

Sparse transition matrix estimation for high-dimensional and locally stationary vector autoregressive models*

Xin Ding

University of British Columbia

e-mail: xin.ding@stat.ubc.ca

Ziyi Qiu

University of Illinois at Urbana-Champaign

University of Chicago

e-mail: ziyiqiu@illinois.edu

and

Xiaohui Chen

University of Illinois at Urbana-Champaign

e-mail: xhchen@illinois.edu

Abstract: We consider the estimation of the transition matrix in the high-dimensional time-varying vector autoregression (TV-VAR) models. Our model builds on a general class of locally stationary VAR processes that evolve smoothly in time. We propose a hybridized kernel smoothing and ℓ^1 -regularized method to directly estimate the sequence of time-varying transition matrices. Under the sparsity assumption on the transition matrix, we establish the rate of convergence of the proposed estimator and show that the convergence rate depends on the smoothness of the locally stationary VAR processes only through the smoothness of the transition matrix function. In addition, for our estimator followed by thresholding, we prove that the false positive rate (type I error) and false negative rate (type II error) in the pattern recovery can asymptotically vanish in the presence of weak signals without assuming the minimum nonzero signal strength condition. Favorable finite sample performances over the ℓ^2 -penalized least-squares estimator and the unstructured maximum likelihood estimator are shown on simulated data. We also provide two real examples on estimating the dependence structures on financial stock prices and economic exchange rates datasets.

MSC 2010 subject classifications: Primary 62M10, 62H12; secondary 91B84.

Keywords and phrases: Vector autoregression, time-varying parameters, locally stationary processes, kernel smoothing, high-dimension, sparsity.

Received November 2016.

*Research partially supported by NSF grant DMS-1404891 and UIUC Research Board Award RB15004.

Contents

| | | |
|-------|---|------|
| 1 | Introduction | 3872 |
| 2 | Method and algorithm | 3874 |
| 3 | Asymptotic properties | 3876 |
| 3.1 | Locally stationary VAR process in L^2 | 3876 |
| 3.2 | Rate of convergence | 3877 |
| 3.3 | Pattern recovery | 3879 |
| 4 | Simulation studies | 3881 |
| 4.1 | Data generation | 3881 |
| 4.2 | Tuning parameter selection | 3881 |
| 4.3 | Comparison results | 3882 |
| 4.3.1 | Estimation errors | 3882 |
| 4.3.2 | Pattern recovery | 3882 |
| 5 | Real data analysis | 3887 |
| 5.1 | Finance data: stock prices | 3887 |
| 5.2 | Economic data: exchange rates | 3889 |
| 6 | Proofs | 3891 |
| A | Detailed setup of the simulation studies | 3897 |
| B | Examples for four graph structures | 3898 |
| C | ROC curves based on the true thresholding parameters $u_{\#}$ | 3899 |
| D | Selected stocks and currencies in real data analysis | 3900 |
| | Acknowledgments | 3900 |
| | References | 3900 |

1. Introduction

Vector autoregression (VAR) is a basic tool in multivariate time series analysis and it has been extensively used to model the cross-sectional and serial dependence in various applications from economics and finance [31, 3, 4, 10, 32, 14, 34, 13, 1]. There are two fundamental limitations of the vector-autoregressive models. First, conventional methods to estimate the transition matrix of a VAR model are based on the least squares (LS) estimator and the maximum likelihood estimator (MLE), in which the parameter estimation is consistent when the sample size increases and the model size is *fixed* [8, 25]. Since the number of parameters grows quadratically in the number of time series variables, the VAR model typically includes no more than ten variables in many real applications [18]. However, due to the recent explosive data enrichment, analysis of panel data with a few hundred variables is often encountered, in which the LS and MLE are not suitable even for a moderate problem size. Second, stationarity plays a major role in the VAR model. Therefore, the stationary VAR does not capture the time-varying underlying data generation structures, which have been observed in a broad regime of applications in economics and finance [2, 27].

Motivated by the limitations of the VAR, this paper studies the estimation problems of the time-varying VAR (TV-VAR) model for high-dimensional time

series data. Let $\mathbf{X}_{d \times n} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ be a sequence of d -dimensional observations generated by a mean-zero TV-VAR of order 1 (TV-VAR(1))

$$\mathbf{x}_i = A(i/n)\mathbf{x}_{i-1} + \mathbf{e}_i, \tag{1.1}$$

where $A(t), t \in [0, 1]$, is a $d \times d$ matrix-valued deterministic function consisting of the transition matrices $A_i := A(i/n)$ at evenly spaced time points and \mathbf{e}_i are independent and identically distributed (iid) mean-zero random errors, i.e. innovations. In this paper, our main focus is to estimate the transition matrices A_i for the TV-VAR(1) and the extension to higher-order VAR is straightforward. Indeed, for a general TV-VAR of order $k \geq 1$

$$\mathbf{x}_i = A_{i,1}\mathbf{x}_{i-1} + A_{i,2}\mathbf{x}_{i-2} + \dots + A_{i,k}\mathbf{x}_{i-k} + \mathbf{e}_i,$$

we can rewrite $\mathbf{z}_i = (\mathbf{x}_i^\top, \dots, \mathbf{x}_{i-k+1}^\top)^\top$ at time $i = k, k + 1, \dots, n$, as a TV-VAR(1) in the augmented space

$$\mathbf{z}_i = \mathbf{A}_i\mathbf{z}_{i-1} + \tilde{\mathbf{e}}_i,$$

where

$$\mathbf{A}_i = \begin{pmatrix} A_{i,1} & A_{i,2} & \dots & A_{i,k-1} & A_{i,k} \\ I_{d \times d} & \mathbf{0}_{d \times d} & \dots & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & I_{d \times d} & \dots & \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} & \dots & I_{d \times d} & \mathbf{0}_{d \times d} \end{pmatrix}, \quad \tilde{\mathbf{e}}_i = \begin{pmatrix} \mathbf{e}_i \\ \mathbf{0}_{d \times 1} \\ \vdots \\ \mathbf{0}_{d \times 1} \end{pmatrix},$$

$I_{d \times d}$ is the $d \times d$ identity matrix and $\mathbf{0}_{d \times d'}$ is the $d \times d'$ zero matrix. Then, we need only to estimate the first d rows in \mathbf{A}_i .

To make the estimation problem feasible for high-dimensional time series when d is large, it is crucial to carefully regularize the coefficient matrices A_i . The main idea is to use certain low-dimensional structures in A_i and the degree of freedom of (1.1) can be dramatically reduced. In our problem, we assume two key structures in A_i . First, since our goal is to estimate a sequence of transition matrices, which can be viewed as the discrete version of a matrix-valued function, it is natural to impose the *smoothness* condition to $A(t)$. In this case, (1.1) is closely related to the *locally stationary processes*, a general class of non-stationary processes proposed by [15]. Examples of other linear and nonlinear locally stationary processes can be found in [11]. In particular, let $i_0 = 1, \dots, n$, be any time point and $t_0 = i_0/n$. Then, under suitable regularity conditions [16], there exists a stationary process $\tilde{\mathbf{x}}_i(t_0)$ such that for all $j = 1, \dots, d$,

$$|X_{ij} - \tilde{X}_{ij}(t_0)| = O_{\mathbb{P}}(|i/n - t_0| + 1/n),$$

where X_{ij} and $\tilde{X}_{ij}(t_0)$ are the j -th element of \mathbf{x}_i and $\tilde{\mathbf{x}}_i(t_0)$ respectively, and

$$\tilde{\mathbf{x}}_i(t_0) = A(t_0)\tilde{\mathbf{x}}_{i-1}(t_0) + \mathbf{e}_i.$$

Therefore, \mathbf{x}_i is an approximately stationary VAR process $\tilde{\mathbf{x}}_i(t_0)$ in a small neighborhood of t_0 .

Second, at each time point $i = 1, \dots, n$, we need to estimate a $d \times d$ matrix. There have been different structural options in literature, such as the sparse [21, 35, 33, 30], banded [19], and low-rank [22] transition matrix, all of which only considered the stationary VAR model. In this paper, we consider a sequence of *sparse* transition matrices and we allow the sparsity patterns to change over time. Note that our problem is also different from the emerging literature for estimating the high-dimensional covariance matrix and its related functionals of time series data [11, 12, 37, 5], since our goal is to directly estimate the data generation mechanism specified by $A(\cdot)$.

To simultaneously address the two issues, we propose a hybridized method of the nonparametric smoothing technique and ℓ^1 -regularization to estimate the sparse transition matrix of the locally stationary VAR. The proposed method is equivalent to solving a sequence of large numbers of linear programs and therefore the estimates can be efficiently obtained by using the high-performance (i.e. parallel) computing technology; see Section 2 for details. In Section 3, we establish the rate of convergence under suitable assumptions on the smoothness of the covariance function and the sparsity of the transition matrix function $A(\cdot)$. Specifically, the dimension d is permitted to increase sub-exponentially fast in the sample size n to obtain consistent estimation, i.e. $d = o(\exp(n))$. In addition, we also prove that when our estimator is followed by thresholding, type I and type II errors in the pattern recovery asymptotically vanish in the presence of weak signals. In contrast with the existing literature on consistent model selection, we do not require the minimum nonzero signal strength to be bounded away from zero. Simulation studies in Section 4 and two real data examples in Section 5 demonstrate favorable performances and more interpretable results of the proposed TV-VAR model. Technical proofs are deferred to Section 6.

We fix the notations that will be used in the rest of the paper. Let M be a generic matrix, $\mathcal{I} \subset \mathbb{N}$ is an index subset, $[M]_{*\mathcal{I}}$ and $[M]_{\mathcal{I}*}$ be the submatrix of M with columns and rows indexed by \mathcal{I} , resp. Denote $J = \{1, \dots, d\}$. Let $\rho(M) = \sup_{|\mathbf{v}|=1} |M\mathbf{v}|$ be the spectral norm of M . $|M|_1$, $|M|_\infty$ and $|M|_F$ are the entry ℓ^1 , ℓ^∞ and Frobenius norm of M , resp. $|M|_{\ell^1} = \max_{j \leq d} \sum_{k=1}^d |M_{jk}|$ and $|M|_{\ell^\infty} = \max_{k \leq d} \sum_{j=1}^d |M_{jk}|$ are the matrix ℓ^1 and ℓ^∞ norm of M , resp. $\text{supp}(M) = \{(\ell, j) : M_{\ell j} \neq 0\}$ is the support of M and $|S|$ is the number of nonzeros in the support set S . For $q > 0$ and a random variable (r.v.) X , $\|X\|_q = (\mathbb{E}|X|^q)^{1/q}$ and we say that $X \in \mathcal{L}^q$ iff $\|X\|_q < \infty$. Write $\|X\| = \|X\|_2$. Denote $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$. If a and b are vectors, then the maximum and minimum are interpreted element-wise.

2. Method and algorithm

For $m \in \mathbb{Z}$, let $\Sigma_{i,m} = \text{Cov}(\mathbf{x}_i, \mathbf{x}_{i+m}) = \mathbb{E}(\mathbf{x}_i \mathbf{x}_{i+m}^\top)$. Then

$$\Sigma_{i-1,1} = \Sigma_{i-1,0} A_i^\top \quad \text{and} \quad \Sigma_{i,-1} = A_i \Sigma_{i-1,0}. \quad (2.1)$$

Therefore, for any estimator, say \hat{A}_i , that is reasonably close to the true coefficient matrices A_i , we must have $(\Sigma_{i-1,1} - \Sigma_{i-1,0}\hat{A}_i^\top) \approx \mathbf{0}_{d \times d}$ and $(\Sigma_{i,-1} - \hat{A}_i\Sigma_{i-1,0}) \approx \mathbf{0}_{d \times d}$. A naive estimator for A_i would be constructed to invert the sample versions of (2.1). Estimators of this kind do not have good statistical properties in high-dimensions because of the ill-conditioning and the dependence information in A_i is not directly used in estimation. If A_i is known to be sparse *a priori*, we may consider the following constrained minimization program

$$\begin{aligned} &\text{minimize}_{\Lambda \in \mathbb{R}^{d \times d}} && |\Lambda|_1 && (2.2) \\ &\text{subject to} && |\Sigma_{i-1,1} - \Sigma_{i-1,0}\Lambda|_\infty \leq \tau \\ &&& |\Sigma_{i,-1} - \Lambda^\top \Sigma_{i-1,0}|_\infty \leq \tau. \end{aligned}$$

Because $\Sigma_{i,1}$ and $\Sigma_{i,0}$ are unknown, the solution of (2.2) is an *oracle* estimator and therefore it is not implementable in practice.

Let $A(\cdot)$ be the continuous version of $A_i, i = 1, \dots, n$, in (1.1), and fix a $t \in (0, 1)$. To estimate $A(t)$, we first estimate $\Sigma_{i,1}$ and $\Sigma_{i,0}$ in (2.1) by their empirical versions. Let

$$\hat{\Sigma}_{t,\ell} = \sum_{m=1}^n w(t, m) \mathbf{x}_m \mathbf{x}_{m+\ell}^\top, \quad \ell = -1, 0, 1, \tag{2.3}$$

be the kernel smoothed estimator of $\Sigma_{i,0}, \Sigma_{i,1}$ and $\Sigma_{i,-1}$, where $w(t, m)$ are the nonnegative weights. Here, we consider the Nadaraya-Watson smoothed estimator

$$w(t, m) = \frac{K(\frac{t-m/n}{b_n})}{\sum_{\ell=1}^n K(\frac{t-\ell/n}{b_n})}, \tag{2.4}$$

where $K(\cdot)$ is a nonnegative bounded kernel function with support in $[-1, 1]$ such that $\int_{-1}^1 K(v)dv = 1$ and b_n is the bandwidth parameter. We assume throughout the paper that the bandwidth satisfies the natural conditions: $b_n = o(1)$ and $n^{-1} = o(b_n)$. Then, our estimator $\hat{A}(t)$ is defined as the transpose of the solution of

$$\begin{aligned} &\text{minimize}_{\Lambda \in \mathbb{R}^{d \times d}} && |\Lambda|_1 && (2.5) \\ &\text{subject to} && |\hat{\Sigma}_{i-1,1} - \hat{\Sigma}_{i-1,0}\Lambda|_\infty \leq \tau \\ &&& |\hat{\Sigma}_{i,-1} - \Lambda^\top \hat{\Sigma}_{i-1,0}|_\infty \leq \tau. \end{aligned}$$

Observe that (2.5) is equivalent to the following d optimization sub-problems

$$\begin{aligned} \hat{\beta}_j(t) &= \arg \min_{\mathbf{u} \in \mathbb{R}^d} |\mathbf{u}|_1 && (2.6) \\ \text{subject to} &&& |[\hat{\Sigma}_{i-1,1}]_{*j} - \hat{\Sigma}_{i-1,0}\mathbf{u}|_\infty \leq \tau \\ &&& |[\hat{\Sigma}_{i,-1}]_{j*} - \mathbf{u}^\top \hat{\Sigma}_{i-1,0}|_\infty \leq \tau \end{aligned}$$

in that $[\hat{A}(t)]_{j*} = \hat{\beta}_j(t)^\top$. Since the d sub-problems (2.6) can be independently solved, we can efficiently compute the solution of (2.5) by parallelizing the optimizations of (2.6). In addition, each sub-problem in (2.6) can be recasted as a

linear program (LP)

$$\begin{aligned} & \text{minimize}_{\mathbf{v}, \mathbf{u} \in \mathbb{R}_+^d} && |\mathbf{v}|_1 + |\mathbf{w}|_1 \\ & \text{subject to} && -\hat{\Sigma}_{i-1,0}\mathbf{v} + \hat{\Sigma}_{i-1,0}\mathbf{w} \leq \tau - \max\{[\hat{\Sigma}_{i-1,1}]_{*j}, [\hat{\Sigma}_{i-1,-1}]_{*j}^\top\} \\ & && \hat{\Sigma}_{i,0}\mathbf{v} - \hat{\Sigma}_{i-1,0}\mathbf{w} \leq \tau + \min\{[\hat{\Sigma}_{i,1}]_{*j}, [\hat{\Sigma}_{i,-1}]_{*j}^\top\} \end{aligned}$$

so that $\hat{\beta}_j(t) = \hat{\mathbf{v}} - \hat{\mathbf{w}}$, where $(\hat{\mathbf{v}}, \hat{\mathbf{w}})$ is the nonnegative solution of the LP. Note that (2.6) can be efficiently solved by the simplex algorithm [26]. Compared with the estimation of sparse transition matrix for the stationary VAR model [21], our estimator requires to solve two sample versions of Yule-Walker equations in (2.1) and (2.5) due to the non-stationarity and therefore two-sided constraints of the estimator $\hat{A}(t)$ are needed in (2.5).

3. Asymptotic properties

3.1. Locally stationary VAR process in L^2

To establish an asymptotic theory for estimating the continuous matrix-valued transition function $A(\cdot)$ of the non-stationary VAR, it is more convenient to model the time series as realizations from a continuous d -dimensional process. Following the framework of [17], $\mathbf{x}_1, \dots, \mathbf{x}_n$ are viewed as realizations of the continuous d -dimensional process $\mathbf{x}(t) = (X_1(t), \dots, X_d(t))^\top$, $t \in [0, 1]$, on the discrete grid $t_i = i/n$, $i = 1, \dots, n$, i.e. $\mathbf{x}(i/n) = A(i/n)\mathbf{x}((i-1)/n) + \mathbf{e}(i/n)$ and the temporal dependence of \mathbf{x}_i is carried out on the rescaled time indices t_i .

Definition 3.1 (Locally stationary VAR process in L^2). A mean-zero non-stationary VAR process $\{\mathbf{x}_i; i = 1, \dots, n\}$ in \mathbb{R}^d of the form (1.1) is said to be *locally stationary in L^2* , or *weakly locally stationary*, if for every $i = 1, \dots, n$ and $t = i/n$, there exists a mean-zero stationary VAR process $\{\tilde{\mathbf{x}}_m(t); m = 1, \dots, n\}$ given by

$$\tilde{\mathbf{x}}_m(t) = A(t)\tilde{\mathbf{x}}_{m-1}(t) + \mathbf{e}_m \quad (3.1)$$

such that for all $m = 1, \dots, n$,

$$\max_{1 \leq j \leq d} \|\tilde{X}_{mj}(t) - X_{mj}\| \leq C \left(\left| \frac{m}{n} - t \right| + \frac{1}{n} \right) \quad (3.2)$$

and the constant $C > 0$ does not depend on d, m , and n .

Note that the approximating stationary VAR process in Definition 3.1 generally depends on t . As suggested by the Yule-Walker equations (2.1), given a fixed time of interest, estimation of $A(t)$ relies on an estimate of $\Sigma_{t,\ell}$. Consider $\ell = 0$ and let $\Sigma(t) = \Sigma_{t,0}$ be the matrix-valued covariance function at lag-zero. With a finite number of observations $\mathbf{x}_1, \dots, \mathbf{x}_n$, it is unclear how to define the covariance function $\Sigma(\cdot)$ off the n points $t_i = i/n$. Nevertheless, the weakly locally stationary VAR processes provide a natural framework for extending $\Sigma(t_i)$ to $\Sigma(t)$ for all $t \in (0, 1)$.

Since $A(\cdot)$ is continuous on $[0, 1]$, the stationary VAR process in (3.1) is defined for all $t \in (0, 1)$. Let $\tilde{\Sigma}(t) = \mathbb{E}[\tilde{\mathbf{x}}_m(t)\tilde{\mathbf{x}}_m(t)^\top]$. By (3.2) and the Cauchy-Schwarz inequality, we have for each $j, k = 1, \dots, d$,

$$\begin{aligned} |\sigma_{jk}(t_i) - \tilde{\sigma}_{jk}(t_i)| &= |\mathbb{E}(X_{ij}X_{ik}) - \mathbb{E}(\tilde{X}_{ij}(t_i)\tilde{X}_{ik}(t_i))| \\ &\leq \|X_{ij}\|_2 \cdot \|X_{ik} - \tilde{X}_{ik}(t_i)\| + \|\tilde{X}_{ik}(t_i)\| \cdot \|X_{ij} - \tilde{X}_{ij}(t_i)\| \\ &\leq (\|X_{ij}\|_2 + \|\tilde{X}_{ik}(t_i)\|) \max_{1 \leq j \leq d} \|X_{ij} - \tilde{X}_{ij}(t_i)\| \\ &\leq Cn^{-1}, \end{aligned}$$

where the constant C here is uniform in $j, k = 1, \dots, d$ and $i = 1, \dots, n$. Therefore, we get

$$\Sigma(t_i) = \tilde{\Sigma}(t_i) + O(n^{-1}).$$

Letting $n \rightarrow \infty$ and by the continuity of $A(\cdot)$, we can extend $\Sigma(t) = \tilde{\Sigma}(t)$ for all $t \in (0, 1)$. Similar extension can be done for $\Sigma_{t,\ell}$ for $\ell = \pm 1$. In Section 3.2, it will be shown that the asymptotic theory of estimating $A(t), t \in (0, 1)$ depends on the smoothness of the weakly locally stationary VAR processes only through the smoothness of $\Sigma(t)$ and therefore $A(t)$.

3.2. Rate of convergence

In this section, we characterize the rate of convergence of our estimator (2.5) under various matrix norms. We assume $d \geq 2$. To study the asymptotic properties of the proposed estimator, we make the following assumptions on model (1.1).

1. The coefficient matrices are sparse: let $0 \leq \alpha < 1$, for each $i = 1, \dots, n$,

$$A_i \in \mathcal{G}_\alpha(s, M_d) = \mathcal{G}_\alpha^+(s, M_d) \cap \mathcal{G}_\alpha^-(s, M_d), \tag{3.3}$$

where

$$\begin{aligned} \mathcal{G}_\alpha^+(s, M_d) &= \left\{ \Theta \in \mathbb{R}^{d \times d} : \max_{j \leq d} \sum_{\ell=1}^d |\Theta_{\ell j}|^\alpha \leq s, |\Theta|_{\ell^1} \leq M_d \right\}, \\ \mathcal{G}_\alpha^-(s, M_d) &= \left\{ \Theta \in \mathbb{R}^{d \times d} : \max_{\ell \leq d} \sum_{j=1}^d |\Theta_{\ell j}|^\alpha \leq s, |\Theta|_{\ell^\infty} \leq M_d \right\}. \end{aligned}$$

2. The marginal and lag-one covariance matrix processes $\{\Sigma_0(t)\}_{t \in [0,1]}$ and $\{\Sigma_1(t)\}_{t \in [0,1]}$ are smooth functions in $\mathbb{R}^{d \times d}$: $\Sigma_{\ell,jk}(t) \in \mathcal{C}^2([0, 1]), \ell = 0, 1$, where $\mathcal{C}^2([0, 1])$ is the class of functions defined on $[0, 1]$ that are twice differentiable with bounded derivatives uniformly in $j, k = 1, \dots, d$.
3. The random innovations $\mathbf{e}_i = (e_{i1}, \dots, e_{id})^\top$ have iid components and sub-Gaussian tail: $\mathbb{P}(|e_{ij}| \geq x) \leq Ce^{-cx^2}$ for all $x > 0$.

Before proceeding, we discuss the above assumptions. (3.3) requires that the transition matrices are sparse in both columns and rows at all time points. A similar matrix class defined by (3.3) was first proposed in [7] for *symmetric* matrices and it has been widely used for estimating high-dimensional covariance and precision matrix; see e.g. [9, 11]. If $\alpha = 0$, then the maximum number of nonzeros in columns and rows of A_i is at most s . Assumption 2 requires that the marginal and lag-one covariance matrices evolve smoothly in time. The smoothness is not defined directly on $A(\cdot)$ for the ease of theorem statements. In view of (3.2), Assumption 2 is implied by the smoothness on A_i under extra regularity conditions. For a generic matrix $M(t)$ parameterized by $t \in (0, 1)$, we let $\dot{M}(t)$ and $\ddot{M}(t)$ be the first two element-wise derivatives of $M(t)$ w.r.t. t .

Lemma 3.1. *Assume that $\sup_{t \in [0,1]} |A(t)|_{\ell^1} < 1$. Then we have*

$$|\dot{\Sigma}(t)|_{\infty} \leq \frac{2|\dot{A}(t)|_{\infty} \cdot |A(t)|_{\ell^1} \cdot |\Sigma(t)|_{\ell^1}}{1 - |A(t)|_{\ell^1}^2} \quad (3.4)$$

and

$$\begin{aligned} |\ddot{\Sigma}(t)|_{\infty} &\leq \frac{8|\dot{A}(t)|_{\infty} \cdot |A(t)|_{\ell^1}^2 \cdot |\dot{A}(t)|_{\ell^1} \cdot |\Sigma(t)|_{\ell^1}}{(1 - |A(t)|_{\ell^1}^2)^2} \\ &+ \frac{4 \max\{|\dot{A}(t)|_{\infty}, |\ddot{A}(t)|_{\infty}\} \cdot \max\{|A(t)|_{\ell^1}, |\dot{A}(t)|_{\ell^1}\} \cdot |\Sigma(t)|_{\ell^1}}{1 - |A(t)|_{\ell^1}^2}. \end{aligned} \quad (3.5)$$

By Lemma 3.1, if $\sup_{t \in [0,1]} |A(t)|_{\ell^1} < 1$ and $A_{jk}(\cdot) \in \mathcal{C}^2([0, 1])$ such that $\sup_{t \in [0,1]} |\dot{A}(t)|_{\ell^1} \vee |\Sigma(t)|_{\ell^1} \leq K$ for some constant K , then $\Sigma_{jk}(\cdot)$ is also a $\mathcal{C}^2([0, 1])$ function and therefore Assumption 2 is fulfilled.

Assumption 3 specifies the tail probability of the innovations \mathbf{e}_i . In [21], \mathbf{e}_i 's follow iid $N(\mathbf{0}, \Psi)$ for some error covariance matrix Ψ . A simple transformation by $\Psi^{-1/2}$ will reduce to the case that \mathbf{e}_i has iid components with the standard normal distribution, a special case of Assumption 3, which covers the sub-Gaussian innovations.

Theorem 3.2. *Let $\bar{\rho} = \sup_{i \geq 0} \rho(A_i)$. Fix an $i \in [nb_n + 1, n(1 - b_n) + 1]$. Suppose that Assumption 1, 2, 3 are satisfied, $\bar{\rho} < 1$ and $A_i \in \mathcal{G}_{\alpha}(s, M_d)$. Then, with probability at least $1 - 2d^{-1}$, we have the estimator \hat{A}_i with tuning parameter $\tau \geq C(1 + M_d)(b_n^2 + \sqrt{(\log d)/(nb_n)})$ obeys*

$$|\hat{A}_i - A_i|_{\infty} \leq 2\tau |\Sigma_{i-1,0}^{-1}|_{\ell^1}, \quad (3.6)$$

$$\rho(\hat{A}_i - A_i) \leq C(\alpha) s (\tau |\Sigma_{i-1,0}^{-1}|_{\ell^1})^{1-\alpha}, \quad (3.7)$$

$$d^{-1} |\hat{A}_i - A_i|_F^2 \leq C(\alpha) s (\tau |\Sigma_{i-1,0}^{-1}|_{\ell^1})^{2-\alpha}. \quad (3.8)$$

From Theorem 3.2, the bandwidth $b_n \asymp ((\log d)/n)^{-1/5}$ gives the optimal rate of convergence and the resulting tuning parameter $\tau \geq C(1 + M_d)((\log d)/n)^{-2/5}$. When d is fixed, it is known that the optimal bandwidth in the nonparametric kernel density estimation is $n^{-1/5}$ for twice continuously differentiable functions under the mean integrated square error. So in the high-dimensional context, the dimension only has a logarithm impact on the choice of optimal bandwidth.

3.3. Pattern recovery

We also study the recovery of time-varying patterns using the estimator (2.5). Let $S_i = \text{supp}(A_i)$ be the nonzero positions of A_i . If the nonzero entries of A_i are small enough, then it is impossible to accurately distinguish the small nonzeros and the zeros. Therefore, the best hope we can expect is that the nonzero entries of A_i with large magnitudes can be well separated from the zeros in A_i . Let

$$u_{\#} = 2\tau |\Sigma_{i-1,0}^{-1}|_{\ell^1}, \tag{3.9}$$

where τ is determined in Theorem 3.2. We use the thresholded version of (2.5) as an estimator of S_i

$$\hat{S}_i = T_{u_{\#}}(\hat{A}_i) = \{\mathbf{1}(|\hat{A}_{i,jk}| > u_{\#})\}_{j,k=1}^d. \tag{3.10}$$

By Theorem 3.2, the maximal fluctuation $|\hat{A}_i - A_i|_{\infty}$ is controlled by $u_{\#}$ with probability at least $1 - 2d^{-1}$. So if we apply a thresholding procedure for \hat{A}_i at the level $u_{\#}$, then we expect that the zeros and non-zeros with magnitudes larger than $2u_{\#}$ in A_i can be identified by the thresholded support of \hat{A}_i . Precisely, we have the instantaneous *partial* recovery consistency.

Theorem 3.3. *Assume that Assumption 1,2,3 are satisfied. Then, for any fixed time points $i \in [nb_n + 1, n(1 - b_n) + 1]$, $\mathbb{P}(\hat{S}_i \subset S_i) \geq 1 - 2d^{-1}$ and $\mathbb{P}(\{(m, k) : |A_{i,mk}| > 2u_{\#}\} \subset \hat{S}_i) \geq 1 - 2d^{-1}$.*

Theorem 3.3 states that, with high probability, the zeros in A_i can be identified and the nonzero entries in A_i with strong signal strength above $2u_{\#}$ can be recovered by \hat{S}_i . Therefore, the false positives (type I error) of the estimator (3.10) are asymptotically controlled; see Theorem 3.5 for precise statement. However, Theorem 3.3 does not provide too much information regarding the false negatives since there is no characterization of the signal strength in $(0, 2u_{\#})$.

Let $\beta > 0$ and $u_0 \in (0, 1)$. We introduce the following $d \times d$ matrix class

$$\begin{aligned} &\mathcal{G}_{\alpha,\beta}(s, M_d, L_d) \\ = &\mathcal{G}_{\alpha}(s, M_d) \cap \left\{ \Theta : \frac{|\{(m, k) : |\Theta_{mk}| \in (0, u)\}|}{|\text{supp}(\Theta)|} \leq L_d u^{\beta}, \forall 0 < u < u_0 \right\}. \end{aligned} \tag{3.11}$$

For $A \in \mathcal{G}_{\alpha,\beta}(s, M_d, L_d)$, the parameters β and L_d control the proportion of small entries in the support of A . If β is large and L_d grows slowly, then the fraction of weak signals in A is small and therefore the false negatives (type II error) can also be well controlled. Below, we shall give such an example.

Example 3.1 (A spatial design). Let $0 < r < 1$. Consider a $d \times d$ symmetric matrix $A = (A_{mk})_{d \times d}$ that is generated by the covariance function of a spatial process Z_1, Z_2, \dots, Z_d , which is a random vector that is observed at sites $h_1^{\circ}, \dots, h_d^{\circ} \in D$ in some spatial domain D . Assume that the covariance between Z_m and $Z_k, m, k = 1, \dots, d$ satisfies

$$A_{mk} := \text{cov}(Z_m, Z_k) = \begin{cases} f(\pi(h_m^{\circ}, h_k^{\circ})) & \pi(h_m^{\circ}, h_k^{\circ}) < \frac{d^{1-r}}{2} \\ 0 & \pi(h_m^{\circ}, h_k^{\circ}) \geq \frac{d^{1-r}}{2} \end{cases}, \tag{3.12}$$

where $\pi(h_m^\circ, h_k^\circ)$ is the distance between the sites h_m° and h_k°

$$\pi(h_m^\circ, h_k^\circ) = \frac{|m - k|}{d^r} \quad (3.13)$$

and f is a real-value covariance function. Here, we consider the rational quadratic covariance function [11, 36]

$$f(\pi(h_m^\circ, h_k^\circ)) = \frac{1}{(1 + \pi(h_m^\circ, h_k^\circ)^2)^\gamma}, \quad \gamma > 1. \quad (3.14)$$

In this example, $A_{mk} = 0$ if $|m - k| \geq d/2$, so A is a banded matrix.

Lemma 3.4. For A in Example 3.1, then there exists a large enough constant $C > 0$ depending only on r, α, γ such that $A \in \mathcal{G}_{\alpha, \beta}(s, M_d, L_d)$, where

$$s = \begin{cases} Cd^r \log d & \text{if } 2\gamma\alpha = 1 \\ Cd^r & \text{if } 2\gamma\alpha > 1 \\ Cd^r d^{(1-r)(1-2\gamma\alpha)} & \text{if } 0 < 2\gamma\alpha < 1 \end{cases}, \quad (3.15)$$

$$M_d = \begin{cases} Cd^r \log d & \text{if } \gamma = 1/2 \\ Cd^r & \text{if } \gamma > 1/2 \\ Cd^r d^{(1-r)(1-2\gamma)} & \text{if } 0 < \gamma < 1/2 \end{cases}, \quad (3.16)$$

and $L_d = Cd^{2(1-r)\gamma\beta}$.

For any fixed distance parameter $r \in (0, 1)$, the weak signal parameter L_d has a natural dependence on γ . If γ is smaller, then the covariance function f decays to zero slower and there is a less fraction of weak signals in A . This can allow L_d to grow slowly in d . Note that class $\mathcal{G}_{\alpha, \beta}(s, M_d, L_d)$ is much less stringent than the widely used condition for support recovery and model selection in literature, which requires that the minimal nonzero signal strength is uniformly bounded away from zero [28]. To quantify the error in the pattern recovery, we use the following two error rate measures.

Definition 3.2. The false positive rate (FPR) and false negative rate (FNR) of \hat{S}_i are defined as

$$\text{FPR}_i = \frac{|\hat{S}_i \cap S_i^c|}{|S_i^c|}, \quad \text{FNR}_i = \frac{|\hat{S}_i^c \cap S_i|}{|S_i|}. \quad (3.17)$$

By convention, if $S_i^c = \emptyset$, then $\text{FPR}_i = 0$; if $S_i = \emptyset$, then $\text{FNR}_i = 0$.

If $\text{FPR}_i = \text{FNR}_i = 0$ with probability tending to one, which is a very strong requirement, then we have the pattern recovery consistency $\mathbb{P}(\hat{S}_i = S_i) \rightarrow 1$. Below, we show that both FPR and FNR are asymptotically controlled in the presence of weak signals.

Theorem 3.5. Assume that Assumption 1,2,3 are satisfied. Fix an $i \in [nb_n + 1, n(1 - b_n) + 1]$. If $A_i \in \mathcal{G}_\alpha(s, M_d)$, then we have $\mathbb{P}(\text{FPR}_i = 0) \geq 1 - 2d^{-1}$. If in addition $A_{i,m} \in \mathcal{G}_{\alpha, \beta}(s, M_d, L_d)$, then we also have $\mathbb{P}(\text{FNR}_i \leq 2^\beta L_d u_i^\beta) \geq 1 - 2d^{-1}$.

Since $u_{\#} = o(1)$, the FNR vanishes with probability tending to one if $L_d = o(u_{\#}^{-\beta})$.

4. Simulation studies

In this section, we present some numerical results on simulated datasets. We compare the following five methods:

- (i) Our TV-VAR estimator (2.5).
- (ii) The stationary vector autoregressive model (Stat. VAR) [21].
- (iii) The time-varying lasso method [23].
- (iv) The time-varying ridge method [20].
- (v) The time-varying maximum likelihood estimator (MLE).

Methods (iii), (iv) and (v) are extensions of the lasso [23], ridge regression [20] and MLE to the time-varying setting by kernel smoothing. We include (ii), i.e. the stationary VAR model, as the baseline method that ignores the dynamic features in the transition matrices. (iii) is a competitor of (i). We solve (iii) with the FISTA [6] and we solve (i) and (ii) with the simplex algorithm [26].

4.1. Data generation

We consider different setups of $n = 100$ and $d = 20, 30, 40$ and 50 . For each setup (n, d) , the data are generated by the following procedure. First, the baseline coefficient matrices A_{01} and A_{02} are generated by using the `sugm.generator()` in `flare` R package [24]. We consider four graph structures defined in `flare`: hub, cluster, band and random for A_{01} and A_{02} . Examples of these four structures are shown in Appendix B. Then we normalize $\rho(A_{01}) = 0.2$, $\rho(A_{02}) = 1$ and smoothly interpolate on the intermediate values

$$A_i = (1 - t_i)^4 A_{01} + t_i^2 A_{02}, \quad t_i = i/n \text{ for } i = 1, \dots, n.$$

Following [21], we specify the innovation covariance matrix $\Psi = \Sigma - A_{01}\Sigma A_{01}^\top$ where $\Sigma = I_d$. [21] showed that the choice of Σ does not affect the numeric performance significantly. In our simulation studies, we use the Epanechnikov kernel $K(v) = 0.75(1 - v^2)\mathbb{I}(|v| \leq 1)$ and fix the model order $k = 1$. The bandwidth b_n is set to be $b_n = 0.8n^{-1/5}$ for (i), (iii) and (iv).

4.2. Tuning parameter selection

For the tuning parameter selection in (i)–(iv), we propose a data-driven procedure that minimizes the one-step-ahead prediction errors as follows.

1. Choose a grid for the tuning parameter (say τ) and the number n_1 of training data points.

2. For each τ , perform the one-step-ahead prediction on the testing set by estimating A_t with $\hat{A}_t(\tau)$ and then predict X_t by $\hat{X}_t = \hat{A}_t(\tau)X_{t-1}$, where $t = n_1 + 1, \dots, n$. Then calculate the prediction error at time t as $Err_t(\tau) = \|X_t - \hat{A}_t(\tau)X_{t-1}\|_2$.
3. Calculate the average errors over the last $n - n_1$ time points

$$\overline{Err}(\tau) = \frac{1}{n - n_1} \sum_{t=n_1+1}^n Err_t(\tau)$$

4. Select $\hat{\tau}$ that minimizes $\overline{Err}(\tau)$.

4.3. Comparison results

4.3.1. Estimation errors

For each setup, we run 100 simulations and for each simulation indexed by nn we estimate the transition matrix at each time point indexed by t , where $t = 1, \dots, n$. Then, we calculate the error $err_{t,nn} = |\hat{A}_{t,nn} - A_t|_\rho$ for $t = 1, \dots, n$ and $nn = 1, \dots, 100$. We also report errors under ℓ_F , ℓ_1 and ℓ_∞ norms. Next, the averaged errors over different time points are aggregated as $\overline{err}_{nn} = \frac{1}{b-a} \sum_{t=a}^b err_{t,nn}$ where $a = [nb_n] + 1$ and $b = [n(1 - b_n)] - 1$. Finally, the averaged errors are reported over 100 simulations as $\overline{err} = \frac{1}{100} \sum_{nn=1}^{100} err_{nn}$. The results are shown in Table 1, 2, 3 and 4.

From Table 1 to Table 4, larger d often results in larger errors. In general, the unstructured MLE performs the worst almost under all matrix norms. Although the ridge method shrinks the coefficients in the transition matrix to zeros, the values of those coefficients are not exactly zeros. Thus, the estimated transition matrices by the ridge method are not sparse which lead to higher estimation errors compared with the TV-VAR, time-varying lasso and stationary VAR. Stationary VAR always perform worse than TV-VAR and time-varying lasso, since it cannot capture the dynamic structure. The time-varying lasso performs better than any other methods except TV-VAR. The proposed TV-VAR performs the best almost under all matrix norms.

4.3.2. Pattern recovery

For the TV-VAR, we also report the $FPR_{t,l,nn}$ and $FNR_{t,l,nn}$ in (3.17), where $l = 1, \dots, 30$ which are the indexes of 30 possible tuning parameters from 0.001 to 0.45. Then, we calculate the averaged $FPR_l = \frac{1}{100} \sum_{nn=1}^{100} \{ \frac{1}{b-a} \sum_{t=a}^b FPR_{t,l,nn} \}$ and $FNR_l = \frac{1}{100} \sum_{nn=1}^{100} \{ \frac{1}{b-a} \sum_{t=a}^b FNR_{t,l,nn} \}$. Following [9], we set the $u_\#$ to 10^{-3} , which is considered to be numerical nonzero. The ROC curves for all possible values of the sparsity control parameter are plotted in Fig. 1a, Fig. 1b, Fig. 1c and Fig. 1d. Based on the ROC curves, the TV-VAR method has better discrimination power for band or random patterns than hub or cluster patterns.

TABLE 1

Comparison of estimation errors under different setups. The standard deviations are shown in parentheses. Here ℓ_ρ , ℓ_F , ℓ_1 and ℓ_∞ are the spectral, Frobenius, ℓ_1 and ℓ_∞ matrix norm resp. The pattern of transition matrix is ‘hub’.

| d=20 | | | | | |
|---------------|-------------------------|------------------|-------------------------|------------------|------------------|
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.407 (0.083) | 1.102 (0.248) | 0.453 (0.068) | 1.530 (0.075) | 2.770 (0.205) |
| ℓ_1 | 0.395 (0.088) | 1.747 (0.516) | 0.446 (0.076) | 1.532 (0.068) | 2.762 (0.255) |
| ℓ_ρ | 0.329 (0.038) | 0.778 (0.147) | 0.348 (0.032) | 0.604 (0.022) | 1.157 (0.08) |
| ℓ_F | 0.239 (0.011) | 0.292 (0.032) | 0.228 (0.011) | 0.342 (0.008) | 0.568 (0.02) |
| d=30 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.643 (0.125) | 1.154 (0.295) | 0.695 (0.111) | 2.116 (0.087) | 4.350 (0.302) |
| ℓ_1 | 0.694 (0.135) | 2.197 (0.652) | 0.732 (0.109) | 2.102 (0.073) | 4.302 (0.302) |
| ℓ_ρ | 0.430 (0.046) | 0.846 (0.141) | 0.443 (0.04) | 0.710 (0.024) | 1.607 (0.112) |
| ℓ_F | 0.245 (0.013) | 0.288 (0.026) | 0.244 (0.014) | 0.393 (0.006) | 0.729 (0.02) |
| d=40 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.711 (0.103) | 1.221 (0.256) | 0.800 (0.095) | 2.594 (0.079) | 6.183 (0.426) |
| ℓ_1 | 0.812 (0.144) | 2.868 (0.775) | 0.872 (0.126) | 2.586 (0.082) | 6.198 (0.406) |
| ℓ_ρ | 0.460 (0.041) | 0.958 (0.124) | 0.486 (0.035) | 0.793 (0.024) | 2.168 (0.142) |
| ℓ_F | 0.253 (0.011) | 0.296 (0.02) | 0.259 (0.011) | 0.429 (0.006) | 0.907 (0.021) |
| d=50 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.827 (0.124) | 1.158 (0.219) | 0.944 (0.102) | 3.043 (0.095) | 8.478 (0.633) |
| ℓ_1 | 1.019 (0.19) | 3.574 (1.021) | 1.072 (0.135) | 3.055 (0.098) | 8.558 (0.65) |
| ℓ_ρ | 0.504 (0.045) | 1.052 (0.157) | 0.531 (0.032) | 0.863 (0.021) | 2.874 (0.224) |
| ℓ_F | 0.266 (0.014) | 0.297 (0.019) | 0.278 (0.012) | 0.454 (0.006) | 1.108 (0.027) |

TABLE 2

Comparison of estimation errors under different setups. The standard deviations are shown in parentheses. Here ℓ_ρ , ℓ_F , ℓ_1 and ℓ_∞ are the spectral, Frobenius, ℓ_1 and ℓ_∞ matrix norm resp. The pattern of transition matrix is 'cluster'.

| d=20 | | | | | |
|---------------|-------------------------|------------------|-------------------------|------------------|------------------|
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.400 (0.073) | 1.093 (0.300) | 0.465 (0.084) | 1.527 (0.081) | 2.736 (0.202) |
| ℓ_1 | 0.383 (0.074) | 1.695 (0.506) | 0.448 (0.077) | 1.527 (0.083) | 2.731 (0.251) |
| ℓ_ρ | 0.324 (0.032) | 0.774 (0.143) | 0.352 (0.033) | 0.601 (0.025) | 1.133 (0.084) |
| ℓ_F | 0.239 (0.011) | 0.291 (0.03) | 0.228 (0.011) | 0.341 (0.008) | 0.562 (0.019) |
| d=30 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.557 (0.096) | 1.127 (0.285) | 0.631 (0.097) | 2.095 (0.083) | 4.305 (0.308) |
| ℓ_1 | 0.589 (0.132) | 2.236 (0.659) | 0.664 (0.110) | 2.120 (0.093) | 4.343 (0.297) |
| ℓ_ρ | 0.399 (0.045) | 0.848 (0.135) | 0.424 (0.038) | 0.710 (0.024) | 1.607 (0.112) |
| ℓ_F | 0.241 (0.009) | 0.289 (0.027) | 0.239 (0.010) | 0.393 (0.006) | 0.729 (0.020) |
| d=40 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.705 (0.115) | 1.139 (0.225) | 0.783 (0.108) | 2.612 (0.106) | 6.179 (0.44) |
| ℓ_1 | 0.803 (0.153) | 2.986 (0.874) | 0.847 (0.133) | 2.596 (0.098) | 6.171 (0.452) |
| ℓ_ρ | 0.460 (0.042) | 0.974 (0.15) | 0.479 (0.037) | 0.794 (0.022) | 2.158 (0.156) |
| ℓ_F | 0.253 (0.011) | 0.295 (0.022) | 0.258 (0.012) | 0.428 (0.006) | 0.906 (0.022) |
| d=50 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.841 (0.121) | 1.116 (0.198) | 0.964 (0.108) | 3.052 (0.106) | 8.593 (0.703) |
| ℓ_1 | 1.023 (0.177) | 3.384 (1.055) | 1.070 (0.152) | 3.043 (0.09) | 8.498 (0.558) |
| ℓ_ρ | 0.504 (0.04) | 1.015 (0.162) | 0.531 (0.036) | 0.863 (0.021) | 2.874 (0.224) |
| ℓ_F | 0.267 (0.013) | 0.292 (0.019) | 0.278 (0.013) | 0.454 (0.006) | 1.108 (0.027) |

TABLE 3

Comparison of estimation errors under different setups. The standard deviations are shown in parentheses. Here ℓ_ρ , ℓ_F , ℓ_1 and ℓ_∞ are the spectral, Frobenius, ℓ_1 and ℓ_∞ matrix norm resp. The pattern of transition matrix is 'band'.

| d=20 | | | | | |
|---------------|-------------------------|------------------|-------------------------|------------------|------------------|
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.402 (0.072) | 1.048 (0.251) | 0.453 (0.068) | 1.532 (0.08) | 2.738 (0.188) |
| ℓ_1 | 0.385 (0.075) | 1.645 (0.476) | 0.437 (0.077) | 1.531 (0.068) | 2.707 (0.183) |
| ℓ_ρ | 0.326 (0.033) | 0.758 (0.137) | 0.347 (0.032) | 0.600 (0.021) | 1.132 (0.072) |
| ℓ_F | 0.237 (0.011) | 0.289 (0.031) | 0.226 (0.01) | 0.341 (0.008) | 0.563 (0.019) |
| d=30 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.574 (0.094) | 1.139 (0.319) | 0.637 (0.101) | 1.537 (0.058) | 4.316 (0.314) |
| ℓ_1 | 0.590 (0.131) | 2.306 (0.727) | 0.654 (0.123) | 1.537 (0.057) | 4.312 (0.323) |
| ℓ_ρ | 0.401 (0.041) | 0.863 (0.154) | 0.422 (0.038) | 0.537 (0.019) | 1.607 (0.112) |
| ℓ_F | 0.240 (0.009) | 0.290 (0.028) | 0.239 (0.011) | 0.315 (0.005) | 0.729 (0.02) |
| d=40 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.800 (0.134) | 1.170 (0.253) | 0.886 (0.113) | 2.590 (0.098) | 6.198 (0.429) |
| ℓ_1 | 0.925 (0.155) | 2.772 (0.825) | 0.964 (0.127) | 2.597 (0.09) | 6.221 (0.457) |
| ℓ_ρ | 0.489 (0.041) | 0.938 (0.138) | 0.510 (0.035) | 0.793 (0.024) | 2.168 (0.142) |
| ℓ_F | 0.262 (0.013) | 0.293 (0.022) | 0.270 (0.012) | 0.429 (0.006) | 0.907 (0.021) |
| d=50 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.839 (0.125) | 1.143 (0.222) | 0.992 (0.118) | 3.056 (0.100) | 8.550 (0.579) |
| ℓ_1 | 1.011 (0.175) | 3.673 (1.212) | 1.105 (0.126) | 3.047 (0.099) | 8.635 (0.678) |
| ℓ_ρ | 0.506 (0.045) | 1.047 (0.176) | 0.541 (0.034) | 0.862 (0.021) | 2.872 (0.223) |
| ℓ_F | 0.267 (0.014) | 0.295 (0.02) | 0.280 (0.013) | 0.454 (0.005) | 1.108 (0.027) |

TABLE 4

Comparison of estimation errors under different setups. The standard deviations are shown in parentheses. Here ℓ_ρ , ℓ_F , ℓ_1 and ℓ_∞ are the spectral, Frobenius, ℓ_1 and ℓ_∞ matrix norm resp. The pattern of transition matrix is 'random'.

| d=20 | | | | | |
|---------------|-------------------------|------------------|-------------------------|------------------|------------------|
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.397 (0.078) | 1.060 (0.244) | 0.447 (0.07) | 1.538 (0.071) | 2.704 (0.188) |
| ℓ_1 | 0.379 (0.080) | 1.661 (0.574) | 0.433 (0.081) | 1.532 (0.077) | 2.762 (0.235) |
| ℓ_ρ | 0.323 (0.036) | 0.757 (0.140) | 0.344 (0.033) | 0.600 (0.021) | 1.132 (0.072) |
| ℓ_F | 0.237 (0.013) | 0.286 (0.031) | 0.227 (0.012) | 0.341 (0.008) | 0.563 (0.019) |
| d=30 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.567 (0.103) | 1.138 (0.264) | 0.641 (0.102) | 1.541 (0.066) | 4.294 (0.292) |
| ℓ_1 | 0.590 (0.132) | 2.280 (0.724) | 0.658 (0.119) | 1.543 (0.062) | 4.292 (0.268) |
| ℓ_ρ | 0.398 (0.044) | 0.865 (0.146) | 0.423 (0.041) | 0.537 (0.019) | 1.607 (0.112) |
| ℓ_F | 0.240 (0.01) | 0.290 (0.026) | 0.239 (0.012) | 0.315 (0.005) | 0.729 (0.020) |
| d=40 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.713 (0.098) | 1.210 (0.268) | 0.810 (0.104) | 2.595 (0.096) | 6.148 (0.413) |
| ℓ_1 | 0.813 (0.136) | 2.904 (0.851) | 0.869 (0.128) | 2.600 (0.09) | 6.205 (0.432) |
| ℓ_ρ | 0.458 (0.036) | 0.951 (0.142) | 0.482 (0.034) | 0.793 (0.024) | 2.168 (0.142) |
| ℓ_F | 0.252 (0.01) | 0.296 (0.021) | 0.260 (0.012) | 0.429 (0.006) | 0.907 (0.021) |
| d=50 | | | | | |
| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
| ℓ_∞ | 0.830 (0.109) | 1.207 (0.234) | 1.002 (0.120) | 3.060 (0.108) | 8.452 (0.553) |
| ℓ_1 | 1.023 (0.187) | 3.529 (1.214) | 1.129 (0.151) | 3.061 (0.092) | 8.457 (0.551) |
| ℓ_ρ | 0.504 (0.039) | 1.034 (0.180) | 0.543 (0.033) | 0.865 (0.022) | 2.861 (0.198) |
| ℓ_F | 0.267 (0.013) | 0.298 (0.021) | 0.282 (0.013) | 0.455 (0.005) | 1.107 (0.027) |

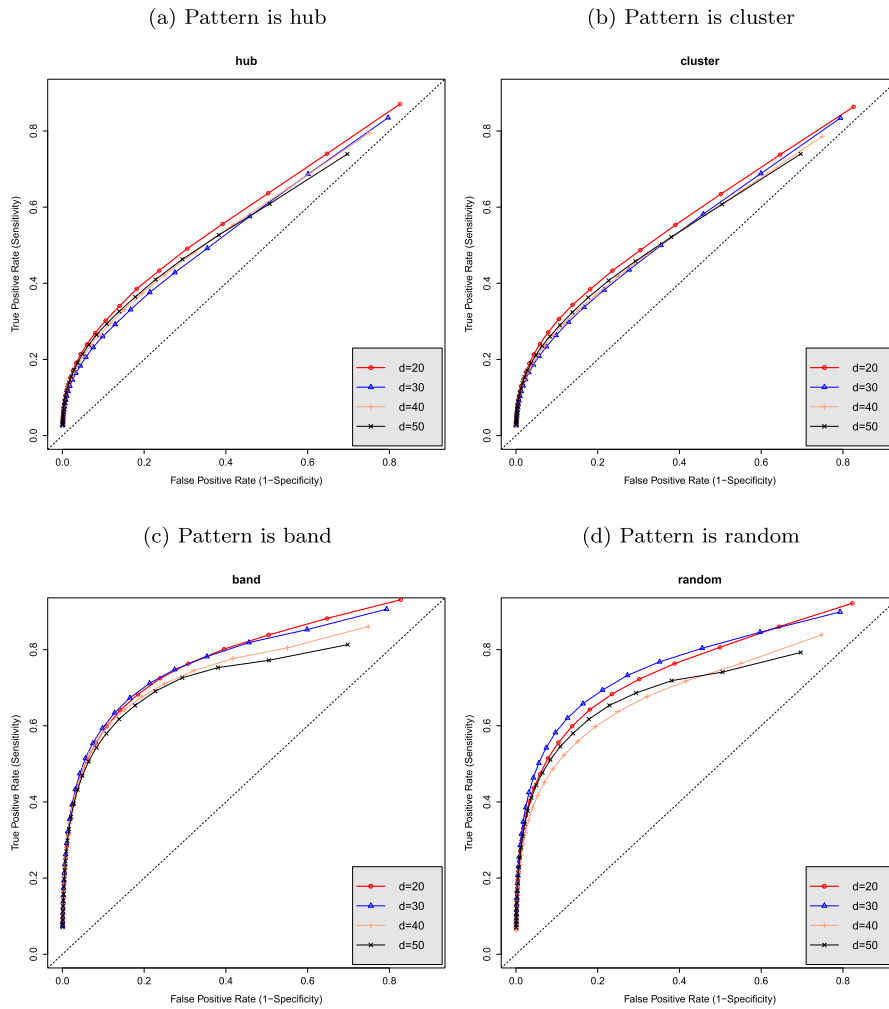


FIG 1. ROC curves under different settings

We also calculate $u_{\hat{\mu}}$ by using the true value of Σ in (3.9) and the resulting ROC curves are shown in Appendix C. Those ROC curves are similar to those based on $u_{\hat{\mu}} = 10^{-3}$.

5. Real data analysis

5.1. Finance data: stock prices

In this section, we compare the aforementioned estimators on a real financial dataset. The dataset is from Yahoo! Finance (finance.yahoo.com). The data

matrix contains daily closing prices of 452 stocks that are consistently in the S&P 500 index between January 1st, 2003 and January 1st, 2008. We choose such time range to avoid the effects of the two financial crises in the year of 2001 and 2008, which could make the stock prices non-smooth with sharp drops and pick-ups. In total, there are 1,258 time points. We first standardize the data to zero mean and unit variance and detrend the data. We then fit the detrended data for those stocks which are most smoothed (without obvious change points) using AR(1) model and perform the Ljung-Box Tests by using `Box.test()` in R. We keep those stocks with nonzero coefficients at significant level 0.05. 30 stocks are finally selected and 10 of them are shown in Table 9. The ten selected stocks in Table 9 are from six sectors including: Consumer Staples, Consumer Discretionary, Industrials, Financials, Utilities and Energy. The resulting data matrix is denoted as $\bar{X} \in \mathbb{R}^{1258 \times 30}$. We use the Priestley-Subba Rao (PSR) stationarity test (such as `stationarity()` in R package `fractal`) to some of the stocks such as Kellogg Co. and Target Corp. to find that these time series are not stationary at significant level 0.05. Thus, it is inappropriate to model the data with the stationary VAR model in [21]. Finally we fit the data \bar{X} into the sparse TV-VAR model with order $k = 1$, stationary VAR [21] with order $k = 1$, time-varying lasso method [23], time-varying ridge method [20] and time-varying MLE.

In order to compare the performance of the five methods, we consider the one-step-ahead prediction for \bar{X}_t , $t = 1159, \dots, 1258$, by using $\bar{X}_{J_t, *}$, where $J_t = \{j : t - 1158 \leq j \leq t - 1\}$ as the training set. We estimate the transition matrices $\hat{A}_t(\lambda)$ where λ is the regularization parameter such as the τ in the TV-VAR or the shrinkage parameter in the ridge method. We use the Epanechnikov kernel $K(v) = 0.75(1 - v^2)\mathbb{I}(|v| \leq 1)$ with bandwidth $b_n = 0.3$, which means that only the last 30% data in the training set are used for prediction for the TV-VAR, time-varying lasso method, time-varying ridge method and time-varying MLE. Since stationary VAR [21] model uses all training data, in order to make fair comparisons, we only treat the last 30% time points (347 observations) in $\bar{X}_{J_t, *}$ as the training set for the stationary VAR. The averaged prediction error for a specific λ is measured by:

$$\overline{Err}(\lambda) = \frac{1}{100} \sum_{t=1159}^{1258} \|\bar{X}_{t, *} - \hat{A}_t(\lambda)\bar{X}_{t-1, *}\|_2.$$

The smallest averaged one-step-ahead prediction errors $\min_{\lambda} \overline{Err}(\lambda)$ for five methods were shown in Table 5. In Figure 2, we show an example for the predicted closing prices from the five methods and the detrended true closing prices of Target Corp. Clearly, the methods with time-varying structures perform similarly but outperform stationary VAR.

One advantage of the TV-VAR compared with the stationary VAR is that it can capture the time-varying data structures. If we treat a transition matrix as an adjacency matrix for an undirected weighted graph, we can visualize the graph structures at different time points. For instance, Fig 3a, 3b, 3c and 3d are four graphical representations of the TV-VAR, which clearly show some time-

evolving patterns. In the graphs, different vertices represent different stocks and the edges represent the cross-sectional dependence between stocks. Bolder lines indicate stronger dependence and the stocks in the same sector are in the same color. In order to illustrate the time-varying structure clearly, we only consider the 10 stocks in Table 9. From these four figures, we observe that stocks in the same and closely related sectors are often connected. For examples, CME Group Inc. and Hartford Financial Svc. GP from Financials sector show a dynamic dependence structure, and Boeing Company in Industrials shows the consistent correlation with Exxon Mobil Corp. in Energy sector.

TABLE 5

The prediction errors for the five methods on the stock price data. The standard deviations are shown in parentheses.

| | TV-VAR | Stat. VAR | Lasso | Ridge | MLE |
|-----------------------------|----------|-----------|----------|----------|----------|
| $\overline{Err}(\lambda^*)$ | 0.4822 | 2.2824 | 0.4580 | 0.4843 | 0.4902 |
| Standard Deviation | (0.1751) | (0.5169) | (0.1810) | (0.1747) | (0.1799) |




FIG 2. Comparison of the predicted closing prices and the true detrended closing prices of Target Corp.

5.2. Economic data: exchange rates

We also apply the exchange rates data on international currencies and study how the exchange rates evolve over time. We use the exchange rates data from the Federal Reserve Bank at St. Louis and choose the date range from January 1st, 2003 to January 1st, 2008 to avoid the financial crises' effects on

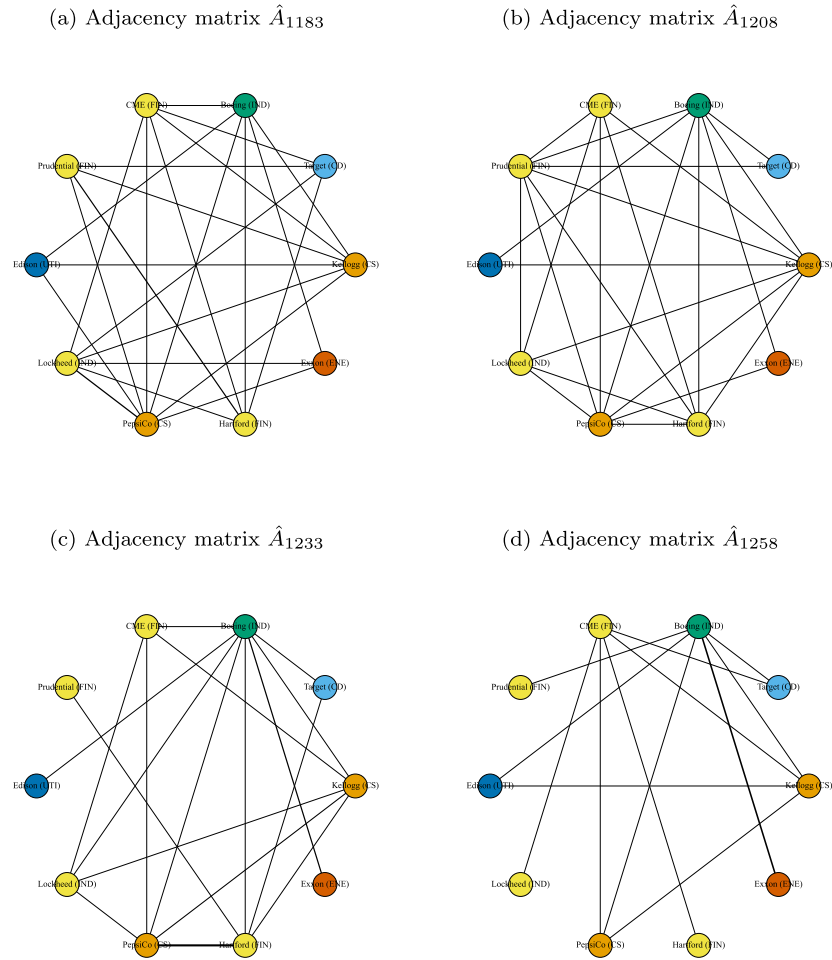


FIG 3. The graphs based on adjacency matrices estimated by TV-VAR

exchange rates. In total, there are 1,303 time points. We choose 15 major currencies from the continents including Europe, Australia, North America, Asian, South America and we normalize each currency by the exchange rate of the U.S. dollar shown in Table 10. As in Section 5.1, we standardize and detrend the data. The exchange rates for all currencies to the U.S. dollars are smoothed and in addition the PSR test rejects the stationarity hypothesis of the exchange rate Euro/U.S. at the significant level 0.05. Thus, stationary VAR is also not appropriate in this application, which is also supported by our empirical findings in Table 6 and Figure 4. We compare the performance of the five methods by the one-step-ahead prediction errors for the last 50 data and use the same prediction metric in Section 5.1. The bandwidth here is $b_n = 0.3$. We do not include the prediction curve of the stationary VAR in Figure 4

TABLE 6
 The prediction errors for the five methods on the exchange rates data. The standard deviations are shown in parentheses.

| | TV-VAR | Stat. Var | Lasso | Ridge | MLE |
|-----------------------------|----------|-----------|----------|----------|----------|
| $\overline{Err}(\lambda^*)$ | 0.2990 | 8.7637 | 0.2554 | 0.2634 | 0.2638 |
| Standard Deviation | (0.0924) | (5.2835) | (0.1037) | (0.1032) | (0.1032) |




FIG 4. Compare the predicted and true detrended exchange rate of U.K./U.S.

since its errors are too large. Table 6 and Figure 4 show that methods with time-varying structure performs similarly but much better than the stationary VAR.

6. Proofs

Proof of Lemma 3.1. Since $\tilde{x}_m(t)$ is a stationary VAR, we have $\tilde{\Sigma}(t) = I_{d \times d} + A(t)\tilde{\Sigma}(t)A(t)^\top$. By the extension, $\Sigma(t) = \tilde{\Sigma}(t)$ for all $t \in (0, 1)$. Applying the chain rule of differentiation w.r.t. t , we get

$$\dot{\Sigma}(t) = \dot{A}(t)\Sigma(t)A(t)^\top + A(t)\dot{\Sigma}(t)A(t)^\top + A(t)\Sigma(t)\dot{A}(t)^\top.$$

By Lemma 6.1 and the triangle inequality,

$$\begin{aligned} & |\dot{\Sigma}(t)|_\infty \\ & \leq |\dot{A}(t)|_\infty |\Sigma(t)|_{\ell^\infty} |A(t)|_{\ell^1} + |A(t)|_{\ell^1} |\dot{\Sigma}(t)|_\infty |A(t)|_{\ell^1} + |A(t)|_{\ell^1} |\Sigma(t)|_{\ell^1} |\dot{A}(t)|_\infty. \end{aligned}$$

Then, (3.4) follows from the assumption that $\sup_{t \in [0,1]} |A(t)|_{\ell^1} < 1$ and that $\Sigma(t)$ is symmetric. (3.5) follows from the differentiation of $\dot{\Sigma}(t)$ w.r.t. t . \square

Lemma 6.1. *Let A, B, C be matrices of compatible dimension for the product ABC . Then*

$$\begin{aligned} |ABC|_\infty &\leq |A|_{\ell^1} |B|_\infty |C|_{\ell^\infty}, \\ |ABC|_\infty &\leq |A|_{\ell^1} |B|_{\ell^1} |C|_\infty, \\ |ABC|_\infty &\leq |A|_\infty |B|_{\ell^\infty} |C|_{\ell^\infty}. \end{aligned}$$

Proof of Lemma 6.1. The first and second inequalities follow from

$$\begin{aligned} |ABC|_\infty &= \max_{i,j} \left| \sum_{k,m} A_{ik} B_{km} C_{mj} \right| \leq \max_{i,j} \sum_k |A_{ik}| \sum_m |B_{km} C_{mj}| \\ &\leq \left(\max_i \sum_k |A_{ik}| \right) \max_{j,k} \sum_m |B_{km} C_{mj}| \\ &\leq |A|_{\ell^1} \min \{ |B|_\infty |C|_{\ell^\infty}, |B|_{\ell^1} |C|_\infty \}. \end{aligned}$$

The third inequality follows from the second one by considering $(ABC)^\top = C^\top B^\top A^\top$ and $|A^\top|_{\ell^1} = |A|_{\ell^\infty}$. \square

Recall the TV-VAR(1) model $\mathbf{x}_i = A_i \mathbf{x}_{i-1} + \mathbf{e}_i$. The following key lemma presents a large deviation bound for the marginal and lag-one autocovariance matrices with sub-Gaussian innovations.

Lemma 6.2. *Suppose $\bar{\rho} = \sup_{i \geq 0} \rho(A_i) < 1$ and*

$$\sup_{t \in [0,1]} |\dot{\Sigma}_{m,jk}(t)| \vee |\ddot{\Sigma}_{m,jk}(t)| \leq C_0, \quad m = 0, 1,$$

for some absolute constant $C_0 < \infty$. If $\mathbf{e}_i = (e_{i1}, \dots, e_{id})$ has iid sub-Gaussian components, i.e. $\|e_{ij}\|_q \leq C_1 q^{1/2}$ for all $q \geq 1$, then, for any fixed $t \in [b_n, 1 - b_n]$, we have with probability at least $1 - 2d^{-1}$ that

$$|\hat{\Sigma}_{t,0} - \Sigma_{t,0}|_\infty \vee |\hat{\Sigma}_{t,1} - \Sigma_{t,1}|_\infty \leq C \left(b_n^2 + \sqrt{\frac{\log d}{nb_n}} \right), \quad (6.1)$$

where the constant $C = C(\bar{\rho}, C_0, C_1)$ does not depend on n, b_n and d .

Proof of Lemma 6.2. First, consider the marginal covariance matrix $m = 0$. For each $j, k = 1, \dots, p$, we write

$$\hat{\Sigma}_{0,jk}(t) - \Sigma_{0,jk}(t) = \{\hat{\Sigma}_{0,jk}(t) - \mathbb{E}[\hat{\Sigma}_{0,jk}(t)]\} + \{\mathbb{E}[\hat{\Sigma}_{0,jk}(t)] - \Sigma_{0,jk}(t)\} = I + II.$$

For $|t_i - t| \leq b_n$, by the assumption $\sup_{t \in [0,1]} |\ddot{\Sigma}_{0,jk}(t)| \leq C_0$, Taylor's expansion yields that $\Sigma_{0,jk}(t_i) = \Sigma_{0,jk}(t) + \dot{\Sigma}_{0,jk}(t)(t_i - t) + O(b_n^2)$ for all $t \in [b_n, 1 - b_n]$. Note that for $\ell = 0, 1, 2$, we have

$$\sup_{t \in [b_n, 1 - b_n]} \left| \sum_{i=1}^n K \left(\frac{t_i - t}{b_n} \right) (t_i - t)^\ell - nb_n \int_{-1}^1 K(v) v^\ell dv \right| = O(1).$$

Then, the bias part is controlled by

$$\begin{aligned} II &= \sum_{i=1}^n w(t, i) \Sigma_{0,jk}(t_i) - \Sigma_{0,jk}(t) \\ &= \dot{\Sigma}_{0,jk}(t) \sum_{i=1}^n w(t, i)(t_i - t) + O(b_n^2) = O((nb_n)^{-1} + b_n^2). \end{aligned} \tag{6.2}$$

Now, we deal with the stochastic part I. Let $Y_{i,jk} = X_{ij}X_{ik} - \mathbb{E}(X_{ij}X_{ik})$. Clearly, $\mathbb{E}(Y_{i,jk}) = 0$ and $I = \sum_{i=1}^n w(t, i)Y_{i,jk}$. Let $B_{i,0} = I_d$ and $B_{i,m} = A_i \cdots A_{i-m+1}$ for $m \geq 1$. Then, $\sup_i \rho(B_{i,m}) \leq \bar{\rho}^m$ and we have the moving-average (MA) representation of $\mathbf{x}_i = \sum_{m=0}^{\infty} B_{i,m} \mathbf{e}_{i-m}$. Let $\boldsymbol{\xi} = (\mathbf{e}_n^\top, \mathbf{e}_{n-1}^\top, \cdots)^\top$, $W_t = \text{diag}(w(t, n), \cdots, w(t, 1))$ be $n \times n$ diagonal matrix, and

$$\tilde{B}^{(j)} = \begin{pmatrix} B_{n,0,j} & B_{n,1,j} & B_{n,2,j} & \cdots & B_{n,n-1,j} & B_{n,n,j} & \cdots \\ 0 & B_{n-1,0,j} & B_{n-1,1,j} & \cdots & B_{n-1,n-2,j} & B_{n-1,n-1,j} & \cdots \\ 0 & 0 & B_{n-2,0,j} & \cdots & B_{n-2,n-3,j} & B_{n-2,n-2,j} & \cdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & B_{1,0,j} & B_{1,1,j} & \cdots \end{pmatrix}.$$

Then, we can write $I = \boldsymbol{\xi}^\top (\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)} \boldsymbol{\xi} - \text{tr}((\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)})$. By the Hanson-Wright inequality [29, Theorem 1.1], there exists a constant $C := C(C_1)$ such that for all $x > 0$

$$\begin{aligned} (*) &:= \mathbb{P} \left(\left| \boldsymbol{\xi}^\top (\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)} \boldsymbol{\xi} - \text{tr}((\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)}) \right| \geq x \right) \\ &\leq 2 \exp \left\{ -C \min \left[\frac{x^2}{|(\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)}|_F^2}, \frac{x}{\rho((\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)})} \right] \right\}. \end{aligned}$$

Observe that $(\tilde{B}^{(j)})^\top W_t \tilde{B}^{(k)}$ has the same nonzero eigenvalues as the $n \times n$ matrix $W_t^{1/2} \tilde{B}^{(k)} (\tilde{B}^{(j)})^\top W_t^{1/2}$. Therefore, by the Cauchy-Schwartz inequality, we get

$$\begin{aligned} |W_t^{1/2} \tilde{B}^{(k)} (\tilde{B}^{(j)})^\top W_t^{1/2}|_F^2 &= \text{tr} \left[W_t^{1/2} \tilde{B}^{(j)} (\tilde{B}^{(k)})^\top W_t \tilde{B}^{(k)} (\tilde{B}^{(j)})^\top W_t^{1/2} \right] \\ &= \text{tr} \left[(\tilde{B}^{(k)})^\top W_t \tilde{B}^{(k)} (\tilde{B}^{(j)})^\top W_t \tilde{B}^{(j)} \right] \\ &\leq |(\tilde{B}^{(k)})^\top W_t \tilde{B}^{(k)}|_F \cdot |(\tilde{B}^{(j)})^\top W_t \tilde{B}^{(j)}|_F \end{aligned}$$

and by the definition of matrix spectral norm

$$\begin{aligned} \rho(W_t^{1/2} \tilde{B}^{(k)} (\tilde{B}^{(j)})^\top W_t^{1/2}) &\leq \rho(W_t^{1/2} \tilde{B}^{(k)}) \rho((\tilde{B}^{(j)})^\top W_t^{1/2}) \\ &= \rho^{1/2}((\tilde{B}^{(k)})^\top W_t \tilde{B}^{(k)}) \rho^{1/2}((\tilde{B}^{(j)})^\top W_t \tilde{B}^{(j)}). \end{aligned}$$

Therefore, (*) is bounded by

$$2 \exp \left\{ -C \min \left[\frac{x^2}{\max_{j \leq p} |W_t^{1/2} \tilde{B}^{(j)} (\tilde{B}^{(j)})^\top W_t^{1/2}|_F^2}, \frac{x}{\max_{j \leq p} \rho(W_t^{1/2} \tilde{B}^{(j)} (\tilde{B}^{(j)})^\top W_t^{1/2})} \right] \right\}.$$

For $i = 1, \dots, n$ and $l = 0, \dots, n - i$, let

$$\gamma_{i,l}^{(j)} = \sum_{m=0}^{\infty} [w(t, i)w(t, i+l)]^{1/2} |B_{i,m,j} \cdot B_{i+l,m+l,j}^\top|.$$

By the Cauchy-Schwartz inequality and the spectral norm bound $\sup_i \rho(B_{i,m}) \leq \bar{\rho}^m$, we have

$$\begin{aligned} \gamma_{i,l}^{(j)} &\leq [w(t, i)w(t, i+l)]^{1/2} \sum_{m=0}^{\infty} \left(\sum_{k=1}^p B_{i,m,jk}^2 \right)^{1/2} \left(\sum_{k=1}^p B_{i+l,m+l,jk}^2 \right)^{1/2} \\ &\leq [w(t, i)w(t, i+l)]^{1/2} \sum_{m=0}^{\infty} \bar{\rho}^m \bar{\rho}^{m+l} = [w(t, i)w(t, i+l)]^{1/2} \frac{\bar{\rho}^l}{1 - \bar{\rho}^2}. \end{aligned}$$

Then, it follows that

$$\begin{aligned} &|W_t^{1/2} \tilde{B}^{(j)} (\tilde{B}^{(j)})^\top W_t^{1/2}|_F^2 \\ &\leq \frac{\sum_{i=1}^n w(t, i)^2}{(1 - \bar{\rho}^2)^2} + 2 \sum_{l=1}^{n-1} \sum_{i=1}^{n-l} [w(t, i)w(t, i+l)]^{1/2} \frac{\bar{\rho}^l}{(1 - \bar{\rho}^2)^2} \\ &\leq \frac{2 \left[\sum_{l=0}^{n-1} \bar{\rho}^{2l} \right] \left[\sum_{i=1}^n w(t, i)^2 \right]}{(1 - \bar{\rho}^2)^2} \\ &\leq C(\bar{\rho}) \sum_{i=1}^n w(t, i)^2 = \frac{C(\bar{\rho})}{nb_n} \end{aligned}$$

and

$$\begin{aligned} &\rho(W_t^{1/2} \tilde{B}^{(j)} (\tilde{B}^{(j)})^\top W_t^{1/2}) \\ &\leq 2 \left[\frac{1}{1 - \bar{\rho}^2} \max_{1 \leq i \leq n} w(t, i) + \sum_{l=1}^{n-1} \frac{\bar{\rho}^l}{1 - \bar{\rho}^2} \max_{1 \leq i \leq n-l} [w(t, i)w(t, i+l)]^{1/2} \right] \\ &\leq \frac{2 \left[\sum_{l=0}^{n-1} \bar{\rho}^l \right] \left[\max_{1 \leq i \leq n} w(t, i) \right]}{1 - \bar{\rho}^2} \leq \frac{C(\bar{\rho})}{nb_n}. \end{aligned}$$

Therefore, we have that for any $x > 0$

$$\sup_{t \in [b_{n-1}, b_n]} \mathbb{P} \left(\left| \sum_{i=1}^n w(t, i) Y_{i,jk} \right| \geq x \right) \leq 2 \exp[-Cnb_n \min(x^2, x)]. \quad (6.3)$$

Now, by (6.2) and using the union bound applied to (6.3), we obtain that there exists a constant C which only depends on $\bar{\rho}, C_0, C_1$ such that

$$|\hat{\Sigma}_{t,0} - \Sigma_{t,0}|_\infty \leq C \left(b_n^2 + \sqrt{\frac{\log d}{nb_n}} \right)$$

holds with probability $\geq 1 - 2d^{-1}$. Similar argument applied to $m = 1$ shows that the lag-one autocovariance matrix obeys the same bound in (6.1). \square

Proof of Lemma 3.4. Denote $l = d/2$. For $0 < \alpha < 1$, note that

$$\max_{1 \leq m \leq d} \sum_{k=1}^d A_{mk}^\alpha = 2 \sum_{k=1}^{l-1} A_{lk}^\alpha + A_{ll}^\alpha = 2 \sum_{k=1}^{l-1} A_{lk}^\alpha + 1.$$

Then, we have

$$\begin{aligned} \sum_{k=1}^{l-1} A_{lk}^\alpha &= \sum_{k=1}^{l-1} \frac{1}{(1 + |l - k|^2/d^{2r})^{\gamma\alpha}} \\ &= \sum_{k=1}^{l-1} \frac{d^{2r\gamma\alpha}}{(d^{2r} + k^2)^{\gamma\alpha}} \\ &\leq 1 + \int_1^l \frac{d^{2r\gamma\alpha}}{(d^{2r} + x^2)^{\gamma\alpha}} dx \\ &= 1 + d^r \int_{d^{-r}}^{d^{1-r}/2} \frac{1}{(1 + x^2)^{\gamma\alpha}} dx \\ &\leq Cd^r \int_1^{d^{1-r}/2} x^{-2\gamma\alpha} dx, \end{aligned}$$

where $C > 0$ is a constant depending only on r, γ, α . Then, it follows that

$$\max_{1 \leq m \leq d} \sum_{k=1}^d A_{mk}^\alpha = \begin{cases} O(d^r) & \text{if } 2\gamma\alpha > 1 \\ O(d^r \log d) & \text{if } 2\gamma\alpha = 1 \\ O(d^{r+(1-r)(1-2\gamma\alpha)}) & \text{if } 1 < 2\gamma\alpha < 1 \end{cases}.$$

So, $A \in \mathcal{G}_\alpha^-(s, M_d)$ for our choice of s and M_d . Since A is symmetric, $A \in \mathcal{G}_\alpha(s, M_d)$. Next, we show that $A \in \mathcal{G}_{\alpha,\beta}(s, M_d, L_d)$. Observe that $|\text{supp}(A)| = (3d - 2)d/4$. By construction, $A_{mk} \leq u$ is equivalent to $\pi(h_m^\circ, h_k^\circ) \geq f^{-1}(u) = (u^{-\frac{1}{\gamma}} - 1)^{\frac{1}{2}}$. Therefore, $A_{mk} \leq u$ if and only if $|m - k| \geq d^r(u^{-\frac{1}{\gamma}} - 1)^{\frac{1}{2}}$. Since $d^r(u^{-\frac{1}{\gamma}} - 1)^{\frac{1}{2}}$ is monotone decreasing with respect to u ,

$$\exists u \in (0, u_0], \text{ s.t. } \frac{d}{2} - 2 < d^r(u^{-\frac{1}{\gamma}} - 1)^{\frac{1}{2}} \leq \frac{d}{2} - 1.$$

So we only need to consider $u \in [u, u_0]$ because if $u \in (0, u)$, then $|\{(m, k) : 0 < A_{mk} \leq u\}| = |\{(m, k) : |m - k| \geq d/2\}| = 0$. Therefore, if we choose L_d such that $L_d u^\beta \gtrsim 1$, then $A \in \mathcal{G}_{\alpha,\beta}(s, M_d, L_d)$. That is, $L_d = C u^{-\beta} = C d^{2(1-r)\gamma\beta}$. \square

Proof of Theorem 3.2. Put $\varepsilon = C(b_n^2 + \sqrt{(\log d)/(nb_n)})$. Recall that $A_i^\top = \Sigma_{i-1,0}^{-1}\Sigma_{i-1,1}$. By Lemma 6.2, we have, with probability at least $1 - 2d^{-1}$, that $|\hat{\Sigma}_{t,0} - \Sigma_{t,0}|_\infty \leq \varepsilon$ and $|\hat{\Sigma}_{t,1} - \Sigma_{t,1}|_\infty \leq \varepsilon$ for any $t \in [b_n, 1 - b_n]$. Then, with such high probability, the true transition matrix A_i is feasible for the minimization program (2.5) with $\tau \geq (1 + M_d)\varepsilon$ because

$$\begin{aligned} |\hat{\Sigma}_{i-1,1} - \hat{\Sigma}_{i-1,0}A_i^\top|_\infty &= |\hat{\Sigma}_{i-1,1} - \hat{\Sigma}_{i-1,0}\Sigma_{i-1,0}^{-1}\Sigma_{i-1,1}|_\infty \\ &\leq |\hat{\Sigma}_{i-1,1} - \Sigma_{i-1,1}|_\infty + |(I - \hat{\Sigma}_{i-1,0}\Sigma_{i-1,0}^{-1})\Sigma_{i-1,1}|_\infty \\ &\leq \varepsilon + |\hat{\Sigma}_{i-1,0} - \Sigma_{i-1,0}|_\infty|A_i^\top|_{\ell^1} \leq \tau. \end{aligned}$$

Hence, $|\hat{A}_i|_1 \leq |A_i|_1$. But then

$$\begin{aligned} &|\hat{A}_i^\top - A_i^\top|_\infty \\ &= |\hat{A}_i^\top - \Sigma_{i-1,0}^{-1}\Sigma_{i-1,1}|_\infty \\ &= |\Sigma_{i-1,0}^{-1}(\Sigma_{i-1,0}\hat{A}_i^\top - \Sigma_{i-1,1})|_\infty \\ &= |\Sigma_{i-1,0}^{-1}(\Sigma_{i-1,0}\hat{A}_i^\top - \hat{\Sigma}_{i-1,0}\hat{A}_i^\top + \hat{\Sigma}_{i-1,0}\hat{A}_i^\top - \hat{\Sigma}_{i-1,1} + \hat{\Sigma}_{i-1,1} - \Sigma_{i-1,1})|_\infty \\ &\leq |\Sigma_{i-1,0}^{-1}(\hat{\Sigma}_{i-1,0} - \hat{\Sigma}_{i-1,1})\hat{A}_i^\top|_\infty + |\Sigma_{i-1,0}^{-1}(\hat{\Sigma}_{i-1,0}\hat{A}_i^\top - \hat{\Sigma}_{i-1,1})|_\infty \\ &\quad + |\Sigma_{i-1,0}^{-1}(\hat{\Sigma}_{i-1,1} - \Sigma_{i-1,1})|_\infty \\ &\leq |\Sigma_{i-1,0}^{-1}|_{\ell^1}(|\hat{\Sigma}_{i-1,0} - \hat{\Sigma}_{i-1,1}|_\infty|\hat{A}_i|_{\ell^1} + \tau + \varepsilon) \\ &\leq |\Sigma_{i-1,0}^{-1}|_{\ell^1}(\varepsilon|A_i|_{\ell^1} + \tau + \varepsilon) \\ &\leq 2\tau|\Sigma_{i-1,0}^{-1}|_{\ell^1}. \end{aligned}$$

Let $u \geq 0$ and $T_j = \{m : |A_{i,jm}| \geq u\}$, $j = 1, \dots, d$. Define $D(u) = \max_{j \leq d} \sum_{m=1}^d (|A_{i,jm}| \wedge u)$. Then, by the triangle inequality and the equivalence between (2.5) and (2.6), we have uniformly in $j = 1, \dots, d$,

$$\begin{aligned} |[\hat{A}_i - A_i]_{j*}|_1 &\leq |[\hat{A}_i]_{j,T_j^c}|_1 + |[A_i]_{j,T_j^c}|_1 + |[\hat{A}_i - A_i]_{j,T_j}|_1 \\ &= |[A_i]_{j*}|_1 - |[\hat{A}_i]_{j,T_j}|_1 + |[A_i]_{j,T_j^c}|_1 + |[\hat{A}_i - A_i]_{j,T_j}|_1 \\ &\leq 2|[A_i]_{j,T_j^c}|_1 + 2|[\hat{A}_i - A_i]_{j,T_j}|_1 \\ &\leq 2|[A_i]_{j,T_j^c}|_1 + 4\tau|\Sigma_{i-1,0}^{-1}|_{\ell^1}|T_j| \\ &\leq 2D(u)(1 + 2\tau|\Sigma_{i-1,0}^{-1}|_{\ell^1}/u). \end{aligned}$$

Choose $u = 2\tau|\Sigma_{i,0}^{-1}|_{\ell^1}$. Since $D(u) \leq su^{1-\alpha}$ for $A_i \in \mathcal{G}_\alpha(s, M_d)$, it follows that

$$|[\hat{A}_i - A_i]_{j*}|_1 \leq 4D(u) \leq 4su^{1-\alpha} = C(\alpha)s(|\Sigma_{i,0}^{-1}|_{\ell^1}\tau)^{1-\alpha}.$$

The analysis of the other constraint is similar. Now, (3.7) and (3.8) are immediate in view of $\rho^2(M) \leq |M|_{\ell^1}|M|_{\ell^\infty}$ and $|M|_F^2 \leq d|M|_\infty|M|_{\ell^1}$ for any $d \times d$ matrix M . \square

Proof of Theorem 3.3. Suppose $(m, k) \notin S_i$, i.e. $A_{i,mk} = 0$. On the event $G = \{|\hat{A}_i - A_i|_\infty \leq u_\sharp\}$, we have $|\hat{A}_{i,mk}| \leq u_\sharp$ and $(m, k) \notin \hat{S}_i$. Therefore, by Theorem 3.2, we have $\mathbb{P}(\hat{S}_i \subset S_i) \geq 1 - 2d^{-1}$. On the other hand, suppose that

$|A_{i,mk}| > 2u_{\#}$, then on the event G , it follows from the triangle inequality that $|A_{i,mk}| - |\hat{A}_{i,mk}| \leq |\hat{A}_{i,mk} - A_{i,mk}| \leq u_{\#}$. Therefore, $|\hat{A}_{i,mk}| > u_{\#}$ and we have $\mathbb{P}(\{(m, k) : |A_{i,mk}| > 2u_{\#}\} \subset \hat{S}_i) \geq 1 - 2d^{-1}$. \square

Proof of Theorem 3.5. If $|S_i| = 0$, then the first claim is trivial. So we may assume that $S_i \neq \emptyset$. Suppose $(m, k) \notin S_i$, then on the event $G = \{|\hat{A}_i - A_i|_{\infty} \leq u_{\#}\}$, we have $|\hat{A}_{i,mk}| \leq u_{\#}$ and $(m, k) \notin \hat{S}_i$. Therefore, $\hat{S}_i \cap S_i^c = \emptyset$ on G and $\mathbb{P}(\text{FPR}_i = 0) \geq 1 - 2d^{-1}$ by Theorem 3.2. On the other hand, let $N_i(u) = \{(m, k) : 0 < |A_{i,mk}| \leq 2u\}$. For any $u > 0$, $N_i(u) \subset S_i$ and $\hat{S}_i^c \cap S_i \subset N_i(u) \cup (\hat{S}_i^c \cap S_i \cap N_i(u)^c)$. Hence, $|\hat{S}_i^c \cap S_i| \leq |N_i(u)| + |\hat{S}_i^c \cap S_i \cap N_i(u)^c|$. If $(m, k) \in \hat{S}_i^c$ and $|A_{i,mk}| > 2u_{\#}$, then $\hat{S}_i^c \cap S_i \cap N_i(u_{\#})^c = \emptyset$ on G . Choose $u = u_{\#}$ and then on the event G , we have

$$\text{FNR}_i = \frac{|\hat{S}_i^c \cap S_i|}{|S_i|} \leq \frac{|N_i(u_{\#})|}{|S_i|} + \frac{|\hat{S}_i^c \cap S_i \cap N_i(u_{\#})^c|}{|S_i|} \leq 2^{\beta} L_d u_{\#}^{\beta},$$

from which the theorem follows. \square

Appendix A: Detailed setup of the simulation studies

We use `sugm.generator()` to generate sparse matrices from multivariate normal distributions with four graph structures- hub, cluster, band and random.

TABLE 7
The meanings of `sugm.generator()`'s arguments

| Arguments | Meanings |
|-------------------|---|
| <code>g</code> | Represents the number of hubs or clusters when the pattern is "hub" or "cluster" |
| <code>prob</code> | Represents the probability that an off-diagonal entry will be non-zero. |
| <code>v</code> | Assigns the values to the off-diagonal entries in the precision matrix and controls the magnitude of partial correlations with <code>u</code> . |
| <code>u</code> | A positive number which is added to the diagonal entries of the precision matrix. |

TABLE 8
Detailed setup of `sugm.generator()`

| Hub | | | | | Cluster | | | | |
|------|----|------|-------|----|---------|------|-------|-------|----|
| d | g | prob | v | u | d | g | prob | v | u |
| 20 | 8 | NULL | 0.001 | 10 | 20 | 8 | NULL | 0.001 | 10 |
| 30 | 10 | NULL | 0.001 | 10 | 30 | 10 | NULL | 0.001 | 10 |
| 40 | 15 | NULL | 0.001 | 10 | 40 | 15 | NULL | 0.001 | 10 |
| 50 | 20 | NULL | 0.001 | 10 | 50 | 20 | NULL | 0.001 | 10 |
| Band | | | | | Random | | | | |
| d | g | prob | v | u | d | g | prob | v | u |
| 20 | 1 | NULL | 0.001 | 10 | 20 | NULL | 0.001 | 0.001 | 10 |
| 30 | 1 | NULL | 0.001 | 10 | 30 | NULL | 0.001 | 0.001 | 10 |
| 40 | 1 | NULL | 0.001 | 10 | 40 | NULL | 0.001 | 0.001 | 10 |
| 50 | 1 | NULL | 0.001 | 10 | 50 | NULL | 0.001 | 0.001 | 10 |

Besides taking the patterns of graphs, the `sugm.generator()` also takes other four critical input arguments which are `g`, `prob`, `v` and `u` shown in Table 7. The detailed setup of these four parameters in our simulation is shown in Table 8.

Appendix B: Examples for four graph structures

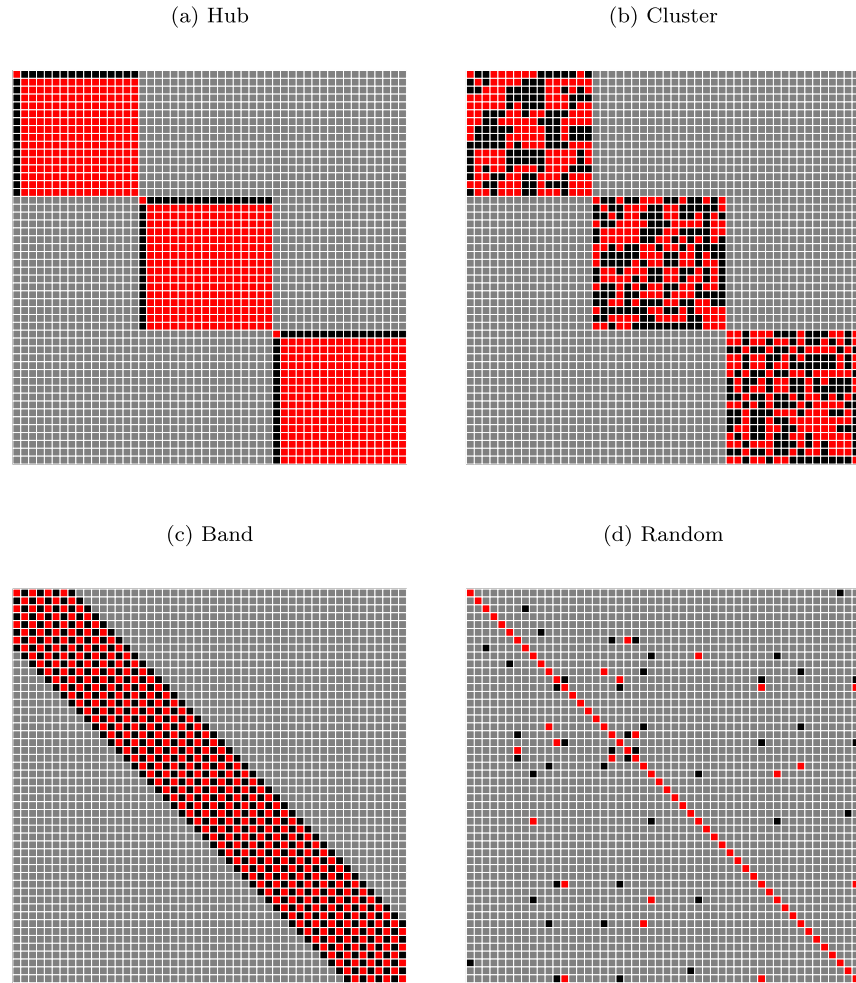


FIG 5. Examples of the baseline coefficient matrices in four graph structures. Zero entries are colored in grey. Positive entries and negative entries are colored in red and black respectively.

Appendix C: ROC curves based on the true thresholding parameters $u_{\#}$

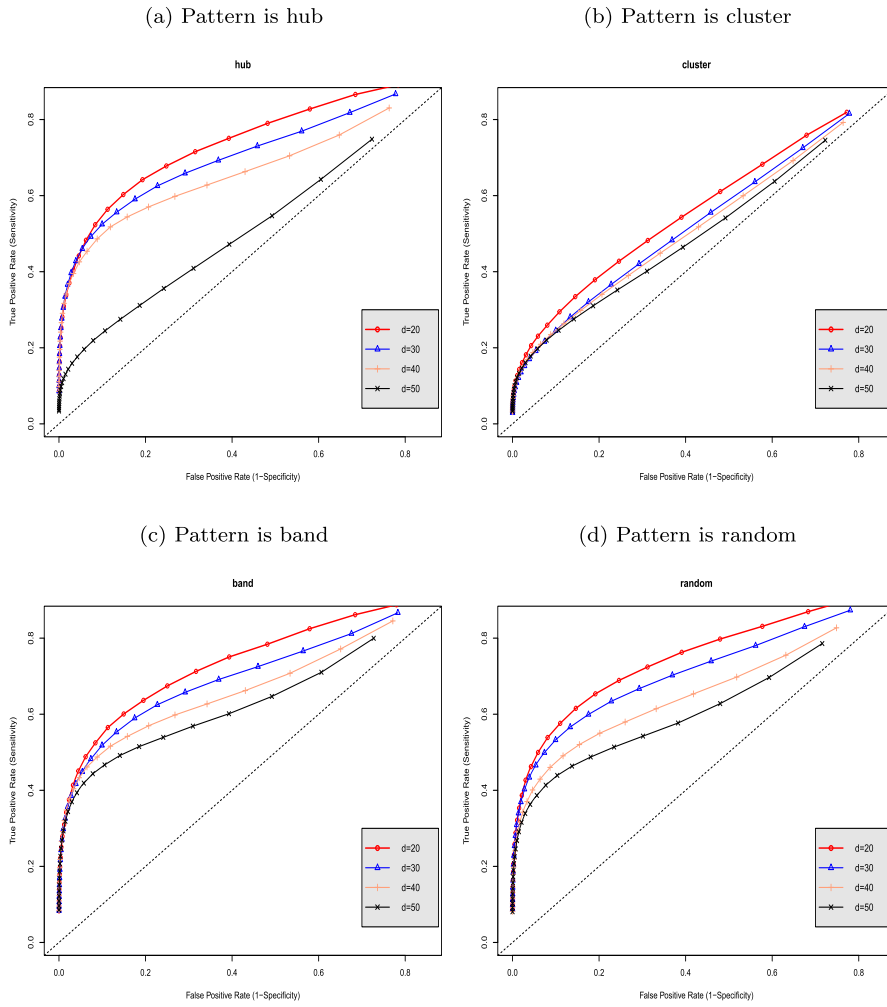


FIG 6. ROC curves under different settings

Appendix D: Selected stocks and currencies in real data analysis

TABLE 9
Part of the 30 selected stocks

| | Stock (Abb.) | Sector (Abb.) |
|----|---------------------------------------|-----------------------------|
| 1 | Kellogg Co. (Kellogg) | Consumer Staples (CS) |
| 2 | Target Corp. (Target) | Consumer Discretionary (CD) |
| 3 | Boeing Company (Boeing) | Industrial (IND) |
| 4 | CME Group Inc. (CME) | Financials (FIN) |
| 5 | Prudential Financial (Prudential) | Financials (FIN) |
| 6 | Edison Int'l (Edison) | Utilities (UTI) |
| 7 | Lockheed Martin Corp. (Lockheed) | Industrial (IND) |
| 8 | PepsiCo Inc. (PepsiCo) | Consumer Staples (CS) |
| 9 | Hartford Financial Svc. GP (Hartford) | Financials (FIN) |
| 10 | Exxon Mobil Corp. (Exxon) | Energy (ENE) |

TABLE 10
Selected currencies

| | Currencies | | Currencies |
|---|--------------------|----|---------------------|
| 1 | Euro / U.S. | 9 | Mexico / U.S. |
| 2 | U.K. / U.S. | 10 | South Korea / U.S. |
| 3 | Swiss / U.S. | 11 | India / U.S. |
| 4 | Australia / U.S. | 12 | Thailand / U.S. |
| 5 | New Zealand / U.S. | 13 | South Africa / U.S. |
| 6 | Canada / U.S. | 14 | Norway / U.S. |
| 7 | Singapore / U.S. | 15 | Sweden / U.S. |
| 8 | Brazil / U.S. | | |

Acknowledgments

We thank Fang Han for providing the stock data in Section 5.1. The high-performance computing work was done on the Illinois Campus Cluster Program (ICCP) at the University of Illinois at Urbana-Champaign.

References

- [1] Lian An and Jian Wang. Exchange rate pass-through: Evidence based on vector autoregression with sign restrictions. *Open Economies Review*, 23(2):359–380, 2012.
- [2] Andrew Ang and Monika Piazzesi. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, 50:745–787, 2003.
- [3] David Backus. The canadian–u.s. exchange rate: Evidence from a vector autoregression. *The Review of Economics and Statistics*, 68(4):628–637, 1986.

- [4] Richard T. Baillie, Robert E. Lippens, and Patrick C. McMahon. Testing rational expectations and efficiency in the foreign exchange market. *Econometrica*, 51(3):553–563, 1983.
- [5] Sumanta Basu and George Michailidis. Regularized estimation in sparse high-dimensional time series models. *Annals of Statistics*, 43(4):1535–1567, 2015. [MR3357870](#)
- [6] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009. [MR2486527](#)
- [7] Peter J. Bickel and Elizaveta Levina. Covariance Regularization by Thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008. [MR2485008](#)
- [8] Peter J. Brockwell and Richard A. Davis. *Time Series: Theory and Methods*. Series in Statistics. Springer, second edition, 1991. [MR1093459](#)
- [9] Tony Cai, Weidong Liu, and Xi Luo. A Constrained ℓ_1 Minimization Approach to Sparse Precision Matrix Estimation. *Journal of American Statistical Association*, 106(494):594–607, 2011. [MR2847973](#)
- [10] John Y. Campbell and Robert J. Shiller. Stock prices, earnings, and expected dividends. *Journal of Finance*, 43(3):661–676, 1988.
- [11] Xiaohui Chen, Mengyu Xu, and Wei Biao Wu. Covariance and precision matrix estimation for high-dimensional time series. *Annals of Statistics*, 41(6):2994–3021, 2013. [MR3161455](#)
- [12] Xiaohui Chen, Mengyu Xu, and Wei Biao Wu. Regularized estimation of linear functionals of precision matrices for high-dimensional time series. *IEEE Transactions on Signal Processing*, 64(24):6459–6470, 2016. [MR3566612](#)
- [13] Alex Coad. Exploring the processes of firm growth: evidence from a vector auto-regression. *Industrial and Corporate Change*, 19(6):1677–1703, 2010.
- [14] John H. Cochrane. Permanent and transitory components of gnp and stock prices. *The Quarterly Journal of Economics*, 109(1):241–265, 1994.
- [15] Rainer Dahlhaus. Fitting Time Series Models to Nonstationary Processes. *The Annals of Statistics*, 25(1):1–37, 1997. [MR1429916](#)
- [16] Rainer Dahlhaus. Locally stationary processes. *arXiv:1109.4174*, 2012.
- [17] Dana Draghicescu, Serge Guillas, and Wei Biao Wu. Quantile Curve Estimation and Visualization for Nonstationary Time Series. *Journal of Computational and Graphical Statistics*, 18(1):1–20, 2009. [MR2511058](#)
- [18] Jianqing Fan, Fang Han, and Han Liu. Challenges of big data analysis. *National Science Review*, (1):293–314,, 2014.
- [19] Shaojun Guo, Yazhen Wang, and Qiwei Yao. High dimensional and banded vector autoregressions. *arXiv:1502.07831*, 2015. [MR3620446](#)
- [20] J.D. Hamilton. *Time Series Analysis*. Cambridge University Press, 2 edition, 1994.
- [21] Fang Han, Huanran Lu, and Han Liu. A direct estimation of high dimensional stationary vector autoregressions. *Journal of Machine Learning Research*, 16:3115–3150, 2015. [MR3450535](#)
- [22] Fang Han, Sheng Xu, and Han Liu. Rate-optimal estimation of high dimensional time series. *Preprint*, 2016.

- [23] N. J. Hsu, H.L. Hung, and Y.M. Chang. Subset selection for vector autoregressive process using lasso. *Computational Statistics and Data Analysis*, 52(7):3645–3657, 2008. [MR2427370](#)
- [24] Xingguo Li, Tuo Zhao, Xiaoming Yuan, and Han Liu. An R Package flare for High Dimensional Linear Regression and Precision Matrix Estimation. *Journal of Machine Learning Research*, to appear, 2014. [MR3335796](#)
- [25] Helmut Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, 2007. [MR2172368](#)
- [26] Katta G Murty. *Linear programming*. John Wiley & Sons, New York, 1983. [MR0720547](#)
- [27] Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005. [MR2148143](#)
- [28] Pradeep Ravikumar, Martin J. Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011. [MR2836766](#)
- [29] Mark Rudelson and Roman Vershynin. Hanson-Wright inequality and sub-Gaussian concentration. *Electronic Communication in Probability*, 18(7):1–9, 2013. [MR3125258](#)
- [30] Ali Shojaie and George Michailidis. Discovering graphical granger causality using the truncating lasso penalty. *Bioinformatics*, 26(18):517–523, 2010.
- [31] Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.
- [32] Christopher A. Sims. Interpreting the macroeconomic time series facts: The effects of monetary policy. *European Economic Review*, 36(5):975–1000, 1992.
- [33] Song Song and Peter J. Bickel. Large vector auto regressions. *arXiv:1106.3915*, 2011.
- [34] J.H. Stock and M.W. Watson. Implications of dynamic factor models for var analysis. *Working Paper No. 11467, National Bureau of Economic Research*, 2005.
- [35] H. Wang, G Li, and C.L. Tsai. Regression coefficient and autoregressive order shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, 69(1):63–78, 2007. [MR2301500](#)
- [36] Christopher K Wikle and Mevin B Hooten. A general science-based framework for dynamical spatio-temporal models. *Test*, 19(3):417–451, 2010. [MR2745992](#)
- [37] Shuheng Zhou, John Lafferty, and Larry Wasserman. Time Varying Undirected Graphs. *Machine Learning*, 80:295–319, 2010. [MR3108169](#)