. Section 3 states the strong factor assumption with a large eigengap, and details the two-step estimation procedure. Simulation studies analysing

the predictive performance of FNIRVAR under various data generating processes, and the effect of violating the large eigengap assumption on FNIRVAR estimation are illustrated in Section 4. Section 5 contains applications to financial returns prediction and a macroeconomic application to forecasting US industrial production using FRED-MD. All code is available in the GitHub repository `bmartin9/fnirvar`.

### 1.3 Related literature

This work contributes to the growing literature combining dense factor models with sparse models [see, for example, 6, 13, 19, 32, 41] and references therein. Factor models have a long history in econometrics starting with the dynamic exact factor models of Geweke [24] and Sargent et al. [44]. Here, exact refers to the assumption of an idiosyncratic error term with diagonal covariance, whereas approximate factor models have cross-sectional dependence in the idiosyncratic term. In particular, the factor model structure considered in this work is the static approximate factor model, introduced by Chamberlain [11] and Chamberlain and Rothschild [12], and further developed by Stock and Watson [45], Stock and Watson [46], Bai [3], and Bai and Ng [4]. A wider class that encompasses all other factor models is the General Dynamic Factor Model proposed by Forni et al. [21] and Forni and Lippi [22]. We note that although we impose dynamics on the factors in this work, we require a finite number of lags, and any dynamic factor model with a finite number of lags can be written as a static factor model [7]. For this reason, we consider our contribution as a static approximate factor model.

As an alternative to factor models, regularised estimation procedures such as the Least Absolute Shrinkage and Selection Operator [LASSO; 47] impose sparsity assumptions to achieve dimensionality reduction. In the context of VAR modelling, there is a large literature on regularised estimation of sparse VAR models [see, for example 1, 8, 28, 31, 43].

The approaches of factor and sparse models have been first combined in the seminal paper of Fan et al. [19]. Next, both Miao et al. [41] and Krampe and Margaritella [32] propose  $\ell_1$ -type regularised estimators for a factor-augmented sparse VAR model. Similar to our framework, Krampe and Margaritella [32] utilise a two-step estimation procedure in which PCA is used in a first step to estimate the factors and the idiosyncratic component is estimated in a second step. In contrast, Miao et al. [41] employ a joint optimisation procedure to estimate the common and idiosyncratic components.

Additionally, there are a number of papers that develop network-based models for the idiosyncratic component, whilst including a factor model to account for the common component. Chen et al. [13] extends the network autoregressive model of Zhu et al. [52] to allow group-wise network effects as well as static factors. The major difference between our work and that of Chen et al. [13] is that we do not assume that the network is observed. Furthermore, their model is based on an expected adjacency matrix, whereas the model we consider is based on a realised adjacency matrix. Barigozzi et al. [6] develop a dynamic factor plus sparse VAR model and propose an  $\ell_1$ -regularised Yule-Walker estimator for the latent VAR process. The variables selected by their  $\ell_1$ -regularised method are then used to construct a network. In contrast, we construct the

network as a first step, and use it to inform the estimation of the latent VAR process.

Finally, we mention the growing use of modern machine learning techniques for time series prediction tasks. Gu et al. [27] provide a comprehensive overview and comparison of machine learning methods for empirical asset pricing, and suggest that the inclusion of nonlinear predictor interactions can lead to predictive gains missed by more traditional regression-based strategies.

## 1.4 Notation

For an integer  $k$ , we let  $[k]$  denote the set  $\{1, \dots, k\}$ . The  $n \times n$  identity matrix is written  $I_n$ , and the indicator function is  $\mathbb{1}\{\mathcal{B}\} = 1$  if the event  $\mathcal{B}$  occurs and 0 otherwise. For vectors,  $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^p$ , we let  $M = (\mathbf{v}_1, \dots, \mathbf{v}_n) \in \mathbb{R}^{p \times n}$  be the matrix whose columns are given by  $\mathbf{v}_1, \dots, \mathbf{v}_n$ . The transpose of  $M$  is denoted  $M'$ . We write  $M_{i,:}$  and  $M_{:,j}$  to denote the  $i$ -th row and  $j$ -th column of  $M$  considered as vectors, respectively. We use  $\odot$  to denote the Hadamard product between two matrices of the same dimensions. For matrices  $M_1, \dots, M_r$ ,  $(M_1 | \dots | M_r)$  and  $(M_1; \dots; M_r)$  denote the column-wise and row-wise concatenation of the matrices, respectively. We write  $\mathcal{S}_{K-1} = \{x \in \mathbb{R}^K : \sum_{k=1}^K x_k = 1, x_k > 0, k \in [K]\}$  to denote the interior of the unit simplex in  $K$  dimensions.

## 2 Model

In this section, we describe the FNIRVAR model, the main contribution of this work. In order to define FNIRVAR, we first introduce the stochastic block model, which is the random graph used to model the idiosyncratic component. For a review of stochastic block models, see Lee and Wilkinson [34]. Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  be a random graph with vertex set  $\mathcal{V} = [N]$  and where  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$  is a set of random edges. In a stochastic block model, each vertex belongs to one of  $K$  communities called blocks and the probability of an edge forming between two vertices depends only on their block memberships.

*Definition 2.1.* (Stochastic block model). Let  $B \in \mathbb{R}^{K \times K}$  be a positive semi-definite matrix of block connection probabilities where  $K \in \mathbb{N}$  is the number of blocks. Let  $\pi = (\pi_1, \dots, \pi_N)' \in \mathcal{S}_{K-1}$  represent the prior probabilities of each node belonging to the  $k$ -th community for  $k \in [K]$ . We say  $A$  is the adjacency matrix for a stochastic block model, written  $A \sim \text{SBM}(B, \pi)$ , if the block assignment map  $z : [N] \rightarrow [K]$  satisfies  $\mathbb{P}\{z(i) = k\} = \pi_k$  for each  $k \in [K]$  and, conditional on  $z$ , the entries  $A_{ij}$  are independently generated as

$$A_{ij} = \begin{cases} \text{Bernoulli}\left(B_{z(i), z(j)}\right) & \text{for } i \neq j \\ 1 & \text{for } i = j. \end{cases}$$

**REMARK 1.** *The condition that  $B$  is positive semi-definite in Definition 2.1 corresponds to an assortative graph, referring to the tendency of similar vertices to be connected to each other [42]. In a network time series context, this amounts to assuming greater co-movement between series within the same community than between communities.*

We now define the FNIRVAR model in which the common component is modelled as a (restricted) dynamic factor model and the idiosyncratic component is modelled as a restricted VAR process whereby the idiosyncratic VAR restrictions are determined by the adjacency matrix of a SBM.

*Definition 2.2.* (FNIRVAR model). Let  $\{X_t\}_{t \in \mathbb{Z}}$  denote a zero mean, second order stationary stochastic process where  $X_t = (X_{1,t}, \dots, X_{N,t})' \in \mathbb{R}^N$ . The FNIRVAR model is defined as:

$$\begin{aligned} X_t &= \Lambda F_t + \xi_t, \\ F_t &= \sum_{k=1}^{l_F} P_k F_{t-k} + \Gamma_u^{1/2} \mathbf{u}_t, \quad \mathbf{u}_t \sim \mathcal{N}(0, I_q), \\ \xi_{i,t} &= \sum_{j=1}^N A_{ij} \tilde{\Phi}_{ij} \xi_{j,t-1} + \epsilon_{i,t}, \quad \epsilon_t \sim \mathcal{N}(0, \Gamma_\epsilon), \\ A &\sim \text{SBM}(B, \pi), \end{aligned} \tag{1}$$

where  $\Lambda$  is a  $N \times r$  matrix of loadings,  $F_t$  is an  $r$ -dimensional vector of factors, where  $r \in \mathbb{N}$  is the number of factors,  $P_k$  are  $r \times r$  coefficient matrices of the factor VAR process with  $l_F \in \mathbb{N}$  lags,  $\Gamma_u$  is an  $r \times q$  matrix, and  $\mathbf{u}_t$  are common shocks. We further define  $P := (P_1; \dots; P_{l_F}) \in \mathbb{R}^{r l_F \times r}$ . The common component is written as  $\chi_t = \Lambda F_t$ , whereas the idiosyncratic component is  $\xi_t$ . The parameters  $\tilde{\Phi}_{ij}$  are taken as fixed. We also define  $\Phi := A \odot \tilde{\Phi}$ .

The NIRVAR idiosyncratic model can be interpreted as a restricted VAR model in which the restrictions on  $\Phi$  are determined by the adjacency matrix of an assortative SBM. As such,  $\Phi$  will have a block diagonal structure with a greater number of non-zero entries within each block than between blocks. The sparsity level of  $\Phi$  is determined by the block probability matrix  $B$ .

For ease of presentation, we consider the standard unweighted SBM and post multiply each realised edge  $A_{ij}$  by a fixed weight  $\tilde{\Phi}_{ij}$ . Equivalently, one could model  $\Phi$  as being the adjacency matrix of a weighted SBM [for a definition, see 23]. In this formulation,  $\Phi_{ij} | z(i), z(j) \sim H\{z(i), z(j)\}$  for some family of distributions  $H$  that depends only on the block assignments  $z(i)$ , and not on the individual vertices.

It is worth pointing out that, in expectation, the idiosyncratic NIRVAR model can itself be viewed as a factor model. Let

$$\mathbb{E}\{\Phi | z(i) = k, z(j) = l\} = \tilde{B}_{kl},$$

where  $\tilde{B} \in \mathbb{R}^{K \times K}$  is the block mean matrix. Then the expected value of the observed weighted adjacency matrix can be written as  $\mathbb{E}(\Phi) = Z \tilde{B} Z'$ , with  $Z \in \{0, 1\}^{N \times K}$  a binary matrix whose entries satisfy  $Z_{ik} = 1$  when vertex  $i$  belongs to community  $k \in [K]$  and zero otherwise. Replacing  $\Phi$  with  $\mathbb{E}(\Phi)$ , the NIRVAR model yields

$$\xi_t = Z \tilde{B} Z' \xi_{t-1} + \epsilon_t = \Lambda_\xi F_t^{(\xi)} + \epsilon_t,$$

where  $F_t^{(\xi)} := Z' \xi_{t-1}$  is a  $K$ -dimensional vector of factors encoding the aggregated response in each community and  $\Lambda_\xi := Z \tilde{B}$  are the loadings of each panel component onto these ‘‘community factors’’. This relation between the expected adjacency matrix and community factors is also discussed in Chen et al. [13].

Before describing the FNIRVAR estimation procedure, we analyse the conditions under which the model is stationary.

**PROPOSITION 1 (STATIONARITY).** *The FNIRVAR process is weakly stationary if and only if  $\det(I_{N l_F} - P^* z) \neq 0$  and  $\det(I_N - \Phi z) \neq 0$*

for  $|z| \leq 1$ , where

$$P^* = \begin{bmatrix} P_1 & P_2 & P_3 & \cdots & P_{l_F} \\ I_N & 0 & 0 & \cdots & 0 \\ 0 & I_N & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & I_N & 0 \end{bmatrix}$$

is the companion matrix of the factor VAR process in (1).

**PROOF.** The VAR processes  $\{F_t\}_{t \in \mathbb{Z}}$  and  $\{\xi_t\}_{t \in \mathbb{Z}}$  are weakly stationary if and only if  $\det(I_r - P_1 z - \cdots - P_{l_F} z^{l_F}) \neq 0$  and  $\det(I_N - \Phi z) \neq 0$  for all  $z \in \mathbb{C}$  such that  $|z| \leq 1$  [see 9, for example]. By Tsay [48, Lemma 2.1],  $\det(I_r - P_1 z - \cdots - P_{l_F} z^{l_F}) = \det(I_{N l_F} - P^* z)$ . By Lütkepohl [36, Theorem 1], if  $\Lambda \neq 0$ , then  $\{\Delta F_t\}_{t \in \mathbb{Z}}$  is weakly stationary whenever  $\det(I_{N l_F} - P^* z) \neq 0$  for  $|z| \leq 1$ . Finally, by Granger and Morris [26], the sum of the processes  $\{\Delta F_t\}_{t \in \mathbb{Z}}$  and  $\{\xi_t\}_{t \in \mathbb{Z}}$  is a stationary vector autoregressive moving average process whenever the individual processes are stationary.  $\square$

### 3 Estimation

We adopt a two-step estimator of the FNIRVAR parameters  $\Lambda$ ,  $F_t$ ,  $P_k$ , and  $\Phi$ . In the first step, the factors  $F_t$  and loadings  $\Lambda$  of the factor model are estimated via PCA. After removing the estimated common component, the latent VAR parameters  $\Phi$  of the idiosyncratic model are then estimated in a second step. The following assumption of a large eigengap is required for this procedure.

**ASSUMPTION 1.** *There exists a positive eigengap  $\Delta := \lambda_{\min}(\Gamma_\chi) - \lambda_{\max}(\Gamma_\xi)$  between the smallest nonzero eigenvalue of the common covariance matrix  $\Gamma_\chi$  and the largest eigenvalue of the idiosyncratic covariance matrix  $\Gamma_\xi$ .*

Consider a set of observations  $X = (x_1, \dots, x_T)$  where  $x_t = (x_{1t}, \dots, x_{Nt})'$  is a realisation of the random variable  $X_t$ . Let  $E := \text{eig}(XX'/T)_{[1:r]}$  be the  $N \times r$  matrix of eigenvectors of the sample covariance matrix corresponding to the  $r$  largest eigenvalues. We estimate the factors and loadings via static principal component analysis as  $\widehat{\Lambda} = E$  and  $\widehat{F}_t = E' x_t$ . For prediction, the next step ahead factor estimate is given by

$$\widehat{F}_{t+1} = \sum_{k=1}^{\ell_F} \widehat{P}_k \widehat{F}_{t-k},$$

where  $P$  is estimated by least-squares regression. In particular, if  $Y_F = (F'_{\ell_F}, \dots, F'_T)' \in \mathbb{R}^{(T-\ell_F) \times r}$  and

$$X_F = \begin{pmatrix} F'_0 & \cdots & F'_{\ell_F-1} \\ \vdots & \ddots & \vdots \\ F'_{T-\ell_F-1} & \cdots & F'_{T-1} \end{pmatrix} \in \mathbb{R}^{(T-\ell_F) \times r \ell_F},$$

then

$$\widehat{P} = (\widehat{X}'_F \widehat{X}_F)^{-1} \widehat{X}'_F Y_F.$$

We choose the number of factors  $r$  using the information criteria of Bai and Ng [4] and the order  $l_F$  using AIC [19]. Once the common component has been estimated as  $\widehat{\chi}_t = \widehat{\Lambda} \widehat{F}_t$ , we estimate the idiosyncratic component as  $\widehat{\xi}_t = X_t - \widehat{\chi}_t$ . Let  $\widehat{\chi} = (\widehat{\chi}_1, \dots, \widehat{\chi}_T)$  and  $\widehat{\xi} = (\widehat{\xi}_1, \dots, \widehat{\xi}_T)$ . The parameters of the NIRVAR model for

#### Algorithm 1 FNIRVAR estimation.

**Require:**  $N \times T$  design matrix  $X$ ; number of factors  $r$ ; factor VAR lag order  $l_F$ ; embedding dimension  $d$ ; number of clusters  $K$ .

- 1: Estimate  $r$  factors and corresponding loadings via static PCA:  $\widehat{\Lambda} = E$  and  $\widehat{F}_t = E' x_t$  where  $E := \text{eig}(XX'/T)_{[1:r]}$ .
- 2: Estimate  $l_F$  factor VAR coefficient matrices  $P_k$  via least-squares.
- 3: Estimate the idiosyncratic component as  $\widehat{\xi}_t = X_t - \widehat{\chi}_t$ .
- 4: Compute the sample covariance matrix  $\widehat{\Gamma}_\xi \leftarrow \widehat{\xi} \widehat{\xi}' / T$ .
- 5: Obtain the  $d$  eigenvectors  $v_1, \dots, v_d$  corresponding to the  $d$  largest eigenvalues in magnitude of  $\widehat{\Gamma}_\xi$ , and form the  $N \times d$  matrix  $V$  with columns  $v_1, \dots, v_d$ .
- 6: Cluster the rows of  $V$  via a  $K$ -component Gaussian mixture model.
- 7: Set  $\widehat{z}(i)$  to the cluster assignment of row  $i$  for each  $i \in [N]$ .
- 8: Define  $\widehat{A}_{ij} = \mathbb{1}\{\widehat{z}(i) = \widehat{z}(j)\}$ .
- 9: Estimate the parameters  $\widehat{\Phi}$  of the restricted VAR model  $\xi_t = \widehat{A} \circ \widehat{\Phi} \xi_{t-1} + \epsilon_t^{(\widehat{A})}$  via OLS.

**Return** Factors  $\widehat{F}_t$ , loadings  $\widehat{\Lambda}$ , factor VAR coefficient matrices  $\widehat{P}_k$ , cluster labels  $\widehat{z}(i)$ ; estimated coefficient matrix  $\widehat{\Phi}$ .

the idiosyncratic component are then estimated via Algorithm 1. Details on the final step of Algorithm 1 of estimating a restricted VAR model via OLS can be found in Chapter 5 of Lütkepohl [37].

Some remarks are in order motivating the NIRVAR estimation procedure in Algorithm 1. Under suitable sparsity and noise levels, spectral clustering of the adjacency matrix of an SBM consistently recovers the blocks of the SBM [35]. In our case, the adjacency matrix is unobserved and we instead use the sample covariance matrix  $\widehat{\Gamma}_\xi$  as the similarity matrix to embed. This is reasonable since the covariance matrix  $\Gamma_\xi$  inherits the block structure of the adjacency matrix through the relation  $\text{vec}(\Gamma_\xi) = (I_{N^2} - \Phi \otimes \Phi)^{-1} \text{vec}(\Gamma_\epsilon)$  [37]. As proven by Martin et al. [39], the hard thresholding given by step (4) above yields a biased estimator whenever the network contains edges between nodes in different communities. On the other hand, the advantages of the hard thresholding are a reduction in variance of the VAR estimates, as well as regularisation which allows for estimation in a high-dimensional setting.

Various methods for choosing the embedding dimension  $d$  have been proposed in the literature, such as the profile likelihood method of Zhu and Ghodsi [51] or the ScreeNOT method of Donoho et al. [16]. We choose  $d$  using tools from random matrix theory that are standard in finance and economics [see 33, for example]. In particular,  $d$  is set equal to the number of eigenvalues of  $\widehat{\Gamma}_\xi$  that are larger than the upper bound of the support of the Marčenko-Pastur distribution [38]. Since the Marčenko-Pastur distribution is the limiting eigenvalue distribution of a Wishart matrix, the interpretation is that those eigenvalues that are greater than the upper bound of its support cannot be attributed to random noise and are thus “informative directions”.

The number of clusters  $K$  could be chosen using an information criterion such as BIC. However, we follow the argument of Martin et al. [39] and set  $K = d$  since, when represented as a latent position random graph [29], the expected adjacency matrix of an SBM has  $K$  distinct latent positions each of dimension  $d = K$ .

	FNIRVAR	Factors Only	Factors + LASSO
MSPE	<b>1.87</b>	1.91	1.95
SE ( $\times 10^{-1}$ )	<b>1.18</b>	1.21	1.23

**Table 1: MSPE and corresponding standard error (SE) for three prediction models when the DGP is FNIRVAR.**

The computational complexity of FNIRVAR is dominated by estimation of the idiosyncratic component, which is  $O(N^3/K^2 + TN^2/K)$  when  $\Gamma_\epsilon$  is homoskedastic and  $O(N^3 + TN^2)$  otherwise (since generalised least-squares estimation must be used in this case) [39]. Estimation of the factors and loadings via PCA on the other hand can be done using a Lanczos type algorithm in  $O(TNr)$  time. Least-squares estimation of the factor VAR process is  $O(Tr^2l_F^2)$  and is therefore subdominant since  $r \ll N$  and  $l_F \ll N$ . We note that least-squares estimation of an unrestricted VAR is  $O(N^3 + TN^2)$ , and thus, FNIRVAR estimation is faster whenever  $\Gamma_\epsilon$  is assumed to be homoskedastic.

## 4 Simulation study

We conduct three simulation studies investigating the relationship between the common and idiosyncratic components of FNIRVAR. Firstly, we generate data from the FNIRVAR model and conduct a backtesting experiment comparing the predictive performance of a factor model with no idiosyncratic model, a factor model with a LASSO estimated sparse VAR model, and the FNIRVAR prediction model outlined in Section 3. As expected, given the data generating process, FNIRVAR achieves the lowest mean squared prediction error (MSPE). Secondly, we look at the predictive performance of FNIRVAR when the generative model is a static factor model. The study highlights the benefits of including the FNIRVAR idiosyncratic model when the estimated number of factors is misspecified. Thirdly, we consider the effect of Assumption 1 on the sample covariance eigenvalues under a FNIRVAR generative model. In particular, the study highlights the necessity of Assumption 1 for the proposed two-step estimation procedure.

### 4.1 Prediction under FNIRVAR

We generate data from an FNIRVAR model with  $N = 100$ ,  $T = 1500$ ,  $r = 5$ ,  $q = 5$ ,  $l_F = 2$ ,  $K = 4$ ,  $B_{ij} = 0.9$  for  $i = j$  and  $B_{ij} = 0.1$  for  $i \neq j$ .  $P_k$  was chosen to have unit diagonal entries and off diagonal entries equal to -0.2, with a subsequent rescaling such that  $\rho(P^*) = 0.7$ . The weights for the NIRVAR coefficient matrix are chosen as  $\tilde{\Phi}_{ij} \sim \mathcal{N}(\mu_i, 1)$  where  $\mu_i := (-1)^{z(i)+1}z(i)$ , and  $\Phi = A \odot \tilde{\Phi}$  is then normalised such that  $\rho(\Phi) = 0.9$ . The error covariance matrices  $\Gamma_u$  and  $\Gamma_\epsilon$  are both set to be the identity. The task is one-step ahead prediction using a rolling window backtest framework with a lookback window of 1000 observations. Table 1 shows the MSPE and corresponding standard error of the MSPE for three different prediction models: FNIRVAR, a static factor model with no idiosyncratic model, and a static factor model with a LASSO estimated sparse VAR model. The LASSO penalty was chosen via BIC. FNIRVAR achieves the lowest MSPE as expected given that the data generating process is the FNIRVAR model.

Metric	MSPE			Standard error		
	1	3	5	1	3	5
Model \ $\hat{r}$						
FNIRVAR	1.41	1.29	<b>1.22</b>	0.09	0.08	<b>0.07</b>
Factors Only	3.20	1.51	<b>1.22</b>	0.22	0.32	<b>0.07</b>


**Table 2: MSPE and corresponding standard error of two different prediction models for a backtesting experiment on data generated by a static factor model with  $r = 5$ .**

### 4.2 Misspecifying the number of factors

We generate data from a static factor model with  $N = 100$ ,  $T = 1500$ ,  $r = 5$  and  $l_F = 2$ . The number of common shocks was  $q = 5$  and  $P_k$  was chosen to have unit diagonal entries and off diagonal entries equal to -0.2, with a subsequent rescaling such that  $\rho(P^*) = 0.7$ . The task is one-step ahead prediction using a rolling window backtest framework with a lookback window of 1000 observations. Table 2 shows the MSPE and corresponding standard error for static factor and FNIRVAR prediction models, assuming different  $r$  values for the predictions. For FNIRVAR, the number of blocks was set as  $K = 4$ . Even though the data generating process has no idiosyncratic component, FNIRVAR outperforms a static factor model when the number of factors is misspecified. The NIRVAR idiosyncratic model captures some of the variance left over from under-specifying  $r$ . Nonetheless, FNIRVAR does not adversely affect the predictive performance when  $r$  is correctly specified.

### 4.3 Violating the positive eigengap assumption

We simulate data from the FNIRVAR model for different values of  $N$ , with  $T = 1000$ , and monitor the growth of the eigenvalues of  $\hat{\Gamma}$  as  $N$  increases. For this simulation, we choose a single factor ( $r = 1$ ) and five blocks ( $K = 5$ ). The loadings  $\Lambda$  are sampled from a mixture of normals:  $\Lambda_{i,j} \sim 0.5 \cdot \mathcal{N}(1, \sigma_\Lambda^2) + 0.5 \cdot \mathcal{N}(-1, \sigma_\Lambda^2)$ . The factor VAR coefficient matrices are sampled independently from  $\mathcal{N}(0, 0.1)$  and then rescaled such that  $\rho(P^*) = 0.6$  to ensure stationarity. The number of common shocks is chosen to be equal to the number of factors ( $q = r$ ) and  $\Gamma_u$  is set to the identity matrix. The number of factor lags is chosen to be  $l_F = 5$ . The block matrix  $B$  is set to have diagonal values of 0.95 and off-diagonal values of 0.01. The weights for the NIRVAR coefficient matrix are chosen as  $\tilde{\Phi}_{ij} \sim \mathcal{N}(\mu_i, 1)$  where  $\mu_i := (-1)^{z(i)+1}z(i)$ , then symmetrized as  $\tilde{\Phi} \leftarrow (\tilde{\Phi} + \tilde{\Phi}')/2$ . The NIRVAR coefficient matrix  $\Phi = A \odot \tilde{\Phi}$  is then normalised such that  $\rho(\Phi) = 0.9$ . The idiosyncratic error covariance matrix is chosen to be the identity:  $\Gamma_\epsilon = I_N$ . For each  $N$ , 20 replicate datasets are simulated using the above hyperparameters. We control the eigengap  $\Delta$  through the choice of  $\sigma_\Lambda$ . For  $\sigma_\Lambda^2 = 9 \times 10^{-3}$ ,  $\Delta > 0$  whereas for  $\sigma_\Lambda^2 = 2.5 \times 10^{-4}$ ,  $\Delta < 0$ . Figure 1 shows the minimum eigenvalue of  $\hat{\Gamma}_\chi$  and maximum eigenvalue of  $\hat{\Gamma}_\xi$  as a function of  $N$  for the two scenarios:  $\Delta > 0$  and  $\Delta < 0$ . In both cases, the eigenvalue of  $\hat{\Gamma}_\chi$  grows linearly in  $N$ , as is expected for a strong factor [see 7, for further discussion on the growth rate of eigenvalues of factor models]. However, when  $\Delta < 0$  (the case where Assumption 1 is violated),  $\lambda_{\max}(\hat{\Gamma}_\chi) < \lambda_{\min}(\hat{\Gamma}_\xi)$ , and therefore, removing the top eigenvalues of  $\hat{\Gamma}$ , as in our two-step estimation procedure, will not



**Figure 1: Growth of eigenvalues of  $\hat{\Gamma}$  with  $N$  under a FNIRVAR model when Assumption 1 (a) holds and (b) is violated.**

pick out the common component. As a further observation, we note that the eigenvalues of  $\hat{\Gamma}$  do not grow linearly with  $N$ , confirming that the idiosyncratic component of FNIRVAR cannot be interpreted as a strong factor model.

## 5 Empirical analysis

The FNIRVAR model is applied to three data sets: (i) a financial dataset of daily excess market returns; (ii) a financial dataset of thirty minutely intra-day excess market returns; (iii) a vintage of macroeconomic indicators typically used for predicting US industrial production. Descriptions of the data sources and transformations applied to each of the raw data sets are given below. Our primary interest is the one-step-ahead out-of-sample predictive performance of FNIRVAR. We compare the results against the factor plus sparse VAR model of Miao et al. [41] and Fan et al. [19] as well as a static factor model with no idiosyncratic model. Fan et al. [19] estimate the sparse VAR model using LASSO whereas Krampe and Margaritella [32] estimate it using adaptive LASSO. We follow Fan et al. [19] and use LASSO, with the tuning parameter selected via BIC. For the financial applications, we examine the economic benefit of an FNIRVAR trading strategy under different portfolio allocation and transaction cost settings.

### 5.1 Application to daily returns prediction

*Data description.* The previous close-to-close (pvCLCL) price returns of  $N = 648$  financial assets between 03/01/2000 and 31/12/2020 ( $T = 5279$ ) were derived from databases provided by the Center for Research in Security Prices, LLC, an affiliate of the University of Chicago Booth School of Business. The assets considered include stocks in the S&P 500 and Nasdaq as well as a small number of exchange-traded funds. Using the S&P 500 index as a proxy for the market, we construct the pvCLCL market excess returns by subtracting the return of SPY, the exchange traded fund which tracks the S&P 500 index.

*Task.* The task is to predict the next day pvCLCL market excess returns. We adopt a rolling window backtesting framework with a fixed four year look-back. The predictive performance between 01/01/2004 and 31/12/2020 of the three models are compared via commonly used metrics in the financial literature [27]. The mean number of factors chosen by PCp2 [4] over the backtest period was 12.16, whilst the mean order of the factor VAR model was 3.97.

Metric	Sharpe ratio			Mean daily PnL		
	Cost			Cost		
Model	0 bpts	1 bpts	2 bpts	0 bpts	1 bpts	2 bpts
FNIRVAR	<b>1.95</b>	1.66	1.37	<b>3.41</b>	2.90	2.40
Factors + LASSO	1.64	1.29	0.94	2.36	1.86	1.35
Factors Only	1.27	1.07	0.86	3.09	2.59	2.09

**Table 3: Sharpe ratio, mean daily PnL (in bpts) for three prediction models under an equal weighted portfolio construction, for different levels of fixed transaction costs.**

Metric	Sharpe ratio			Mean daily PnL		
	Cost			Cost		
Decile	0 bpts	1 bpt	2 bpts	0 bpts	1 bpt	2 bpts
100	1.95	1.66	1.37	3.41	2.90	2.40
75	1.99	1.76	1.53	4.35	3.85	3.34
50	<b>2.02</b>	1.83	1.65	5.50	4.99	4.49
25	1.90	1.76	1.62	<b>6.81</b>	6.30	5.80

**Table 4: Sharpe ratio and mean daily PnL (in bpts) for equal weighted decile portfolios under different transaction costs.**

*Performance metrics.* Each day we measure the profit and loss as  $\text{PnL}_t = \sum_{i=1}^N \omega_i^{(t)} \text{sign}(\hat{s}_i^{(t)}) s_i^{(t)}$  where  $\hat{s}_i^{(t)}$  is the predicted return of asset  $i$  on day  $t$ ,  $s_i^{(t)}$  is the realised return of asset  $i$  on day  $t$ , and  $\omega_i^{(t)}$  are portfolio weights such that  $\sum_{i=1}^N \omega_i^{(t)} = 1$  for each  $t \in [T]$ . Letting  $\text{PnL} := \{\text{PnL}_t\}_{t=1, \dots, T}$ , the annualised Sharpe Ratio (SR) is then given by  $\sqrt{252} \times \text{mean}(\text{PnL}) / \text{stdev}(\text{PnL})$ . Furthermore, we compute the mean daily PnL (given in basis points, denoted *bpts*, with 1% = 100bpts). Fixed transaction costs of 1bpts and 2bpts are also considered, whereby a fixed amount is subtracted from the PnL every time we flip our position in a given asset from long (short) to short (long). This is a realistic level of transaction cost that one might expect from a broker.

*Equal-weighted portfolio.* The simplest possible portfolio construction strategy is an equal weighted portfolio defined by uniform weights,  $\omega_i^{(t)} = 1/N$  for all  $i \in [N], t \in [T]$ . Table 3 gives the Sharpe ratios and mean daily PnL for each model for different levels of fixed transaction costs. FNIRVAR has the highest Sharpe ratios and mean daily PnL values across all transaction cost levels, followed by the LASSO estimated factor plus sparse VAR model. This suggests that modelling the idiosyncratic component on top of the common component can lead to significant economic gains.

We also consider decile portfolios, in which solely the top  $x\%$  largest in magnitude signals are traded each day [17]. The Sharpe ratio and mean daily PnL (using FNIRVAR as the prediction model) for equal weighted decile portfolios with different transaction costs are shown in Table 4. The table shows that lower decile portfolio have higher PnL, with the Sharpe ratios maintaining a similar level.

*Value-weighted portfolio.* A value-weighted portfolio accounts for the liquidity in each asset and is defined by  $\omega_i^{(t)} = \min(\alpha v_{i,t}, \beta)$ , where  $v_{i,t} = \text{median}\{V_{i,t}, \dots, V_{i,1}\}$ , with  $V_{i,t}$  being the dollar volume traded for asset  $i$  on day  $t$ . The parameters  $\alpha$  and  $\beta$  control

Metric		Sharpe ratio			Mean std. of weights		
$\alpha \backslash \beta$		100,000	200,000	500,000	100,000	200,000	500,000
		0.001	1.14	1.05	1.00	0.00165	0.00203
0.005	1.34	1.26	1.14	0.00101	0.00124	0.00165	
0.010	<b>1.44</b>	1.34	1.24	0.00082	0.00101	0.00133	

**Table 5: Sharpe ratio and mean daily standard deviation of portfolio weights for different values of  $\alpha$  and  $\beta$ .**

Metric		Total bet size ( $\times 10^6$ )			Mean bet size per asset ( $\times 10^3$ )		
$\alpha \backslash \beta$		100,000	200,000	500,000	100,000	200,000	500,000
		0.001	23.7	32.2	41.5	36.6	49.7
0.005	40.8	67.7	118.5	63.0	104.5	182.9	
0.010	46.8	81.6	156.9	72.3	126.0	242.1	

**Table 6: Bet sizes per asset for different values of  $\alpha$  and  $\beta$ .**

Metric		Sharpe Ratio			Mean Daily PnL (bpts)		
Decile	Cost	0 bpts	1 bpt	2 bpts	0 bpts	1 bpt	2 bpts
		100	1.00	0.74	0.47	1.88	1.38
75	<b>1.07</b>	0.86	0.66	2.60	2.09	1.58	
50	1.05	0.89	0.73	3.30	2.78	2.27	
25	0.87	0.76	0.65	<b>3.84</b>	3.34	2.83	


**Table 7: Sharpe ratio and mean daily PnL (in basis points) for value weighted decile portfolios under different transaction costs ( $\alpha = 0.001$ ,  $\beta = 500,000$ ).**

the imbalance between portfolio weights: as  $\alpha$  increases and  $\beta$  decreases we get closer to the equal weighted portfolio setting. Table 5 shows that as  $\alpha$  increases and  $\beta$  decreases, the Sharpe ratio increases and portfolio variance decreases, respectively. The mean portfolio weight for every choice of  $(\alpha, \beta)$  was 0.00154. We also show the total book value and mean bet size in dollars for different values of  $(\alpha, \beta)$  in Table 6. We note that the largest book size considered is \$156.9 million which is much smaller than the total market. Therefore, market impact will be negligible. For fixed  $\alpha = 0.001$  and  $\beta = 500,000$ , Table 7 shows the Sharpe ratio and mean daily PnL for various decile portfolios and different transaction costs. Note that the choice of  $(\alpha, \beta)$  in Table 7 is the furthest from an equal-weighted portfolio out of all pairs of  $(\alpha, \beta)$  considered. Again, lower decile portfolios give larger PnL.

## 5.2 Intraday returns prediction

*Data description.* The minutely best-bid and ask prices for a universe of  $N = 516$  assets in the Nasdaq between 27/06/2007 and 31/12/2021 ( $T = 3656$ ) were extracted from the LOBSTER database<sup>1</sup>. From this, minutely mid-market prices were constructed. More specifically, 30-minute log returns were then computed, leading to 13 observations per day (overnight returns were not considered). The raw data contained some erroneous values which were dealt

<sup>1</sup>LOBSTER: Limit Order Book System - The Efficient Reconstructor at Humboldt Universität zu Berlin, Germany. <https://lobsterdata.com/>



**Figure 2: Cumulative PnL in bpts for three models backtest on 30 minutely intraday data.**

with by setting any outliers, defined as a thirty minutely return of over 25 percent, to zero.

*Task.* The task is to predict the next 30-minutes market excess returns. Again, we adopt a rolling window backtesting framework with a fixed look-back of six months. Unlike the daily returns application, however, we do not recompute the model parameters at each backtesting step due to computational expense. Instead, we recompute the factors, loadings, and covariance matrix of the estimated idiosyncratic component every month. The universe of assets considered each month is defined as the set of assets for which there exists data at every time point of the look-back window for that month. If an asset enters the market mid-month, we include it in our universe at the beginning of the following month. We compare the same three models as in the application to daily returns prediction in Section 5.1.

*Performance metrics.* As with the daily returns prediction, we consider the Sharpe ratio and mean daily PnL. Note, however, that in this case,  $PnL_t$  is the sum of the intra-day returns for day  $t$ . Figure 2 shows the cumulative PnL in bpts for the three prediction models under an equal weighted portfolio allocation with no transaction costs. FNIRVAR yielded a Sharpe ratio of 2.20, the LASSO estimated factor plus sparse VAR model gave a Sharpe ratio of 1.14, whilst the static factor model alone gave a Sharpe ratio of 0.58. A passive strategy whereby one waits for an order to be filled without crossing the spread is realistic in an intra-day setting. As such, we do not consider fixed transaction costs here. In fact, for a passive strategy, we are likely to earn half-spread each time our order is filled, though one needs to take into account adverse selection as well.

## 5.3 US Industrial Production

*Data description.* FRED-MD<sup>2</sup> [40] is a publicly accessible database of monthly observations of macroeconomic variables, updated in real-time. We choose the August 2022 vintage of the FRED-MD database which extends from January 1960 until December 2019 ( $T = 719$ ), preprocessed as discussed in [40]. Only variables with all observations in the time period are used ( $N = 122$ ).

*Task.* The task is one-step ahead prediction of the first-order difference of the logarithm of the monthly industrial production

<sup>2</sup>Data available at: <https://research.stlouisfed.org/econ/mccracken/fred-databases>.

	FNIRVAR	Factors Only	Factors + LASSO
Overall MSE	<b>0.0074</b>	0.0077	0.0077

**Table 8: Overall MSE for the task of forecasting US IP.**

(IP) index. We backtest each model from January 2000 until December 2019 using a rolling window with a look-back window of 480 observations. We set  $r = 8$  for the entire backtest as suggested by McCracken and Ng [40]. The order  $l_F$  of the factor VAR model changed over the course of the backtesting period, having a mean value of 4.17 over the period.

*Performance metrics.* Table 8 shows that using FNIRVAR reduces the overall MSE compared with a factor model with no idiosyncratic model. In contrast, combining a factor model with a LASSO estimated sparse VAR model provides no forecasting improvement over the factor model alone.

## 6 Conclusion and discussion

In this work, we introduced a factor-driven network-based VAR model called FNIRVAR for modelling the common and idiosyncratic components of a high dimensional multivariate time series. In FNIRVAR, a static factor model is combined with a network VAR model whose coefficient matrix is the weighted adjacency matrix of a SBM. The regularisation due to the network makes FNIRVAR an appropriate model for high dimensional settings. Under the assumption of strong factors and a large eigengap, we estimate the factors and loadings via PCA. NIRVAR estimation is employed on the remaining idiosyncratic component leading to a block-sparse estimated VAR coefficient matrix. As the FNIRVAR estimator does not require an observed network, it is especially suited to applications in finance and economics where human-defined networks may not be available. To test the predictive performance of FNIRVAR, we constructed a backtesting framework for three datasets of daily returns, intraday returns, and macroeconomic variables. In all three cases FNIRVAR out-performed a static factor model as well as a static factor plus LASSO estimated sparse VAR model. FNIRVAR showed particularly compelling results on the financial applications, achieving a Sharpe ratio of 1.95 for the task of daily returns prediction, and 2.20 for the task of thirty minutely intraday returns prediction. These results highlight the benefits of including sparse estimation along with dense factor modelling.

The FNIRVAR framework is especially suited to datasets containing groups of highly co-moving variables. From one perspective, the FNIRVAR estimator can be viewed as picking up block-weak factors [see 7, for a discussion of block-weak factors]. Theoretical analysis of the relation between FNIRVAR and weak factors is an interesting area for future research. Another future theoretical avenue is to investigate the asymptotic properties of the FNIRVAR estimator. Fan et al. [19] develop theory on combining dense and sparse time series models; the challenge for FNIRVAR is to show that the idiosyncratic model is consistent. This would require extending existing SBM results [for example, 10, 14, 35] to the setting of dependent data. Also, the static factor model of FNIRVAR could be modified to allow for a dynamic factor model. In this case, dynamic PCA would be required to estimate the factor space [see 6,

for example]. Furthermore, the assumption of strong factors and a large eigengap could be replaced by a sparsity assumption on the idiosyncratic covariance matrix, with the optimisation-based approach of Agarwal et al. [2] being used for estimation.

## Acknowledgments

BM acknowledges funding from the Engineering and Physical Sciences Research Council (EPSRC), grant number EP/S023151/1. FSP acknowledges funding from the EPSRC, grant no. EP/Y002113/1.

## References

- [1] Robert Adamek, Stephan Smeekes, and Ines Wilms. 2023. Lasso inference for high-dimensional time series. *Journal of Econometrics* 235, 2 (2023), 1114–1143.
- [2] Alekh Agarwal, Sahand Negahban, and Martin J Wainwright. 2012. Noisy matrix decomposition via convex relaxation: optimal rates in high dimensions. *The Annals of Statistics* 40, 2 (2012), 1171–1197.
- [3] Jushan Bai. 2003. Inferential theory for factor models of large dimensions. *Econometrica* 71, 1 (2003), 135–171.
- [4] Jushan Bai and Serena Ng. 2002. Determining the number of factors in approximate factor models. *Econometrica* 70, 1 (2002), 191–221.
- [5] Matteo Barigozzi, Giuseppe Cavaliere, and Graziano Moramarco. 2025. Factor network autoregressions. *Journal of Business & Economic Statistics* (2025), 1–14.
- [6] Matteo Barigozzi, Haeran Cho, and Dom Owens. 2024. FNETS: Factor-adjusted network estimation and forecasting for high-dimensional time series. *Journal of Business & Economic Statistics* 42, 3 (2024), 890–902.
- [7] Matteo Barigozzi and Marc Hallin. 2024. The Dynamic, the Static, and the Weak: Factor Models and the Analysis of High-Dimensional Time Series. *Journal of Time Series Analysis* (2024).
- [8] Sumanta Basu and George Michailidis. 2015. Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics* 43, 4 (2015), 1535–1567.
- [9] Peter J Brockwell and Richard A Davis. 1991. *Time series: theory and methods*. Springer science & business media.
- [10] J Cape, M Tang, and CE Priebe. 2019. Signal-plus-noise matrix models: eigenvector deviations and fluctuations. *Biometrika* 106, 1 (2019), 243–250.
- [11] Gary Chamberlain. 1983. Funds, factors, and diversification in arbitrage pricing models. *Econometrica: Journal of the Econometric Society* (1983), 1305–1323.
- [12] Gary Chamberlain and Michael Rothschild. 1983. Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets. *Econometrica: Journal of the Econometric Society* (1983), 1281–1304.
- [13] Elynn Y Chen, Jianqing Fan, and Xuening Zhu. 2023. Community network autoregression for high-dimensional time series. *Journal of Econometrics* 235, 2 (2023), 1239–1256.
- [14] Yuxin Chen, Yuejie Chi, Jianqing Fan, Cong Ma, et al. 2021. Spectral methods for data science: A statistical perspective. *Foundations and Trends® in Machine Learning* 14, 5 (2021), 566–806.
- [15] Victor Chernozhukov, Christian Hansen, and Yuan Liao. 2017. A lava attack on the recovery of sums of dense and sparse signals. *The Annals of Statistics* 45, 1 (2017), 39–76.
- [16] David Donoho, Matan Gavish, and Elad Romanov. 2023. ScreeNOT: Exact MSE-optimal singular value thresholding in correlated noise. *The Annals of Statistics* 51, 1 (2023), 122–148.
- [17] Eugene F Fama and Kenneth R French. 1992. The cross-section of expected stock returns. *the Journal of Finance* 47, 2 (1992), 427–465.
- [18] Jianqing Fan, Yuan Ke, and Kaizheng Wang. 2020. Factor-adjusted regularized model selection. *Journal of econometrics* 216, 1 (2020), 71–85.
- [19] Jianqing Fan, Ricardo P Masini, and Marcelo C Medeiros. 2023. Bridging factor and sparse models. *The Annals of Statistics* 51, 4 (2023), 1692–1717.
- [20] Mario Forni, Domenico Giannone, Marco Lippi, and Lucrezia Reichlin. 2009. Opening the black box: Structural factor models with large cross sections. *Econometric Theory* 25, 5 (2009), 1319–1347.
- [21] Mario Forni, Marc Hallin, Marco Lippi, and Lucrezia Reichlin. 2000. The generalized dynamic-factor model: Identification and estimation. *Review of Economics and Statistics* 82, 4 (2000), 540–554.
- [22] Mario Forni and Marco Lippi. 2001. The generalized dynamic factor model: representation theory. *Econometric theory* 17, 6 (2001), 1113–1141.
- [23] Ian Gallagher, Andrew Jones, Anna Bertiger, Carey E Priebe, and Patrick Rubin-Delanchy. 2024. Spectral embedding of weighted graphs. *J. Amer. Statist. Assoc.* 119, 547 (2024), 1923–1932.
- [24] John Geweke. 1977. The dynamic factor analysis of economic time series. *Latent variables in socio-economic models* (1977).
- [25] Domenico Giannone, Michele Lenza, and Giorgio E Primiceri. 2021. Economic predictions with big data: The illusion of sparsity. *Econometrica* 89, 5 (2021),

- 2409–2437.
- [26] Clive WJ Granger and Michael J Morris. 1976. Time series modelling and interpretation. *Journal of the Royal Statistical Society Series A: Statistics in Society* 139, 2 (1976), 246–257.
- [27] Shihao Gu, Bryan Kelly, and Dacheng Xiu. 2020. Empirical asset pricing via machine learning. *The Review of Financial Studies* 33, 5 (2020), 2223–2273.
- [28] Fang Han, Huanran Lu, and Han Liu. 2015. A direct estimation of high dimensional stationary vector autoregressions. *Journal of Machine Learning Research* (2015).
- [29] Peter D Hoff, Adrian E Raftery, and Mark S Handcock. 2002. Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* 97, 460 (2002), 1090–1098.
- [30] Marina Knight, Kathryn Leeming, Guy Nason, and Matthew Nunes. 2020. Generalized network autoregressive processes and the GNAR package. *Journal of Statistical Software* 96 (2020), 1–36.
- [31] Marina Iuliana Knight, MA Nunes, and GP Nason. 2016. Modelling, detrending and decorrelation of network time series. *arXiv preprint arXiv:1603.03221* (2016).
- [32] Jonas Krampe and Luca Margaritella. 2025. Factor Models With Sparse Vector Autoregressive Idiosyncratic Components. *Oxford Bulletin of Economics and Statistics* (2025).
- [33] Laurent Laloux, Pierre Cizeau, Marc Potters, and Jean-Philippe Bouchaud. 2000. Random matrix theory and financial correlations. *International Journal of Theoretical and Applied Finance* 3, 03 (2000), 391–397.
- [34] Clement Lee and Darren J Wilkinson. 2019. A review of stochastic block models and extensions for graph clustering. *Applied Network Science* 4, 1 (2019), 1–50.
- [35] Jing Lei and Alessandro Rinaldo. 2015. Consistency of spectral clustering in stochastic block models. *The Annals of Statistics* 43, 1 (2015), 215–237.
- [36] Helmut Lütkepohl. 1984. Linear transformations of vector ARMA processes. *Journal of Econometrics* 26, 3 (1984), 283–293.
- [37] Helmut Lütkepohl. 2005. *New introduction to multiple time series analysis*. Springer Science & Business Media.
- [38] Vladimir A Marčenko and Leonid Andreevich Pastur. 1967. Distribution of eigenvalues for some sets of random matrices. *Mathematics of the USSR-Sbornik* 1, 4 (1967), 457.
- [39] Brendan Martin, Francesco Sanna Passino, Mihai Cucuringu, and Alessandra Luati. 2024. NIRVAR: Network Informed Restricted Vector Autoregression. *arXiv preprint arXiv:2407.13314* (2024).
- [40] Michael W McCracken and Serena Ng. 2016. FRED-MD: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics* 34, 4 (2016), 574–589.
- [41] Ke Miao, Peter CB Phillips, and Liangjun Su. 2023. High-dimensional VARs with common factors. *Journal of Econometrics* 233, 1 (2023), 155–183.
- [42] Mark EJ Newman. 2003. Mixing patterns in networks. *Physical review E* 67, 2 (2003), 026126.
- [43] William B Nicholson, Ines Wilms, Jacob Bien, and David S Matteson. 2020. High dimensional forecasting via interpretable vector autoregression. *Journal of Machine Learning Research* 21, 166 (2020), 1–52.
- [44] Thomas J Sargent, Christopher A Sims, et al. 1977. Business cycle modeling without pretending to have too much a priori economic theory. *New methods in business cycle research* 1 (1977), 145–168.
- [45] James H Stock and Mark W Watson. 2002. Forecasting using principal components from a large number of predictors. *J. Amer. Statist. Assoc.* 97, 460 (2002), 1167–1179.
- [46] James H Stock and Mark W Watson. 2002. Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics* 20, 2 (2002), 147–162.
- [47] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 58, 1 (1996), 267–288.
- [48] Ruey S Tsay. 2013. *Multivariate time series analysis: with R and financial applications*. John Wiley & Sons.
- [49] Ganapathy Vidyamurthy. 2004. *Pairs Trading: quantitative methods and analysis*. John Wiley & Sons.
- [50] Hang Yin, Abolfazl Safikhani, and George Michailidis. 2023. A general modeling framework for network autoregressive processes. *Technometrics* 65, 4 (2023), 579–589.
- [51] Mu Zhu and Ali Ghodsi. 2006. Automatic dimensionality selection from the scree plot via the use of profile likelihood. *Computational Statistics & Data Analysis* 51, 2 (2006), 918–930.
- [52] X Zhu, R Pan, G Li, Y Liu, and H Wang. 2017. Network vector autoregression. *The Annals of Statistics* 45, 3 (2017), 1096–1123.