

UC Berkeley

UC Berkeley Previously Published Works

Title

Statistical Inference for High-Dimensional Vector Autoregression with Measurement Error.

Permalink

<https://escholarship.org/uc/item/3s80c2g7>

Journal

Statistica Sinica, 33(Spec Issue)

ISSN

1017-0405

Authors

Lyu, Xiang

Kang, Jian

Li, Lexin

Publication Date

2023-05-01

DOI

10.5705/ss.202021.0151

Peer reviewed



Published in final edited form as:

Stat Sin. 2023 May ; 33(Spec Iss): 1435–1459. doi:10.5705/ss.202021.0151.

Statistical Inference for High-Dimensional Vector Autoregression with Measurement Error

Xiang Lyu[†], Jian Kang[‡], Lexin Li[†]

[†] University of California at Berkeley

[‡] University of Michigan

Abstract

High-dimensional vector autoregression with measurement error is frequently encountered in a large variety of scientific and business applications. In this article, we study statistical inference of the transition matrix under this model. While there has been a large body of literature studying sparse estimation of the transition matrix, there is a paucity of inference solutions, especially in the high-dimensional scenario. We develop inferential procedures for both the global and simultaneous testing of the transition matrix. We first develop a new sparse expectation-maximization algorithm to estimate the model parameters, and carefully characterize their estimation precisions. We then construct a Gaussian matrix, after proper bias and variance corrections, from which we derive the test statistics. Finally, we develop the testing procedures and establish their asymptotic guarantees. We study the finite-sample performance of our tests through intensive simulations, and illustrate with a brain connectivity analysis example.

Keywords

Brain connectivity analysis; Covariance inference; Expectation-maximization algorithm; Simultaneous testing; Global testing; Vector autoregression

1 Introduction

In this article, we study statistical inference for high-dimensional vector autoregression (VAR) with measurement error. More specifically, we consider the model,

$$\begin{aligned} \mathbf{y}_t &= \mathbf{x}_t + \boldsymbol{\epsilon}_t, \\ \mathbf{x}_{t+1} &= \mathbf{A}_* \mathbf{x}_t + \boldsymbol{\eta}_t, \end{aligned} \tag{1}$$

where $\mathbf{y}_t = (y_{t,1}, \dots, y_{t,p})^\top \in \mathbb{R}^p$ is the observed multivariate time series, $\mathbf{x}_t = (x_{t,1}, \dots, x_{t,p})^\top \in \mathbb{R}^p$ is the multivariate latent signal that admits an autoregressive structure, $\boldsymbol{\epsilon}_t = (\epsilon_{t,1}, \dots, \epsilon_{t,p})^\top \in \mathbb{R}^p$ is the measurement error for the observed time series, $\boldsymbol{\eta}_t = (\eta_{t,1}, \dots, \eta_{t,p})^\top \in \mathbb{R}^p$ is the white noise of the latent signal, and $\mathbf{A}_* = (A_{*,ij}) \in \mathbb{R}^{p \times p}$ is the sparse transition matrix that encodes the directional relations among the latent signal variables of \mathbf{x}_t . Furthermore, we focus on the scenario $\|\mathbf{A}_*\|_2 < 1$ such that the VAR model of \mathbf{x}_t is stationary. The error terms $\boldsymbol{\epsilon}_t$ and $\boldsymbol{\eta}_t$

are i.i.d. multivariate normal with mean zero and covariance $\sigma_{\epsilon_*}^2 \mathbf{I}_p$ and $\sigma_{\eta_*}^2 \mathbf{I}_p$, respectively, and are independent of \mathbf{x}_t . Here we focus on the lag-1 autoregressive structure and homoscedastic errors. We later discuss potential extensions in Section 7.

Model like (1) is frequently employed in a variety of scientific and business applications, e.g., finance, engineering and neuroscience. Our motivation is brain effective connectivity analysis based on functional magnetic resonance imaging (fMRI). The brain is a highly interconnected dynamic system, in which the activity and temporal evolution of neural elements are triggered and influenced by the activities of other elements (Garg et al., 2011). Of great interest in neuroscience is to understand the directional relations among the neural elements through fMRI, which measures synchronized blood oxygen level dependent signals at different brain locations. VAR model is an important tool to model such directional relations, which are encoded by the transition matrix \mathbf{A}_* , while the stationarity is often assumed (Bullmore and Sporns, 2009; Chen et al., 2011). However, unlike a typical VAR, the observed time series \mathbf{y}_t is the contaminated version of the true signal \mathbf{x}_t , added with a measurement error ϵ_t (Zhang et al., 2015; Cao et al., 2019).

We address the statistical inference problem of the transition matrix \mathbf{A}_* under model (1), and we aim at a high-dimensional setting where p^2 exceeds the length of series T . We first test the global hypotheses,

$$H_0 : A_{*,ij} = A_{0,ij}, \text{ for all } (i, j) \in \mathcal{S} \text{ versus } H_1 : A_{*,ij} \neq A_{0,ij}, \text{ for some } (i, j) \in \mathcal{S}, \quad (2)$$

for a given $\mathbf{A}_0 = (A_{0,ij}) \in \mathbb{R}^{p \times p}$ and $\mathcal{S} \subseteq [p] \times [p]$, where $[p] = \{1, \dots, p\}$. The most common choice is $\mathbf{A}_0 = \mathbf{0}_{p \times p}$ and $\mathcal{S} = [p] \times [p]$. We next test the simultaneous hypotheses,

$$H_{0,ij} : A_{*,ij} = A_{0,ij}, \text{ versus } H_{1,ij} : A_{*,ij} \neq A_{0,ij}, \text{ for all } (i, j) \in \mathcal{S}. \quad (3)$$

There has been a large body of literature studying sparse estimation of \mathbf{A}_* in VAR models (Hsu et al., 2008; Song and Bickel, 2011; Negahban and Wainwright, 2011; Basu and Michailidis, 2015; Han et al., 2015, among others). However, they all assumed that there is no measurement error ϵ_t , or equivalently, that \mathbf{x}_t is fully observed. Moreover, while both estimation and inference can produce a sparse representation of \mathbf{A}_* , they are utterly different problems. Sparse estimation usually does not explicitly control the false discovery rate (type I error), and does not produce an explicit significance quantification (p -value). There has been a relative paucity of inference methods for \mathbf{A}_* in VAR models. Existing inference solutions mostly focused on the low-dimensional VAR setting; see Reinsel (2003); Tsay and Chen (2018); Shao (2015) for a review. More recently, for the high-dimensional VAR setting, Krampe et al. (2018) proposed to bootstrap the de-biased Lasso estimator, while Zheng and Raskutti (2019) extended the de-correlated score test of Ning and Liu (2017). However, they only addressed the global testing problem (2), but not the simultaneous testing problem (3). Besides, it is unclear how to adapt their tests to accommodate additional

measurement error. To the best of our knowledge, there is no existing solution to directly address both global and simultaneous testing problems under the high-dimensional VAR setting with error.

Our proposal is built upon two key ingredients: a sparse expectation-maximization (EM) algorithm, and the high-dimensional covariance inference. The first ingredient, the EM algorithm, offers a way to estimate model parameters in the existence of measurement error. Early EM methods, however, only justified the convergence to a local optimum and did not consider sparsity. Recently, a seminal work of Balakrishnan et al. (2017) provided sufficient conditions to guarantee the convergence of standard EM to the global optimum but only in a low-dimensional setting, while Cai et al. (2019) extended the guarantee to a high-dimensional sparse Gaussian mixture model. See also Wang et al. (2015); Yi and Caramanis (2015). On the other hand, they all worked with i.i.d. observations, whereas our problem involves temporally highly dependent data. Extension from independent to dependent observations is far from trivial. The second ingredient, the high-dimensional covariance inference, has been intensively studied in recent years, including both global testing (Cai and Jiang, 2011; Xiao and Wu, 2013; Chen et al., 2010) and simultaneous testing (Liu, 2013; Cai et al., 2013). See also Cai (2017); Cai and Sun (2017) for reviews. However, they all assumed the data which the covariance is constructed from are fully observed. By contrast, our inference is about the transition matrix \mathbf{A}_* of the latent unobserved \mathbf{x}_t , and the covariance of the observed \mathbf{y}_t is a nonlinear transformation of \mathbf{A}_* , making it difficult to trace back to \mathbf{A}_* . Consequently, there is a considerable gap before we can apply the existing covariance inference tools to our setting.

In this article, we develop inferential procedures for both the global and simultaneous testing problems (2) and (3) under the high-dimensional VAR model with error. Our proposal includes three main steps. First, we develop a new sparse EM algorithm to estimate relevant model parameters. Next, we construct a Gaussian matrix on the domain of transition matrix, from which we derive the test statistics. Finally, we develop the global and simultaneous testing procedures with proper theoretical guarantees.

In the first step, we develop a new sparse EM algorithm to estimate both the transition matrix \mathbf{A}_* and the error variances $\sigma_{\epsilon_t}^2$ and $\sigma_{\eta_t}^2$. In particular, the maximization step is done via a generalized Dantzig selector for Yule-Walker equation, which can be efficiently solved by parallel linear programming (Candes and Tao, 2007; Han et al., 2015). We then establish the convergence of our sparse EM estimators to the true parameters, within the statistical precision required for the test statistics and the transition matrix inferences in later steps. We note that, the existing EM theory adopts the log-likelihood in an infinite-sample scheme as the key analytical tool, which becomes an expectation at a single observation given i.i.d. observations (Balakrishnan et al., 2017; Cai et al., 2019). However, the temporal dependence in our model makes the expectation of the log-likelihood change with the sample size. To tackle the issue, in our theoretical analysis, we consider the expectation in a finite-sample scheme instead, which introduces additional technical difficulty. We then derive several new concentration inequalities to establish the statistical error under some weak sparsity assumptions.

In the second step, we construct a Gaussian matrix as the test statistic for our transition matrix inference. This is built on a key observation that the inference on \mathbf{A}_* is equivalent to the inference on the lagged auto-covariance of some noise term. Since this noise is not directly observed, we employ the sparse EM algorithm in the first step to reconstruct the noise. We then study the non-asymptotic behavior of the sample lagged auto-covariance of the reconstructed noise, and explicitly characterize its bias and variance. This in turn leads to the construction of the test statistic matrix whose entries marginally follow a standard Gaussian distribution under the null hypothesis.

In the third step, we develop a global testing procedure based on the extreme distribution of the maximal entry of the test statistic matrix from the second step, and develop a simultaneous testing procedure by thresholding at a level that controls false discovery rate (FDR). Theoretically, we obtain the asymptotic size and power of the global test, which together establish the consistency of our test. We also show that our simultaneous test achieves a consistent FDR control. Our testing procedures are extensions of the covariance inference methods such as Cai and Jiang (2011); Liu (2013); Cai et al. (2013). But unlike the existing methods that are built on the sample covariance of fully observed data, our tests are obtained from the sample lagged auto-covariance of the reconstructed noise. This difference requires us to derive new concentration inequalities and Gaussian approximations to disentangle the reconstruction error, lag effect, and temporal dependence. These new theoretical results themselves may be of independent interest.

We employ the following notation throughout this article. Let $|\mathcal{S}|$ denote the cardinality of a set \mathcal{S} . For a scalar $a \in \mathbb{R}$, let $\lceil a \rceil$ and $\lfloor a \rfloor$ denote the smallest and largest integer greater than or smaller than a . For two scalars $a, b \in \mathbb{R}$, let $a \vee b$ and $a \wedge b$ denote the maxima and minima. For a vector $\mathbf{a} = (a_1, \dots, a_p)^\top \in \mathbb{R}^p$, define $\|\mathbf{a}\|_1 = \sum_{i=1}^p |a_i|$, $\|\mathbf{a}\|_2 = (\sum_{i=1}^p |a_i|^2)^{1/2}$, and $\|\mathbf{a}\|_\infty = \max_{1 \leq i \leq p} |a_i|$. For an index set $\mathcal{S} \subseteq [p]$, let $\mathbf{a}_{\mathcal{S}}$ denote the sub-vector of \mathbf{a} containing only the coordinates indexed by \mathcal{S} . For a matrix $\mathbf{M} = (M_{ij}) \in \mathbb{R}^{p_1 \times p_2}$, define $\|\mathbf{M}\|_1 = \sum_{ij} |M_{ij}|$, $\|\mathbf{M}\|_2 = \lambda_{\max}^{1/2}(\mathbf{M}^\top \mathbf{M})$, $\|\mathbf{M}\|_F = (\sum_{ij} M_{ij}^2)^{1/2}$, $\|\mathbf{M}\|_{\max} = \max_{ij} |M_{ij}|$, $\|\mathbf{M}\|_{l_1} = \max_{j \in [p_2]} \sum_{i=1}^{p_1} |M_{ij}|$, $\|\mathbf{M}\|_{l_\infty} = \max_{i \in [p_1]} \sum_{j=1}^{p_2} |M_{ij}|$, $\|\mathbf{M}\|_{l_\infty} = \max_{i \in [p_1]} \sum_{j=1}^{p_2} |M_{ij}|$, and $\|\mathbf{M}\|_{r,2} = \max_{i \in [p_1]} \sqrt{\sum_{j=1}^{p_2} |M_{ij}|^2}$ to be its element-wise ℓ_1 norm, spectral norm, Frobenius norm, max norm, maximum absolute column sum, maximum absolute row sum, and maximal row-wise Euclidean norm, respectively. Let $\mathbf{M}_{i \cdot}$ and $\mathbf{M}_{\cdot j}$ denote the i th row and j th column. Let $\lambda_{\min}(\mathbf{M})$ and $\lambda_{\max}(\mathbf{M})$ denote its smallest and the largest eigenvalue, $\text{tr}(\mathbf{M})$ the trace, and $|\mathbf{M}|$ the determinant. Define $D(\mathbf{M})$ as a diagonal matrix whose diagonal elements are the same as those of \mathbf{M} .

The rest of the article is organized as follows. Section 2 presents the sparse EM algorithm, Section 3 constructs the test statistic matrix, and Section 4 develops the global and simultaneous testing procedures. Section 5 presents the simulations, and Section 6 illustrates with a brain connectivity analysis example. Section 7 concludes the paper with a discussion. All proofs are relegated to the Supplementary Appendix.

2 Sparse EM Estimation

2.1 Sparse EM algorithm

Let $\{\mathbf{y}_t, \mathbf{x}_t\}_{t=1}^T$ denote the complete data, where T is the total number of observations, \mathbf{y}_t is observed but \mathbf{x}_t is latent. Let $\Theta = \{\mathbf{A}, \sigma_{\eta}^2, \sigma_{\epsilon}^2\}$ collect all the parameters of interest in model (1), and $\Theta_* = \{\mathbf{A}_*, \sigma_{\eta,*}^2, \sigma_{\epsilon,*}^2\}$ denote the true parameters. The goal is to estimate Θ_* by maximizing the log-likelihood function of the observed data, $\ell(\Theta | \{\mathbf{y}_t\}_{t=1}^T)$, with respect to Θ . The computation of $\ell(\Theta | \{\mathbf{y}_t\}_{t=1}^T)$, however, is highly nontrivial. The standard EM algorithm then turns to an auxiliary function, named the finite-sample Q -function,

$$Q_y(\Theta | \Theta') = \mathbb{E}[\ell(\Theta | \{\mathbf{y}_t, \mathbf{x}_t\}_{t=1}^T) | \{\mathbf{y}_t\}_{t=1}^T, \Theta'],$$

which is defined as the expectation of the log-likelihood function for the complete data $\ell(\Theta | \{\mathbf{y}_t, \mathbf{x}_t\}_{t=1}^T)$, conditioning on a parameter set Θ' and the observed data \mathbf{y}_t , and the expectation is taken with respect to the latent data \mathbf{x}_t . The Q -function can be computed efficiently, and provides a lower bound of the target log-likelihood function $\ell(\Theta | \{\mathbf{y}_t\}_{t=1}^T)$ for any Θ . The equality $\ell(\Theta' | \{\mathbf{y}_t\}_{t=1}^T) = Q_y(\Theta' | \Theta')$ holds if $\Theta = \Theta'$. Maximizing Q -function provides an uphill step of the likelihood. Starting from an initial set of parameters $\hat{\Theta}_0$, the EM algorithm then alternates between the expectation step (E-step), where the Q -function $Q_y(\Theta | \hat{\Theta}_k)$ conditioning on the parameters $\hat{\Theta}_k$ of the k th iteration is computed, and the maximization step (M-step), where the parameters are updated by maximizing the Q -function $\hat{\Theta}_{k+1} = \arg\max_{\Theta} Q_y(\Theta | \hat{\Theta}_k)$.

For our problem, we carry out the E-step via the standard Kalman filter and smoother (Ghahramani and Hinton, 1996). For the M-step, the maximizer of $Q_y(\Theta | \hat{\Theta}_k)$ satisfies that $(T-1)^{-1} \sum_{t=1}^{T-1} \mathbf{E}_{t,t+1,k} = (T-1)^{-1} \sum_{t=1}^{T-1} \mathbf{E}_{t,t,k} \mathbf{A}^\top$, where $\mathbf{E}_{t,s,k} = \mathbb{E}\{\mathbf{x}_t \mathbf{x}_s^\top | \{\mathbf{y}_t\}_{t=1}^T, \hat{\Theta}_{k-1}\}$ for $s, t \in [T]$ is obtained from the E-step. The standard EM algorithm directly inverts the matrix involving $\mathbf{E}_{t,t,k}$'s, which is computationally challenging when the dimension p is high. In addition, it yields a dense estimator of \mathbf{A}_* , leading to a divergent statistical error. To overcome these issues, we propose a sparse EM algorithm to deal with the high dimensionality and to produce a sparse estimate of the transition matrix. Specifically, we consider a generalized Dantzig selector for Yule-Walker equation (Candes and Tao, 2007; Han et al., 2015),

$$\hat{\mathbf{A}}_k = \underset{\mathbf{A} \in \mathbb{R}^{p \times p}}{\operatorname{argmin}} \|\mathbf{A}\|_1, \text{ such that } \left\| \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbf{E}_{t,t+1,k} - \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbf{E}_{t,t,k} \mathbf{A}^\top \right\|_{\max} \leq \tau_k, \quad (4)$$

where τ_k is the tolerance parameter that is tuned via cross-validation. The optimization problem (4) is solved using linear programming in a row-by-row parallel fashion. We next update the variance estimates as,

$$\begin{aligned}\hat{\sigma}_{\eta,k}^2 &= \frac{1}{p(T-1)} \sum_{t=1}^{T-1} \left\{ \text{tr}(\mathbf{E}_{t+1,t+1;k}) - \text{tr}(\widehat{\mathbf{A}}_k \mathbf{E}_{t,t+1;k}) \right\}, \\ \hat{\sigma}_{\epsilon,k}^2 &= \frac{1}{pT} \sum_{t=1}^T \left\{ \mathbf{y}_t^\top \mathbf{y}_t - 2\mathbf{y}_t^\top \mathbf{E}_{t,k} + \text{tr}(\mathbf{E}_{t,t;k}) \right\},\end{aligned}\tag{5}$$

where $\mathbf{E}_{t,k} = \mathbb{E}\{\mathbf{x}_t \mid \{\mathbf{y}_t\}_{t=1}^T, \hat{\Theta}_{k-1}\}$ for $t \in [T]$, and (5) comes from taking derivative on $Q_y(\Theta \mid \hat{\Theta}_k)$. We terminate our sparse EM algorithm when the estimates are close enough in two consecutive iterations, e.g., $\min\{\|\hat{\mathbf{A}}_k - \hat{\mathbf{A}}_{k-1}\|_F, |\hat{\sigma}_{\eta,k} - \hat{\sigma}_{\eta,k-1}|, |\hat{\sigma}_{\epsilon,k} - \hat{\sigma}_{\epsilon,k-1}|\} \leq 10^{-3}$.

We summarize our sparse EM procedure in Algorithm 1.

Algorithm 1

Sparse EM algorithm for parameter estimation in model (1).

Initialization: $\hat{\Theta}_0 = \{\widehat{\mathbf{A}}_0, \hat{\sigma}_{\eta,0}^2, \hat{\sigma}_{\epsilon,0}^2\}$, and set $k = 1$.

repeat

1. E-step: Obtain $\mathbf{E}_{t,k}$, $\mathbf{E}_{t,t;k}$, and $\mathbf{E}_{t,t+1;k}$ via Kalman filter and smoothing, conditional on $\{\mathbf{y}_t\}_{t \in [T]}$ and $\hat{\Theta}_{k-1}$.
2. M-step:
 - 2.1. Compute $\hat{\mathbf{A}}_k$ by (4).
 - 2.2. Compute $\hat{\sigma}_{\eta,k}^2$ and $\hat{\sigma}_{\epsilon,k}^2$ by (5).
3. Collect $\hat{\Theta}_k = \{\hat{\mathbf{A}}_k, \hat{\sigma}_{\eta,k}^2, \hat{\sigma}_{\epsilon,k}^2\}$, and set $k = k + 1$.

until the stopping criterion is met.

2.2 Estimation consistency

We next establish the estimation consistency of our sparse EM estimators, and show that they achieve estimation errors within statistical precision required for the construction of test statistic and inference methods in subsequent steps.

Similar to Balakrishnan et al. (2017); Cai et al. (2019), we first introduce a concept, the population Q -function, as

$$Q(\Theta \mid \Theta^*) = \mathbb{E}\{Q_y(\Theta \mid \Theta^*) \mid \Theta^*\},$$

where the expectation is with respect to the observed data \mathbf{y}_t , and this population Q -function depends on the true parameters Θ^* and the sample size T . On the other hand, our definition is not exactly the same as that of Balakrishnan et al. (2017); Cai et al. (2019). They considered the limit of infinite i.i.d. observations, which naturally leads to the expectation of Q_y at a single observation by the law of large numbers. However, the temporal dependence in our problem makes the expectation of Q_y change with the sample size T . So we define

the population Q -function as the expectation at a finite T . This change brings additional technical difficulty for subsequent theoretical analysis.

We next introduce a sequence of intermediate estimators, $\Theta_{k+1} = \operatorname{argmax}_{\Theta} Q(\Theta | \hat{\Theta}_k)$. Note that Θ_{k+1} is obtained by maximizing the population Q -function $Q(\cdot | \hat{\Theta}_k)$ and can be viewed as a population-level estimator, whereas $\hat{\Theta}_{k+1}$ in the sparse EM algorithm is obtained by maximizing the finite-sample Q -function $Q_n(\cdot | \hat{\Theta}_k)$ and can be viewed as a perturbation to its population counterpart. Meanwhile, both Q -functions are conditioning on the sparse EM estimator $\hat{\Theta}_k$ from the previous iteration. We then break our theoretical analysis into two steps. We first characterize the contraction behavior of the intermediate estimator Θ_{k+1} at the population level. We next quantify the perturbation of the sparse EM estimator $\hat{\Theta}_{k+1}$ from Θ_{k+1} . The resulting error bound of the sparse EM estimator consists of two errors, a computational error and a statistical error. The first comes from the population behavior, and the latter measures the perturbation due to finite samples.

Our first step of theoretical analysis is to characterize the contraction behavior of Θ_{k+1} . Define the oracle auxiliary function $q(\Theta) = Q(\Theta | \Theta_*)$, and the maximizer $M(\Theta) = \operatorname{argmax}_{\Theta'} Q(\Theta' | \Theta)$. Also define a local neighborhood of the true parameters Θ_* , for some constants $\lambda \in (0, 1)$ and $r > 0$,

$$\mathcal{B}(\lambda, r) = \{ \{ \mathbf{A}, \sigma_\eta^2, \sigma_\epsilon^2 \} : |\sigma_\eta^2 - \sigma_{\eta,*}^2| \leq \lambda \sigma_{\eta,*}^2, |\sigma_\epsilon^2 - \sigma_{\epsilon,*}^2| \leq \lambda \sigma_{\epsilon,*}^2, \| \mathbf{A} - \mathbf{A}_* \|_{\max} \leq r \}. \quad (6)$$

Our first key insight is that, in a local neighborhood of the true parameter Θ_* , if $q(\cdot)$ is strongly concave, $Q(\cdot | \hat{\Theta}_k)$ is geometrically similar to $q(\cdot)$, and $M(\cdot)$ well behaves, then Θ_{k+1} is closer to the truth Θ_* than $\hat{\Theta}_k$. Denote the first and second-order partial derivatives of $q(\Theta)$ at any parameter entry θ in Θ as $\partial_\theta q(\Theta)$ and $\partial_\theta^2 q(\Theta)$. Then, for any $\Theta \in \mathcal{B}(\lambda, r)$, $\partial_\theta^2 q(\Theta)$ is upper bounded by a negative constant. Therefore, $q(\Theta)$ is strongly concave in $\mathcal{B}(\lambda, r)$. The next assumption characterizes the geometric similarity between $Q(\cdot | \hat{\Theta}_k)$ and $q(\cdot)$, as well as the behavior of $M(\cdot)$ in $\mathcal{B}(\lambda, r)$.

Assumption 1.—For any $\Theta \in \mathcal{B}(\lambda, r)$ and any entry of the parameter θ in Θ , assume that $|\partial_{\theta_M} q\{M(\Theta)\}| \leq \kappa |\partial_{\theta_*}^2 q(\Theta_*)| |\theta - \theta_*|$, for some constant $0 < \kappa < 1$, where θ_M and θ_* denote the corresponding parameter in $M(\Theta)$ and Θ_* , respectively.

We note that the above inequality always holds when $\Theta = \Theta_*$, even with $\kappa = 0$. When κ is strictly positive, intuitively, it is reasonable to extend this inequality over a local region $\mathcal{B}(\lambda, r)$ around Θ_* with some positive λ and r . A similar condition was also imposed in Balakrishnan et al. (2017). By Assumption 1, the strong concavity of $q(\Theta)$, and the fact that $\partial_{\theta_*} q(\Theta_*) = 0$, we have that,

$$|\partial_{\theta_*}^2 q(\Theta_*)| |\theta_{k+1} - \theta_*| \leq |\partial_{\theta_*} q(\Theta_*) - \partial_{\theta_{k+1}} q(\Theta_{k+1})| \leq \kappa |\partial_{\theta_*}^2 q(\Theta_*)| |\hat{\theta}_k - \theta_*|,$$

for any corresponding entries $\theta_{k+1}, \hat{\theta}_k$ of $\Theta_{k+1}, \hat{\Theta}_k$. Therefore, the population update Θ_{k+1} is closer to Θ_* than $\hat{\Theta}_k$ at rate κ .

Next, we need to ensure the above contraction property holds for all population updates Θ_k for any $k \geq 1$. That is, once the initial estimator $\hat{\Theta}_0$ locates in $\mathcal{B}(\lambda, r)$, then the population updates $\{\Theta_k\}_{k \geq 1}$ all locate in $\mathcal{B}(\lambda, r)$, and thus the contraction property applies to any Θ_k with $k \geq 1$. To achieve that, we need to ensure that $M(\Theta) \in \mathcal{B}(\lambda, r)$ for any $\Theta \in \mathcal{B}(\lambda, r)$. Let $\Sigma_y(\Theta) \in \mathbb{R}^{pT \times pT}$ be the covariance matrix of the stacked vector $\mathbf{y}_{[T]} = (\mathbf{y}_1^\top, \dots, \mathbf{y}_T^\top)^\top \in \mathbb{R}^{pT}$ conditioning on the parameter set Θ . Define $\Sigma_1(\Theta) = (T-1)^{-1} \mathbb{E} \left\{ \sum_{t=1}^{T-1} \mathbb{E}(\mathbf{x}_t \mathbf{x}_{t+1}^\top \mid \mathbf{y}_{[T]}, \Theta) \mid \Theta_* \right\}$, $\Sigma_0(\Theta) = (T-1)^{-1} \mathbb{E} \left\{ \sum_{t=1}^{T-1} \mathbb{E}(\mathbf{x}_t \mathbf{x}_t^\top \mid \mathbf{y}_{[T]}, \Theta) \mid \Theta_* \right\}$, $\Sigma_0^{(1)}(\Theta) = (T-1)^{-1} \mathbb{E} \left\{ \sum_{t=2}^T \mathbb{E}(\mathbf{x}_t \mathbf{x}_t^\top \mid \mathbf{y}_{[T]}, \Theta) \mid \Theta_* \right\}$. In a sense, the three terms can be viewed as expectations of $\mathbf{E}_{t,t+1,k}$ and $\mathbf{E}_{t,t,k}$ conditioning on true parameter if $\hat{\Theta}_{k-1} = \Theta$.

Assumption 2.—For any $\Theta = \{\mathbf{A}, \sigma_n^2, \sigma_c^2\} \in \mathcal{B}(\lambda, r)$, assume that

$$\begin{aligned} & \left\| \{\Sigma_1(\Theta)\}^\top \{\Sigma_0(\Theta)\}^{-1} - \mathbf{A}_* \right\|_{\max} \leq r, \left| \text{tr} \left[\{\Sigma_0^{(1)}(\Theta) - \{\Sigma_1(\Theta)\}^\top \{\Sigma_0(\Theta)\}^{-1} \Sigma_1(\Theta) \right] - p \lambda \sigma_n^{2*} \right| \leq p \lambda \sigma_n^{2*}, \text{ and} \\ & \left| p T \sigma_c^2 + \sigma_c^4 \text{tr} \left[\{\Sigma_y(\Theta)\}^{-1} \Sigma_y(\Theta_*) \{\Sigma_y(\Theta)\}^{-1} - \{\Sigma_y(\Theta)\}^{-1} \right] - p T \lambda \sigma_c^{2*} \right| \leq p T \lambda \sigma_c^{2*}. \end{aligned}$$

This assumption trivially holds if $\Theta = \Theta_*$ when $\lambda = r = 0$. When $\lambda, r > 0$, intuitively, it is reasonable to expect the assumption remains valid over a proper local region around Θ_* . Taking \mathbf{A}_* for instance, note that $\{\Sigma_1(\hat{\Theta}_k)\}^\top \{\Sigma_0(\hat{\Theta}_k)\}^{-1}$ is the population update of \mathbf{A}_* from $\hat{\Theta}_k$ and falls in the region $\mathcal{B}(\lambda, r)$ if $\hat{\Theta}_k \in \mathcal{B}(\lambda, r)$, thus the identity $\Sigma_1(\Theta) = \Sigma_0(\Theta) \mathbf{A}_*^\top + (T-1)^{-1} \mathbb{E} \left\{ \sum_{t=1}^{T-1} \mathbb{E}(\mathbf{x}_t \boldsymbol{\eta}_t^\top \mid \{\mathbf{y}_{t'}\}_{t'=1}^T, \Theta) \mid \Theta_* \right\}$ implies that this assumption essentially bounds the reminder $(T-1)^{-1} \{\Sigma_0(\Theta)\}^{-1} \mathbb{E} \left\{ \sum_{t=1}^{T-1} \mathbb{E}(\mathbf{x}_t \boldsymbol{\eta}_t^\top \mid \{\mathbf{y}_{t'}\}_{t'=1}^T, \Theta) \mid \Theta_* \right\}$, which should be small for a range of Θ_* , since $\mathbb{E}(\mathbf{x}_t \boldsymbol{\eta}_t^\top \mid \Theta_*) = 0$. A similar condition was imposed in Cai et al. (2019, Condition C1) on the initialization too. The explicit forms of λ, r are difficult to obtain in our case though, given the temporal dependence and the complicated high-dimensional matrices and their inverses.

Together, Assumptions 1 and 2 ensure that the population update Θ_{k+1} gets closer to the truth Θ_* than finite-sample update $\hat{\Theta}_k$ at a contraction rate κ . This subsequently leads to a geometrically decreasing computational error at rate κ in the error bound of the sparse EM estimator.

Our second step of theoretical analysis is to quantify the perturbation of the sparse EM estimator $\hat{\Theta}_{k+1}$ from its population counterpart Θ_{k+1} . We introduce the next two assumptions that characterize the temporal dependence and the sparsity of the model. In Assumption 3, $[\mathbf{M}]_{mn}$ denotes the (m, n) th $p \times p$ block matrix in $\mathbf{M} \in \mathbb{R}^{Tp \times Tp}$.

Assumption 3.—Assume $\sup_{\Theta \in \mathcal{B}(\lambda, r)} \max_{m \in [T]} \sum_{l=1}^T \left\| \left[\{\Sigma_y(\Theta)\}^{-1} \Sigma_y(\Theta_*) \{\Sigma_y(\Theta)\}^{-1} \right]_{ml} \right\|_{\max} < \infty$.

In addition, assume that

$$\sup_{\Theta \in \mathcal{B}(\lambda, r)} \max_{m \in [T]} \sum_{n=1}^T \left(\left\| \left[\Sigma_y(\Theta_*) \{\Sigma_y(\Theta)\}^{-1} \right]_{mn} \right\|_{\max} + \left\| \left[\Sigma_y(\Theta_*) \{\Sigma_y(\Theta)\}^{-1} \right]_{mn} \right\|_{\max} \right) < \infty.$$

This assumption constrains the temporal dependence, where the matrices in this assumption are the covariance matrices of the quadratic forms of \mathbf{y}_t 's. Note that the randomness in $\widehat{\Theta}_{k+1}$ comes from the average of the quadratic forms of \mathbf{y}_t 's, and the law of large numbers holds as long as the temporal dependence between the quadratic forms of \mathbf{y}_t 's is bounded. This assumption is reasonable as the inequalities that bound the spectral of the quadratic forms comply with the existing concentration theory (Negahban and Wainwright, 2011).

Next, consider a weakly sparse matrix space, $\mathcal{M}(r_q, R_q, r_1, R_1)$, defined as,

$$\left\{ \mathbf{M} \in \mathbb{R}^{p \times p}: \max_{j \in [p]} \sum_{i=1}^p |M_{ij}|^q \leq r_q, \sum_{i,j \in [p]} |M_{ij}|^q \leq R_q, \|\mathbf{M}\|_{l_1} \leq r_1, \|\mathbf{M}\|_1 \leq R_1 \right\}, \quad (7)$$

for some constants $0 \leq q < 1, r_q > 0, R_q > 0, r_1 > 0$, and $R_1 > 0$.

Assumption 4.—*There exist constants $q \in [0, 1), r_q > 0, R_q > 0$, and $r_1 > 0$, such that, for any $\Theta \in \mathcal{B}(\lambda, r)$, $\{\Sigma_0(\Theta)\}^{-1} \Sigma_1(\Theta) \in \mathcal{M}(r_q, R_q, r_1, R_1)$.*

This assumption imposes a weak sparsity constraint on the matrix $\{\Sigma_0(\Theta)\}^{-1} \Sigma_1(\Theta)$, the population update of \mathbf{A}_* , in that the matrix can be dense as long as there are only a few dominant entries and the rest entries are small. Besides, we allow r_q, R_q in $\mathcal{M}(r_q, R_q, r_1, R_1)$ to diverge in the subsequent theoretical development. This is much weaker than requiring the population update to be strictly sparse with only a few nonzero entries. This assumption is similar in spirit as the sparsity assumption in Cai et al. (2019), except that it involves a more complicated form due to the temporal dependence of the time series model.

Together, Assumptions 3 and 4 ensure that the sparse EM estimator $\widehat{\Theta}_{k+1}$ is not too far away from its population counterpart Θ_{k+1} , which contributes to the statistical error in the error bound of the sparse EM estimator.

Now we are ready to present the main theorem regarding the computational and statistical errors of our sparse EM estimator. The key idea is that, for any entry of the parameter space, we have that $|\theta_{k+1} - \theta_*| \leq \kappa |\hat{\theta}_k - \theta_*|$. Besides, define $\delta_\theta = \sup_{\hat{\theta}_k \in \mathcal{B}(\lambda, r)} |\theta_{k+1} - \hat{\theta}_{k+1}|$ as the distance between the population and finite-sample estimator of θ . Then, we have,

$$|\hat{\theta}_{k+1} - \theta_*| \leq |\theta_{k+1} - \theta_*| + |\hat{\theta}_{k+1} - \theta_{k+1}| \leq \kappa |\hat{\theta}_k - \theta_*| + \delta_\theta \leq \kappa^{k+1} |\hat{\theta}_0 - \theta_*| + \frac{1}{1 - \kappa} \delta_\theta,$$

in which the first term is the geometrically decaying computational error, and the second term is the statistical error. The next theorem gives a more precise summary.

Theorem 1.—*Suppose the following conditions hold.*

- a. *The initial parameter set $\widehat{\Theta}_0 = \{\widehat{\mathbf{A}}_0, \hat{\sigma}_{n,0}^2, \hat{\sigma}_{c,0}^2\}$ are in a neighborhood $\mathcal{B}(\lambda, r)$ that satisfies Assumptions 1, 2, 3, and 4 for some $\lambda \in (0, 1)$ and $r > 0$.*

- b. The tolerance parameter $\tau_l = c_l(r_l + 1)\sqrt{\log(p)/T}$ for some positive constant c_l , $\forall l \leq k$.
- c. The dimension of time series p and the length of series T satisfy that $C \log p \leq T$ for some positive constant C .

Then, the sparse EM estimator $\hat{\Theta}_k = \{\hat{\mathbf{A}}_k, \hat{\sigma}_{\eta,k}^2, \hat{\sigma}_{\epsilon,k}^2\}$ at the k th iteration satisfies that, for any constant $c_0 > 0$, there exist positive constants c_1 to c_5 such that the event

$$\begin{aligned} |\hat{\sigma}_{\epsilon,k}^2 - \sigma_{\epsilon,*}^2| &\leq \kappa^k |\hat{\sigma}_{\epsilon,0}^2 - \sigma_{\epsilon,*}^2| + \frac{c_1}{1 - \kappa} \|\boldsymbol{\Sigma}_y(\Theta_*)\|_2 \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_y(\Theta)\}^{-1}\|_2 \sqrt{\frac{\log p}{Tp}}, \\ |\hat{\sigma}_{\eta,k}^2 - \sigma_{\eta,*}^2| &\leq \kappa^k |\hat{\sigma}_{\eta,0}^2 - \sigma_{\eta,*}^2| + \frac{c_2}{1 - \kappa} \left[\|\boldsymbol{\Sigma}_y(\Theta_*)\|_2 \left(1 \vee \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_y(\Theta)\}^{-1}\|_2 \right) \sqrt{\frac{\log p}{Tp}} \right. \\ &\quad \left. + \frac{1}{p} \sqrt{\frac{\log p}{T}} \left[(r_1 \vee 1) R_1 + R_q \left((r_1 \vee 1) \sqrt{\frac{\log p}{T}} \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_0(\Theta)\}^{-1}\|_{l_1} \right)^{1-q} \right] \right], \\ \|\hat{\mathbf{A}}_k - \mathbf{A}_*\|_{\max} &\leq \kappa^k \|\hat{\mathbf{A}}_0 - \mathbf{A}_*\|_{\max} + \frac{c_3}{1 - \kappa} (r_1 \vee 1) \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_0(\Theta)\}^{-1}\|_{l_1} \sqrt{\frac{\log p}{T}}, \end{aligned}$$

$$\begin{aligned} \|\hat{\mathbf{A}}_k - \mathbf{A}_*\|_{l_\infty} &\leq \kappa^k \|\hat{\mathbf{A}}_0 - \mathbf{A}_*\|_{l_\infty} + \frac{c_4}{1 - \kappa} r_q \left[(r_1 \vee 1) \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_0(\Theta)\}^{-1}\|_{l_1} \sqrt{\frac{\log p}{T}} \right]^{1-q}, \\ \|\hat{\mathbf{A}}_k - \mathbf{A}_*\|_{r,2} &\leq \kappa^k \|\hat{\mathbf{A}}_0 - \mathbf{A}_*\|_{r,2} + \frac{c_5}{1 - \kappa} \sqrt{r_q} \left[(r_1 \vee 1) \sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_0(\Theta)\}^{-1}\|_{l_1} \sqrt{\frac{\log p}{T}} \right]^{1-\frac{q}{2}}, \end{aligned}$$

happens with probability at least $1 - p^{-c_0}$.

We make some remarks. First, the non-asymptotic error bound portrays the estimation error of sparse EM at each iteration, and reveals the interplay between the computational efficiency and the statistical rate of convergence. After a sufficient number of iterations, the computational error is to be dominated by the statistical error. Second, the statistical errors are all vanishing if $\log p$ scales with T , and they decay sufficiently fast in terms of p and T for subsequent statistical inference, even when r_q and R_q in $\mathcal{M}(r_q, R_q, r_1, R_1)$ in (7) diverge. Third, the statistical errors for \mathbf{A}_* do not explicitly display the sparsity. This information is hidden in $\mathcal{M}(r_q, R_q, r_1, R_1)$. Moreover, since the update of $\sigma_{\eta,*}^2$ involves the update of \mathbf{A}_* , the statistical error of $\sigma_{\eta,*}^2$ is more complicated than that of $\sigma_{\epsilon,*}^2$. Finally, we observe the phenomenon of ‘‘blessing of dimensionality’’, in that the statistical errors of $\sigma_{\epsilon,*}^2$ and $\sigma_{\eta,*}^2$ decrease when the dimension p grows under a fixed sample size T . In general, we allow p to diverge at an exponential rate of T as both approach infinity.

3 Test Statistics

We next construct a Gaussian matrix as our test statistic for the transition matrix inference in our high-dimensional VAR with measurement error. Given model (1), we observe a time series of \mathbf{y}_t that follows an autoregressive structure, $\mathbf{y}_{t+1} = \mathbf{A}_* \mathbf{y}_t + \boldsymbol{\epsilon}_t$, with the error term $\boldsymbol{\epsilon}_t = -\mathbf{A}_* \boldsymbol{\epsilon}_t + \boldsymbol{\epsilon}_{t+1} + \boldsymbol{\eta}_t$. Then the lag-1 auto-covariance of the error $\boldsymbol{\epsilon}_t$ is of the form,

$$\Sigma_e = \text{Cov}(\mathbf{e}_t, \mathbf{e}_{t-1}) = -\sigma_{e,*}^2 \mathbf{A}_*.$$

This suggests that we can apply the covariance testing methods on Σ_e to infer transition matrix \mathbf{A}_* . However, \mathbf{e}_t is not directly observed. Define generic estimators of Θ_* by $\{\hat{\mathbf{A}}, \hat{\sigma}_{e,*}^2, \hat{\sigma}_\eta^2\}$. We use them to reconstruct this error, and obtain the sample lag-1 auto-covariance estimator,

$$\hat{\Sigma}_e = \frac{1}{T-2} \sum_{t=2}^{T-1} \hat{\mathbf{e}}_t \hat{\mathbf{e}}_{t-1}^\top, \quad \text{where } \hat{\mathbf{e}}_t = \mathbf{y}_{t+1} - \hat{\mathbf{A}} \mathbf{y}_t - \frac{1}{T-1} \sum_{t'=1}^{T-1} (\mathbf{y}_{t'+1} - \hat{\mathbf{A}} \mathbf{y}_{t'}).$$

This sample estimator $\hat{\Sigma}_e$, nevertheless, involves some bias due to the reconstruction of the error term, and also an inflated variance due to the temporal dependence of the time series data. We next explicitly quantify such bias and variance, by characterizing the non-asymptotic behavior of $\hat{\Sigma}_e$, which eventually leads to our Gaussian matrix test statistic.

Denote the maximal row-wise ℓ_1 estimation error as $\Delta_1 = \|\mathbf{A}_* - \hat{\mathbf{A}}\|_{\ell_1}$, and the maximal row-wise Euclidean estimation error as $\Delta_2 = \|\mathbf{A}_* - \hat{\mathbf{A}}\|_{r,2}$. The next proposition characterizes the non-asymptotic behavior of $\hat{\Sigma}_e$.

Proposition 1.

For any constant $c > 0$, there exist positive constants c_1 to c_3 , such that, when $T \geq c_1 \log p$,

$$\mathbb{P} \left\{ \left\| \hat{\Sigma}_e + (\sigma_{\eta,*}^2 + \sigma_{e,*}^2) \hat{\mathbf{A}} - \frac{1}{T-2} \sum_{t=2}^{T-1} \mathbf{e}_t \mathbf{e}_{t-1}^\top \right\|_{\max} \leq c_2 \left(\Delta_{1,s_r} \sqrt{\frac{\log p}{T}} + \Delta_2^2 + \frac{\log p}{T} \right) \right\} \geq 1 - c_3 p^{-c},$$

where $s_r = \max_{i \in [p]} |\{j : A_{*,ij} \neq 0\}|$ is the maximal row-wise sparsity of \mathbf{A}_* .

This proposition suggests using $\sqrt{T-2} \hat{\Sigma}_e$ to construct the Gaussian matrix test statistic, since $(T-2)^{-1/2} \sum_{t=2}^{T-1} (\mathbf{e}_t \mathbf{e}_{t-1}^\top - \mathbb{E} \mathbf{e}_t \mathbf{e}_{t-1}^\top)$ converges to a zero-mean Gaussian matrix by the central limit theorem. The max norm error of the sparse EM estimator and the fact $\mathbb{E} \mathbf{e}_t \mathbf{e}_{t-1}^\top = -\sigma_{e,*}^2 \mathbf{A}_*$ further imply that the non-vanishing bias of $\sqrt{T-2} \hat{\Sigma}_e$ is $\sqrt{T-2} \{-(\sigma_{\eta,*}^2 + \sigma_{e,*}^2) \hat{\mathbf{A}} + \sigma_{\eta,*}^2 \mathbf{A}_*\}$, which can be estimated by $\sqrt{T-2} \{-(\hat{\sigma}_\eta^2 + \hat{\sigma}_e^2) \hat{\mathbf{A}} + \hat{\sigma}_\eta^2 \mathbf{A}_0\}$ under the null hypothesis. Meanwhile, after the bias correction and some direct calculation of the entry-wise variance of $(T-2)^{-1/2} \sum_{t=2}^{T-1} (\mathbf{e}_t \mathbf{e}_{t-1}^\top - \mathbb{E} \mathbf{e}_t \mathbf{e}_{t-1}^\top)$, the entry-wise limit variance of $\sqrt{T-2} \hat{\Sigma}_e$ is,

$$\begin{aligned} \sigma_{e,ij}^2 &= (\sigma_{e,*}^2 + \sigma_{\eta,*}^2)^2 + \sigma_{e,*}^4 A_{2,ij}^2 + 2\sigma_{e,*}^4 A_{*,ii} A_{*,jj} + \sigma_{e,*}^4 \|\mathbf{A}_{*,i}\|_2 \|\mathbf{A}_{*,j}\|_2 \\ &\quad + (\sigma_{e,*}^4 + \sigma_{e,*}^2 \sigma_{\eta,*}^2) (\|\mathbf{A}_{*,i}\|_2 + \|\mathbf{A}_{*,j}\|_2), \quad i, j \in [p]. \end{aligned}$$

Plugging in the estimators $\{\hat{\mathbf{A}}, \hat{\sigma}_{e,*}^2, \hat{\sigma}_\eta^2\}$ into the above equation, we obtain the corresponding estimator $\hat{\sigma}_{ij}^2$. We also comment that, one can use any generic estimators $\{\hat{\mathbf{A}}, \hat{\sigma}_{e,*}^2, \hat{\sigma}_\eta^2\}$ to estimate the bias and variance of $\sqrt{T-2} \hat{\Sigma}_e$. Later, we present the sufficient conditions on the

estimation precision of the generic estimators, so to achieve the desired theoretical properties of inference. We then show that the estimators from our sparse EM algorithm satisfy those conditions.

Now, we construct the Gaussian matrix test statistic \mathbf{H} , whose (i, j) th entry is,

$$H_{ij} = \frac{\sum_{t=2}^{T-1} \{\hat{e}_{t,i}\hat{e}_{t-1,j} + (\hat{\sigma}_\eta^2 + \hat{\sigma}_\epsilon^2)\hat{A}_{ij} - \hat{\sigma}_\eta^2 A_{0,ij}\}}{\sqrt{T-2}\hat{\sigma}_{ij}}, \quad i, j \in [p]. \quad (8)$$

Denote the estimation errors, $\Delta_\epsilon = |\hat{\sigma}_\epsilon^2 - \sigma_{\epsilon,*}^2|$, $\Delta_\eta = |\hat{\sigma}_\eta^2 - \sigma_{\eta,*}^2|$, and $\Delta_\sigma = \max_{i,j \in [p]} |\hat{\sigma}_{ij}^2 - \sigma_{i,j}^2|$. The next theorem provides the sufficient conditions to guarantee the asymptotic standard normality of H_{ij} under the null hypothesis.

Theorem 2.

Suppose the following conditions hold.

- a. *The estimation errors satisfy that $\Delta_1 = o_p\{s_r^{-1}(\log p)^{-1/2}\}$, $\Delta_2 = o_p(T^{-1/4})$, $\Delta_\epsilon = o_p(T^{-1/2})$, $\Delta_\eta = o_p(T^{-1/2})$, and $\Delta_\sigma = o_p(1)$.*
- b. *The dimension of time series p and the length of series T satisfy that $\log p = o(T^{1/2})$. Then*

$$\frac{\sum_{t=2}^{T-1} \{\hat{e}_{t,i}\hat{e}_{t-1,j} + (\hat{\sigma}_\eta^2 + \hat{\sigma}_\epsilon^2)\hat{A}_{ij} - \hat{\sigma}_\eta^2 A_{*,ij}\}}{\sqrt{T-2}\hat{\sigma}_{ij}} \xrightarrow{d} N(0, 1)$$

uniformly for $i, j \in [p]$ as $p, T \rightarrow \infty$.

Here the normality holds when the dimension p grows at the exponential rate of \sqrt{T} . The matrix \mathbf{H} is to serve as the test statistic for the subsequent inference procedures.

4 Transition Matrix Inference

4.1 Global inference

We first develop a testing procedure for the global hypotheses (2). The key observation is that the squared maximum entry of a zero mean normal vector converges to a Gumbel distribution (Cai and Jiang, 2011). Specifically, we construct the global test statistic as,

$$G_{\mathcal{S}} = \max_{(i,j) \in \mathcal{S}} H_{ij}^2$$

The next theorem states that the asymptotic null distribution of $G_{\mathcal{S}}$ is Gumbel. We again state the sufficient conditions required for generic estimators $\{\hat{\mathbf{A}}, \hat{\sigma}_\epsilon^2, \hat{\sigma}_\eta^2\}$ first, and show later that the sparse EM estimators satisfy these conditions.

Theorem 3.—*Suppose the following conditions hold.*

- a. The estimation errors satisfy that $\Delta_1 = o_p\{(s, \log p)^{-1}\}$, $\Delta_2 = o_p\{(T \log p)^{-1/4}\}$, $\Delta_\epsilon = o_p\{(T \log p)^{-1/2}\}$, $\Delta_\eta = o_p\{(T \log p)^{-1/2}\}$ and $\Delta_\sigma = o_p\{(\log p)^{-1}\}$.
- b. The dimension of time series p and the length of series T satisfy that $\log p = o(T^{1/7})$. Then, under the global null hypothesis in (2), for any $\mathcal{S} \subseteq [p] \times [p]$ and any $x \in \mathbb{R}$,

$$\lim_{|\mathcal{S}| \rightarrow \infty} \mathbb{P}\left\{G_{\mathcal{S}} - 2 \log |\mathcal{S}| + \log \log |\mathcal{S}| \leq x\right\} = \exp\left\{-\exp(-x/2)/\sqrt{\pi}\right\}.$$

We note that the condition (a) about the estimation consistency in Theorem 3 is stronger than that in Theorem 2 for the asymptotic normality. This is because the Gumbel convergence is built upon the normality property that needs to be guaranteed first. Based on this limiting null distribution, we define the asymptotic α -level test as,

$$\Psi_\alpha = \mathbb{1}\{G_{\mathcal{S}} > 2 \log |\mathcal{S}| - \log \log |\mathcal{S}| - \log \pi - 2 \log\{-\log(1 - \alpha)\}\}.$$

We reject the global null if $\Psi_\alpha = 1$.

Next, we study the asymptotic power of the test Ψ_α . Toward that end, we introduce a parameter class of alternatives,

$$\mathcal{A}(c, \mathcal{S}) = \left\{ \{\mathbf{A}_*, \sigma_{\eta_*}^2, \sigma_{\epsilon_*}^2\} : \max_{(i,j) \in \mathcal{S}} \frac{\sigma_{\eta_*}^2 \delta_{ij}}{\sigma_{\epsilon_*}^2} \geq c \sqrt{\frac{\log |\mathcal{S}|}{T}} \right\}, \quad (9)$$

where $\delta_{ij} = |A_{*,ij} - A_{0,ij}|$ is the distance between the null and the true transition matrix. The class $\mathcal{A}(c, \mathcal{S})$ requires at least one entry in \mathcal{S} has a proper signal-to-noise ratio against the null. Note that this is a very large class, because the imposed magnitude $\sqrt{\log |\mathcal{S}|/T}$ is vanishing, and it only requires one entry to satisfy. The next theorem shows that Ψ_α has the power converging to one uniformly over $\mathcal{A}(2\sqrt{2}, \mathcal{S})$. Together, Theorems 3 and 4 establish the asymptotic size and power, and thus the consistency of the global test Ψ_α .

Theorem 4.—*Suppose the same conditions in Theorem 3 hold. Then*

$$\inf_{\{\mathbf{A}_*, \sigma_{\eta_*}^2, \sigma_{\epsilon_*}^2\} \in \mathcal{A}(2\sqrt{2}, \mathcal{S})} \mathbb{P}(\Psi_\alpha = 1) \rightarrow 1, \quad \text{as } |\mathcal{S}| \rightarrow \infty.$$

Next, we show that, when we employ the sparse EM estimators developed in Section 2, we can obtain the same desired results as in Theorems 3 and 4. Recall the sparse EM estimators at iteration k are denoted as $\{\hat{\mathbf{A}}_k, \hat{\sigma}_{\eta,k}^2, \hat{\sigma}_{\epsilon,k}^2\}$. Plugging in these estimators yields the corresponding sparse EM estimator $\hat{\sigma}_{ij,k}^2$ of σ_{ij}^2 . Denote the global test statistic and the α -level test based on these sparse EM estimators as $G_{\mathcal{S}, \text{sEM}}$ and $\Psi_{\alpha, \text{sEM}}$, respectively. The next proposition establishes their size and power properties.

Proposition 2.—*Suppose the following conditions hold.*

1. *The initial parameter set $\hat{\Theta}_0 = \{\hat{\mathbf{A}}_0, \hat{\sigma}_{n,0}^2, \hat{\sigma}_{e,0}^2\}$ are in a neighborhood $\mathcal{B}(\lambda, r)$ that satisfies Assumptions 1, 2, 3, and 4 for some $\lambda \in (0, 1)$ and $r > 0$.*
2. *The parameters in (7) satisfy that $R_1 < \infty, r_q = o[T^{\frac{1}{2}} - \frac{q}{2} / \{s_r(\log p)^{\frac{3}{2}} - \frac{q}{2}\}]$ and $R_q = o[pT^{\frac{1}{2}} - \frac{q}{2} / (\log p)^{\frac{3}{2}} - \frac{q}{2}]$. Moreover, $\sup_{\Theta \in \mathcal{B}(\lambda, r)} \|\{\boldsymbol{\Sigma}_y(\Theta)\}^{-1}\|_2 < \infty$, and $\|\boldsymbol{\Sigma}_y(\Theta_*)\|_2 < \infty$.*
3. *The tolerance parameter $\tau_l = c_l(r_l + 1)\sqrt{\log(p)/T}$ for some positive constant $c_l, \forall l \leq k$.*
4. *The iteration $k \geq \lceil C \log\{(T \log p) \vee (s_r \log p)\} \rceil$ for some positive constant C .*
5. *The dimension of time series p and the length of series T satisfy that $\log p = o(T^{1/7})$.*

Then, under the global null hypothesis in (2), for any $\mathcal{S} \subseteq [p] \times [p]$,

$$\lim_{|\mathcal{S}| \rightarrow \infty} \mathbb{P}\left(G_{\mathcal{S}, \text{sEM}} - 2 \log |\mathcal{S}| + \log \log |\mathcal{S}| \leq x\right) = \exp\{-\exp(-x/2)/\sqrt{\pi}\} \text{ for any } x \in \mathbb{R},$$

$$\inf_{\{\mathbf{A}^*, \sigma_{n,*}^2, \sigma_{e,*}^2\} \in \mathcal{A}(2\sqrt{2}, \mathcal{S})} \mathbb{P}(\Psi_{\alpha, \text{sEM}} = 1) \rightarrow 1, \text{ as } |\mathcal{S}| \rightarrow \infty.$$

The conditions for this proposition essentially combine those of Theorems 1 and 3. When the number of iterations k is large enough, the statistical error is to dominate the error bound of the sparse EM estimators, and this bound decays sufficiently fast to ensure the properties of the global testing procedure. The requirements on the statistical error of sparse EM are reasonable, in that r_q and R_q are allowed to diverge at certain rates. Moreover, we consider a finite R_1 here for simplicity, though it is possible to relax the sufficient condition to let R_1 diverge too.

4.2 Simultaneous inference with FDR control

We next develop a testing procedure for the simultaneous hypotheses (3) with a proper FDR control. Let $\mathcal{H}_0 = \{(i, j) : A_{*,ij} = A_{0,ij}, (i, j) \in \mathcal{S}\}$ denote the set of true null hypotheses, and $\mathcal{H}_1 = \{(i, j) : (i, j) \in \mathcal{S}, (i, j) \notin \mathcal{H}_0\}$ denote the set of true alternatives. The test statistic H_{ij} follows a standard normal distribution when $H_{0,ij}$ holds, and as such, we reject $H_{0,ij}$ if $|H_{ij}| > t$ for some thresholding value $t > 0$. Let $R_{\mathcal{S}}(t) = \sum_{(i,j) \in \mathcal{S}} \mathbb{1}\{|H_{ij}| > t\}$ denote the number of rejections at t . Then the false discovery proportion (FDP) and the false discovery rate (FDR) in our simultaneous testing problem are,

$$\text{FDP}_{\mathcal{S}}(t) = \frac{\sum_{(i,j) \in \mathcal{H}_0} \mathbb{1}\{|H_{ij}| > t\}}{R_{\mathcal{S}}(t) \vee 1}, \text{ and } \text{FDR}_{\mathcal{S}}(t) = \mathbb{E}\{\text{FDP}_{\mathcal{S}}(t)\}.$$

An ideal choice of the threshold t is to reject as many true positives as possible, while controlling the false discovery at the pre-specified level β . That is, we choose

$\inf\{t > 0 : \text{FDP}_{\mathcal{S}}(t) \leq \beta\}$ as the threshold. However, \mathcal{H}_0 in $\text{FDP}_{\mathcal{S}}(t)$ is unknown. Observing that $\mathbb{P}(|H_{ij}| > t) \approx 2\{1 - \Phi(t)\}$ by Theorem 2, where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution, we estimate the false rejections $\sum_{(i,j) \in \mathcal{H}_0} \mathbb{1}\{|H_{ij}| > t\}$ in $\text{FDP}_{\mathcal{S}}(t)$ using $\{2 - 2\Phi(t)\}|\mathcal{S}|$. Moreover, we restrict the search of t to the range $(0, \sqrt{2 \log |\mathcal{S}|}]$, since $\mathbb{P}(\hat{t} \text{ exists in } (0, \sqrt{2 \log |\mathcal{S}|}]) \rightarrow 1$ as we show later in the proof of Theorem 5. We summarize our simultaneous testing procedure in Algorithm 2.

Next, we study the asymptotic FDR control of Algorithm 2. We need two assumptions.

Assumption 5.—*There exist positive constants u_1 and u_2 , such that*

$$\left| \left\{ (i, j) : (i, j) \in \mathcal{H}_1, \frac{\sigma_n^2 \delta_{ij}}{\sigma_{n,ij}} > (4 + u_1) \sqrt{\frac{\log p}{T}} \right\} \right| \geq u_2 \sqrt{\log \log |\mathcal{S}|}.$$

Algorithm 2

Simultaneous inference with FDR control.

1. Calculate H_{ij} for all $(i, j) \in \mathcal{S}$.
2. Compute the thresholding value,

$$\hat{t} = \inf \left\{ 0 < t \leq \sqrt{2 \log |\mathcal{S}|} : \frac{\{2 - 2\Phi(t)\}|\mathcal{S}|}{R_{\mathcal{S}}(t) \vee 1} \leq \beta \right\}.$$

If \hat{t} does not exist, set $\hat{t} = \sqrt{2 \log |\mathcal{S}|}$.

3. For all $(i, j) \in \mathcal{S}$, reject $H_{0,ij}$ if $|H_{ij}| > \hat{t}$.

This assumption requires a reasonable number of alternatives in \mathcal{S} . Intuitively, if the number of alternatives is too small, then $\sum_{(i,j) \in \mathcal{H}_0} \mathbb{1}\{|H_{ij}| > t\} \approx R_{\mathcal{S}}(t)$ for any t , and the resulting FDR is close to one regardless thresholding value. This assumption is rather mild, since the required number is logarithm of logarithm of $|\mathcal{S}|$. Liu and Shao (2014) showed that this assumption is nearly necessary in the sense that the FDR control for large-scale simultaneous testing would fail if the number of true alternatives is fixed.

Assumption 6.—*For some constants $0 < \nu < (1 - \bar{\sigma})/(1 + \bar{\sigma})$, $\gamma > 0$, and $u > 0$, we have $|\{(i_1, j_1), (i_2, j_2)\} : |\tilde{\sigma}_{i_1 j_1, i_2 j_2}| > (\log |\mathcal{S}|)^{-2 - \gamma}; (i_1, j_1) \neq (i_2, j_2); (i_1, j_1), (i_2, j_2) \in \mathcal{H}_0\}| \leq u |\mathcal{S}|^{1 + \nu}$, where $\tilde{\sigma}_{i_1 j_1, i_2 j_2}$ is the limit covariance between $H_{i_1 j_1}$ and $H_{i_2 j_2}$, for $(i_1, j_1) \neq (i_2, j_2) \in \mathcal{S}$, and $\bar{\sigma} = \max_{(i_1, j_1) \neq (i_2, j_2); (i_1, j_1), (i_2, j_2) \in \mathcal{H}_0} |\tilde{\sigma}_{i_1 j_1, i_2 j_2}|$.*

This assumption bounds the number of strongly correlated entries in the null hypotheses. The bound, $|\mathcal{S}|^{1 + \nu}$, is weak, since there are $|\mathcal{S}|^2$ pairs in total and the majority of them are allowed to be strongly correlated. A similar assumption was adopted in Xia et al. (2018) to ensure the FDR control consistency. The explicit expression of $\tilde{\sigma}_{i_1 j_1, i_2 j_2}$ is given in the proof of Theorem 5.

The next theorem shows that the simultaneous testing procedure in Algorithm 2 controls both FDR and FDP. We again state the sufficient conditions required for any estimators $\{\hat{\mathbf{A}}, \hat{\sigma}_\epsilon^2, \hat{\sigma}_\eta^2\}$ first, then show that the sparse EM estimators satisfy these conditions.

Theorem 5.—*Suppose the following conditions hold.*

- a. *Suppose Assumptions 5 and 6 hold.*
- b. *The estimation errors satisfy the precision requirements in (a) of Theorem 3.*
- c. *Suppose $|\mathcal{H}_0| \geq c_1 |\mathcal{S}|$ for some positive constant c_1 .*
- d. *The dimension of time series p and the length of series T satisfy that $p \leq T^{c_2}$ for some positive constant c_2 .*

Then, for simultaneous hypotheses (3), for any $\mathcal{S} \subseteq [p] \times [p]$,

$$\lim_{|\mathcal{S}| \rightarrow \infty} \frac{FDR_{\mathcal{S}}(\hat{t})}{\beta |\mathcal{H}_0| / |\mathcal{S}|} = 1, \text{ and } \frac{FDP_{\mathcal{S}}(\hat{t})}{\beta |\mathcal{H}_0| / |\mathcal{S}|} \xrightarrow{p} 1 \text{ as } |\mathcal{S}| \rightarrow \infty.$$

Compared to the global testing, the estimation consistency condition (b) is the same for the simultaneous testing. Meanwhile, the simultaneous testing places some additional requirements on the number of alternatives as in condition (c) and Assumption 5, and the entry dependence as in Assumption 6. In addition, the dimension p grows at the polynomial rate of the sample size T , as in condition (d). These requirements are reasonable because, intuitively, the global testing only deals with the maximum entry, whereas the simultaneous testing tackles every individual entry. As such, the simultaneous testing relies more on the dependence structure among the entries, and needs a larger sample size than the global testing. Finally, the slight deflation $\beta |\mathcal{H}_0| / |\mathcal{S}|$ in the limiting FDR comes from substituting $|\mathcal{H}_0|$ with $|\mathcal{S}|$ in the false rejection approximation.

Next, we show that, when we employ the sparse EM estimators developed in Section 2, we can obtain the same properties as in Theorem 5.

Proposition 3.—*Suppose the following conditions hold.*

- a. *Suppose Assumptions 5 and 6 hold.*
- b. *Suppose the conditions (a) to (d) in Proposition 2 hold.*
- c. *Suppose $|\mathcal{H}_0| \geq c_1 |\mathcal{S}|$ for some positive constant c_1 .*
- d. *The dimension of time series p and the length of series T satisfy that $p \leq T^{c_2}$ for some positive constant c_2 .*

Then, for simultaneous hypotheses (3), for any $\mathcal{S} \subseteq [p] \times [p]$,

$$\lim_{|\mathcal{S}| \rightarrow \infty} \frac{FDR_{\mathcal{S}}(\hat{t})}{\beta |\mathcal{H}_0| / |\mathcal{S}|} = 1, \text{ and } \frac{FDP_{\mathcal{S}}(\hat{t})}{\beta |\mathcal{H}_0| / |\mathcal{S}|} \xrightarrow{p} 1 \text{ as } |\mathcal{S}| \rightarrow \infty.$$

The conditions for this proposition essentially combine those of Theorems 1 and 5. The requirement (b) on the sparse EM algorithm is the same as that for the global testing.

5 Simulations

5.1 Setup

We carry out intensive simulations to study the finite-sample performance of our proposed method. We generate the data following model (1). We consider four common network structures for the transition matrix \mathbf{A}_* : banded, Erdős-Rényi, stochastic block, and hub, as shown in Figure 1. We first fix $\sigma_{\epsilon,*} = \sigma_{\eta,*} = 0.2$, $\|\mathbf{A}_*\|_2 = 0.97$, and vary the dimension and sample size $(p, T) = (30, 500), (50, 500), (50, 1000), (70, 1000)$. Next, we fix $p = 50, T = 1000, \sigma_{\epsilon,*} = \sigma_{\eta,*} = 0.2$ and vary the signal strength $\|\mathbf{A}_*\|_2 = 0.7, 0.8, 0.9, 0.97$. Finally, we fix $p = 50, T = 1000, \|\mathbf{A}_*\|_2 = 0.97$, and vary the noise level $(\sigma_{\epsilon,*}, \sigma_{\eta,*}) = (0.1, 0.1), (0.2, 0.2), (0.3, 0.3), (0.4, 0.4)$.

5.2 Parameter estimation

We first report the estimation accuracy of our sparse EM. The tuning of the tolerance parameter τ_k in (4) is done by cross-validation, where we use the first 25% of data points for testing, the last 60% for training, and the middle 15% discarded to reduce the temporal dependence between the training and testing samples. We choose the value of τ_k that minimizes the average prediction error of the testing samples. We find our algorithm converges fast, usually within 10 iterations.

We compare our method with three alternative solutions, including the standard EM without sparsity constraint, the Lasso estimator (Hsu et al., 2008), and the Dantzig estimator (Han et al., 2015), both of which were designed for VAR without measurement error. We evaluate the estimation accuracy by the Frobenius error $\|\hat{\mathbf{A}} - \mathbf{A}_*\|_F$. Figure 2 reports the average estimation accuracy out of 1000 data replications for the varying (p, T) , the varying signal strength $\|\mathbf{A}_*\|_2$, and the varying noise level $(\sigma_{\epsilon,*}, \sigma_{\eta,*})$, respectively. It is seen that our proposed sparse EM achieves the smallest estimation error across all settings, except when the noise level is close to 0. For the case, the model reduces to a standard VAR model with little measurement error, and the Lasso and Dantzig estimators should work the best. Moreover, our method performs similarly under different network structures, reflecting its robustness with respect to the connectivity patterns. We also consider other error norms for \mathbf{A}_* and the estimation accuracy for $\sigma_{\epsilon,*}, \sigma_{\eta,*}$. The results show the same qualitative patterns as Figure 2, and are thus omitted.

5.3 Global and simultaneous inference

We next evaluate the performance of our global and simultaneous inference procedures. Table 1 reports the empirical size and power based on 1000 data replications, with the significant level set at $\alpha = 5\%$. It is seen that our global test maintains a reasonable control of the size, and at the same time achieves a good power. Table 2 reports the average false discovery proportion and the average true positive rate based on 1000 data replications, with

the FDR level set at 5%. It is seen that our simultaneous test achieves both a high true positive rate and a low false discovery proportion.

6 Brain Connectivity Analysis

We illustrate the proposed method with a brain connectivity study based on task-evoked fMRI. The data is part of the Human Connectome Project (HCP, Van Essen et al., 2013), whose overarching objective is to understand brain connectivity patterns of healthy adults. We study the fMRI scans of two individual subjects of the same age and sex and both participating the same story-math task. The task consists of blocks of auditory stories and addition-subtraction calculations, and requires the participant to answer a series of questions. An accuracy score is given at the end based on the participant's answers. The performance of the two subjects differ considerably, with one achieving the perfect score and the other getting only about half correct. We aim to estimate and infer the brain connectivity networks of the two subjects and compare between them. We have preprocessed the fMRI data following the pipeline of Glasser et al. (2013). The resulting data for each subject are $p = 264$ time series, corresponding to 264 brain regions-of-interest following the brain atlas of Power et al. (2011). The length of each time series is $n = 316$. The 264 brain regions have been further grouped into 14 functional modules (Smith et al., 2009): auditory (AD), cerebellar (CR), cingulo-opercular task control (CO), default mode (DM), dorsal attention (DAT), fronto-parietal task control (FP), memory retrieval (MR), salience (SA), sensory/somatomotor hand (SMH), sensory/somatomotor mouth (SMM), subcortical (SC), uncertain (UN), ventral attention (VA), and visual (VS). Each module possesses a relatively autonomous functionality, and complex brain tasks are believed to perform through coordinated collaborations among the modules.

We begin with the global test for each subject separately. The p -values for the global test for both subjects are smaller than 10^{-15} , indicating that at least one pair of brain regions have statistically significant connectivity. We then apply the simultaneous test, with the the FDR set at 0.001. First of all, we have identified more within-module connections than the between-module connections (294 out of 7700 or 3.8% versus 961 out of 61936 or 1.6% for the high-accuracy subject, and 376 out of 7700 or 4.9% versus 1350 out of 61936 or 2.2% for the low-accuracy subject). The partition of the brain regions to the functional modules has been fully based on the biological knowledge, and our finding lends some numerical support to this partition. Second, the majority of within-module connections are concentrated on eight functional modules. Moreover, when comparing between the two subjects among those modules, we find that the high-accuracy subject has more within-module connections than the low-accuracy subject for the following functional modules: visual (118 versus 27 out of 961), salience (29 versus 11 out of 324), cingulo-opercular task control (17 versus 3 out of 196), and memory retrieval (6 versus 2 out of 25) modules. Such findings suggest that the high-accuracy subject has exhibited more intensive neural activities for processing visual imagery, memory retrieval, tonic alertness and executive control when performing the story-math task, which agrees with the literature (Sadaghiani and D'Esposito, 2015; Luo et al., 2014). On the other hand, we find that the high-accuracy subject has fewer connections than the low-accuracy subject for the following functional modules: default mode (25 versus 200 out of 3364), fronto-parietal task control (15 versus

37 out of 625), auditory (2 versus 8 out of 169), and subcortical (19 versus 49 out of 169) modules. These findings again agree with the literature, in that these modules have been found strongly associated with the language and reasoning type tasks (Schultz and Cole, 2016), and the high-accuracy subject has exhibited less brain activity interplay related to auditory processing and mind wandering (Van Praag et al., 2017). Figure 3 shows the identified connectivity patterns for the two subjects, and Figure 4 shows the corresponding brain regions visualized using BrainNet Viewer (Xia et al., 2013).

7 Discussion

In this article, we study both global and simultaneous inferences of the transition matrix under the high-dimensional vector autoregression model with measurement error. There is no existing solution, and our proposal makes a useful contribution for scientific applications such as brain connectivity analysis and others. The technical tools we develop are also of independent interest, and can facilitate the development of inferential procedures for other models involving latent variables or correlated observations. We next make some remarks regarding our model assumptions, potential limitations, and possible extensions.

We have primarily focused on a lag-1 autoregressive structure in this article. Meanwhile, our proposal can be extended in a relatively straightforward fashion to a more general lag structure. Specifically, suppose the number of lags is d . Then the latent process in model (1) becomes $\mathbf{x}_t = \sum_{l=1}^d \mathbf{A}_{l,*} \mathbf{x}_{t-l} + \boldsymbol{\eta}_{t-1}$, and the problem of interest becomes testing $\mathbf{A}_{1,*}, \dots, \mathbf{A}_{d,*}$. This lag- d VAR model can be equivalently rewritten as a lag-1 model, such that $\tilde{\mathbf{x}}_t = \tilde{\mathbf{A}}_* \tilde{\mathbf{x}}_{t-1} + \tilde{\boldsymbol{\eta}}_{t-1}$, $\tilde{\mathbf{x}}_t = (\mathbf{x}_t^\top, \dots, \mathbf{x}_{t-d+1}^\top)^\top \in \mathbb{R}^{pd}$, $\tilde{\boldsymbol{\eta}}_t = (\boldsymbol{\eta}_t^\top, \mathbf{0}_p^\top, \dots, \mathbf{0}_p^\top)^\top \in \mathbb{R}^{pd}$, and

$$\tilde{\mathbf{A}}_* = \begin{pmatrix} \mathbf{A}_{1,*} & \mathbf{A}_{2,*} & \dots & \mathbf{A}_{d,*} \\ \mathbf{I}_p & \mathbf{0}_{p \times p} & \dots & \mathbf{0}_{p \times p} \\ \mathbf{0}_{p \times p} & \mathbf{I}_p & \dots & \mathbf{0}_{p \times p} \\ \mathbf{0}_{p \times p} & \mathbf{0}_{p \times p} & \dots & \mathbf{I}_p \end{pmatrix}. \quad \text{We can then apply our test to the first block row of } \tilde{\mathbf{A}}_*, \text{ which in}$$

turn tests $\mathbf{A}_{1,*}, \dots, \mathbf{A}_{d,*}$.

We have assumed a homoscedastic and independent error structure for both error terms $\boldsymbol{\epsilon}_t$ and $\boldsymbol{\eta}_t$. This is essentially a tradeoff. Under such an error structure, the individual variables in \mathbf{x}_t are still non-identically distributed and highly correlated given the autoregressive structure of the model. In applications such as brain connectivity analysis, it is often reasonable to keep a simplified error structure (Zhang et al., 2015). In the VAR literature, more general error structures have been considered. However, when estimating the transition matrix, none of existing methods directly estimated this error structure. By contrast, our inference hinges on a good estimate of the error terms. A more general form of the error structure would introduce more unknown parameters, and requires a considerable amount of extra work to characterize the estimation precision. We thus keep a simple error structure in this first work on statistical inference, and leave the more general form of the error terms as future research.

In brain connectivity analysis, the early experiments usually focus on a single experiment subject or a very small number of subjects (Friston, 2011). More recently, data involving a large number of subjects are emerging. It is of interest to extend our modeling framework of

a single subject to multiple subjects. The key is to capture the subject-to-subject variability by incorporating the subject-specific covariates, meanwhile integrating common information shared across different subjects. A full pursuit of this topic is beyond the scope of this article, and we leave it as future research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

References

- Balakrishnan S, Wainwright MJ and Yu B (2017). Statistical guarantees for the em algorithm: From population to sample-based analysis. *The Annals of Statistics* 45 77–120.
- Basu S and Michailidis G (2015). Regularized estimation in sparse high-dimensional time series models. *The Annals of Statistics* 43 1535–1567.
- Bullmore E and Sporns O (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews. Neuroscience* 10 186–198. [PubMed: 19190637]
- Cai TT (2017). Global testing and large-scale multiple testing for high-dimensional covariance structures. *Annual Review of Statistics and Its Application* 4 423–446.
- Cai TT and Jiang T (2011). Limiting laws of coherence of random matrices with applications to testing covariance structure and construction of compressed sensing matrices. *The Annals of Statistics* 39 1496–1525.
- Cai TT, Liu W and Xia Y (2013). Two-sample covariance matrix testing and support recovery in high-dimensional and sparse settings. *J. Amer. Statist. Assoc.* 108 265–277.
- Cai TT, Ma J and Zhang L (2019). Chime: Clustering of high-dimensional gaussian mixtures with em algorithm and its optimality. *The Annals of Statistics* 47 1234–1267.
- Cai TT and Sun W (2017). Large-scale global and simultaneous inference: Estimation and testing in very high dimensions. *Annual Review of Economics* 9 411–439.
- Candes E and Tao T (2007). The dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics* 35 2313–2351.
- Cao X, Sandstede B and Luo X (2019). A functional data method for causal dynamic network modeling of task-related fmri. *Frontiers in Neuroscience* 13 127. [PubMed: 30872989]
- Chen G, Glen D, Saad Z, Hamilton JP, Thomason M, Gotlib I and Cox R (2011). Vector autoregression, structural equation modeling, and their synthesis in neuroimaging data analysis. *Computers in biology and medicine* 41 1142–55. [PubMed: 21975109]
- Chen SX, Zhang L-X and Zhong P-S (2010). Tests for high-dimensional covariance matrices. *Journal of the American Statistical Association* 105 810–819.
- Friston KJ (2011). Functional and effective connectivity: A review. *Brain Connectivity* 1 13–36. [PubMed: 22432952]
- Garg R, Cecchi G and Rao R (2011). Full-brain auto-regressive modeling (farm) using fmri. *NeuroImage* 58 416–41. [PubMed: 21439388]
- Ghahramani Z and Hinton GE (1996). Parameter estimation for linear dynamical systems. Tech. rep, University of Toronto.
- Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR et al. (2013). The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80 105–124. [PubMed: 23668970]
- Han F, Lu H and Liu H (2015). A direct estimation of high dimensional stationary vector autoregressions. *The Journal of Machine Learning Research* 16 3115–3150.
- Hsu N-J, Hung H-L and Chang Y-M (2008). Subset selection for vector autoregressive processes using lasso. *Computational Statistics & Data Analysis* 52 3645–3657.
- Krampe J, Kreiss J and Paparoditis E (2018). Bootstrap based inference for sparse high-dimensional time series models. arXiv preprint arXiv:1806.11083.

- Liu W (2013). Gaussian graphical model estimation with false discovery rate control. *The Annals of Statistics* 41 2948–2978.
- Liu W and Shao Q-M (2014). Phase transition and regularized bootstrap in large-scale t-tests with false discovery rate control. *The Annals of Statistics* 42 2003–2025.
- Luo Y, Qin S, Fernandez G, Zhang Y, Klumpers F and Li H (2014). Emotion perception and executive control interact in the salience network during emotionally charged working memory processing. *Human Brain Mapping* 35 5606–5616. [PubMed: 25044711]
- Negahban S and Wainwright MJ (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *The Annals of Statistics* 39 1069–1097.
- Ning Y and Liu H (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *The Annals of Statistics* 45 158–195.
- Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL et al. (2011). Functional network organization of the human brain. *Neuron* 72 665–678. [PubMed: 22099467]
- Reinsel GC (2003). *Elements of multivariate time series analysis*. Springer Science & Business Media.
- Sadaghiani S and D’Esposito M (2015). Functional characterization of the cingulo-opercular network in the maintenance of tonic alertness. *Cerebral Cortex* 25 2763–2773. [PubMed: 24770711]
- Schultz DH and Cole MW (2016). Higher intelligence is associated with less task-related brain network reconfiguration. *Journal of neuroscience* 36 8551–8561. [PubMed: 27535904]
- Shao X (2015). Self-normalization for time series: a review of recent developments. *Journal of the American Statistical Association* 110 1797–1817.
- Smith SD, Fox PT, Miller K, Glahn D, Fox P, Mackay CE, Filippini N, Watkins KE, Toro R, Laird A and Beckmann CF (2009). Correspondence of the brain; functional architecture during activation and rest. *Proceedings of the National Academy of Sciences of the United States of America* 106 13040–5. [PubMed: 19620724]
- Song S and Bickel PJ (2011). Large vector auto regressions. arXiv preprint arXiv:1106.3915.
- Tsay RS and Chen R (2018). *Nonlinear time series analysis*, vol. 891. John Wiley & Sons.
- Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Ugurbil K, Consortium W-MH et al. (2013). The wu-minn human connectome project: an overview. *Neuroimage* 80 62–79. [PubMed: 23684880]
- Van Praag CDG, Garfinkel SN, Sparasci O, Mees A, Philippides AO, Ware M, Ottaviani C and Critchley HD (2017). Mind-wandering and alterations to default mode network connectivity when listening to naturalistic versus artificial sounds. *Scientific Reports* 7 45273. [PubMed: 28345604]
- Wang Z, Gu Q, Ning Y and Liu H (2015). High dimensional em algorithm: Statistical optimization and asymptotic normality. In *Advances in neural information processing systems*.
- Xia M, Wang J and He Y (2013). Brainnet viewer: A network visualization tool for human brain connectomics. *PLOS ONE* 8 1–15.
- Xia Y, Cai T and Cai TT (2018). Multiple testing of submatrices of a precision matrix with applications to identification of between pathway interactions. *Journal of the American Statistical Association* 113 328–339. [PubMed: 29881130]
- Xiao H and Wu WB (2013). Asymptotic theory for maximum deviations of sample covariance matrix estimates. *Stochastic Processes and their Applications* 123 2899–2920.
- Yi X and Caramanis C (2015). Regularized em algorithms: A unified framework and statistical guarantees. In *Advances in Neural Information Processing Systems*.
- Zhang T, Wu J, Li F, Caffo B and Boatman-Reich D (2015). A dynamic directional model for effective brain connectivity using electrocorticographic (ECoG) time series. *Journal of the American Statistical Association* 110 93–106. [PubMed: 25983358]
- Zheng L and Raskutti G (2019). Testing for high-dimensional network parameters in auto-regressive models. *Electronic Journal of Statistics* 13 4977–5043.




Figure 1:
Structures of the transition matrix A . Black dots represent the nonzero entries.

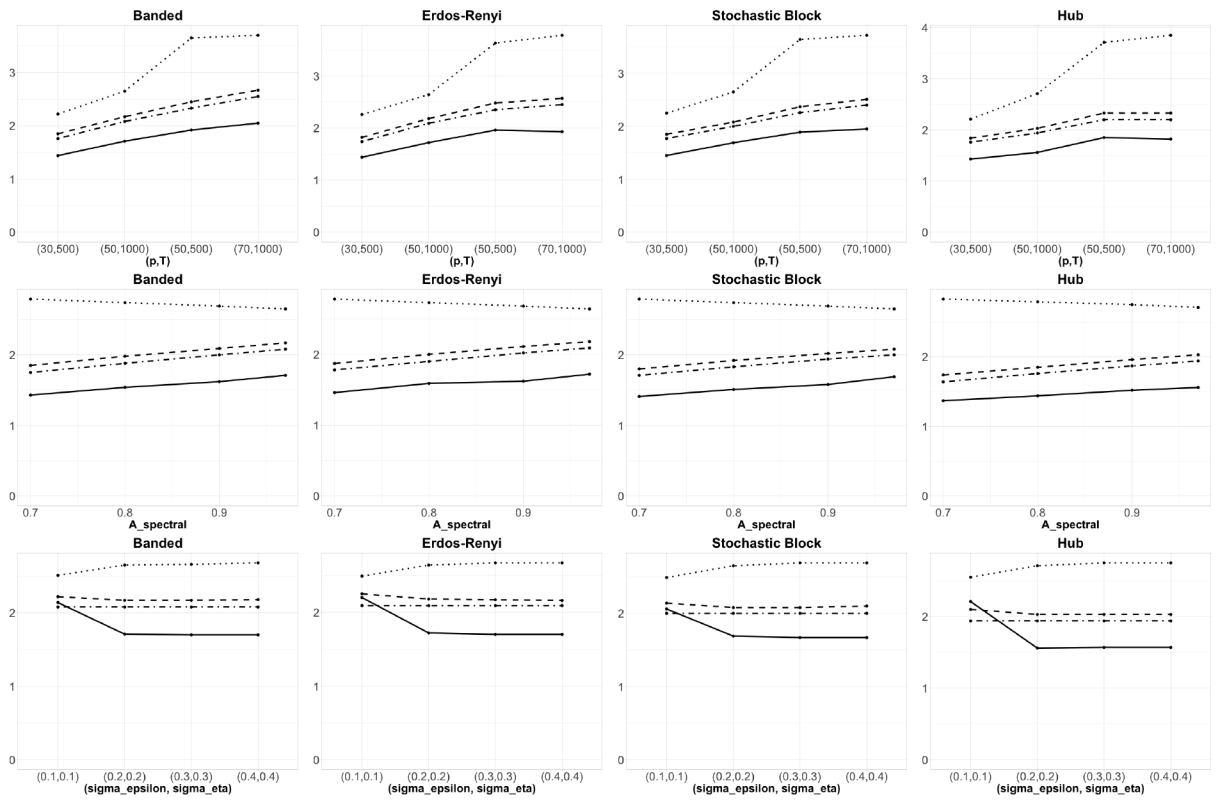


Figure 2: Estimation error of the transition matrix A_* for four network structures, and the varying (p, T) (top row), the varying signal strength $\|A_*\|_2$ (middle row), and the varying noise level $(\sigma_{\epsilon,*}, \sigma_{\eta,*})$ (bottom row). Four methods are compared: the proposed sparse EM (solid line), the standard EM (dotted line), the Lasso estimator (dot-dashed line), and the Dantzig estimator (dashed line).

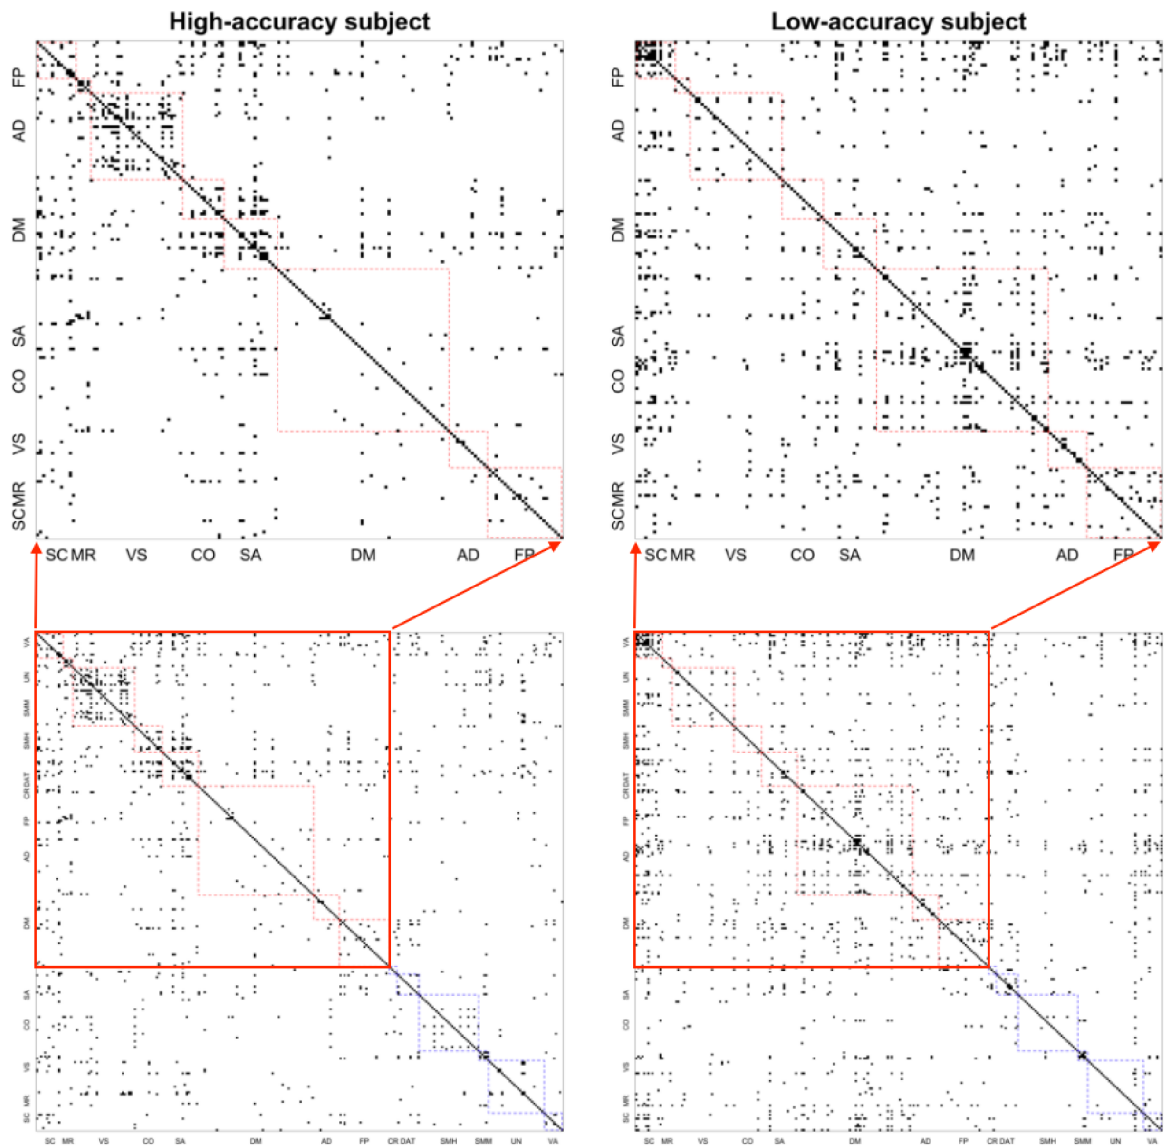


Figure 3: The identified brain connectivity patterns for the high-accuracy subject (left column) and low-accuracy subject (right column). The 14 functional modules are indicated by the blocks (bottom row), and the 8 modules that demonstrate the most within-module connections are highlighted and amplified (top row).

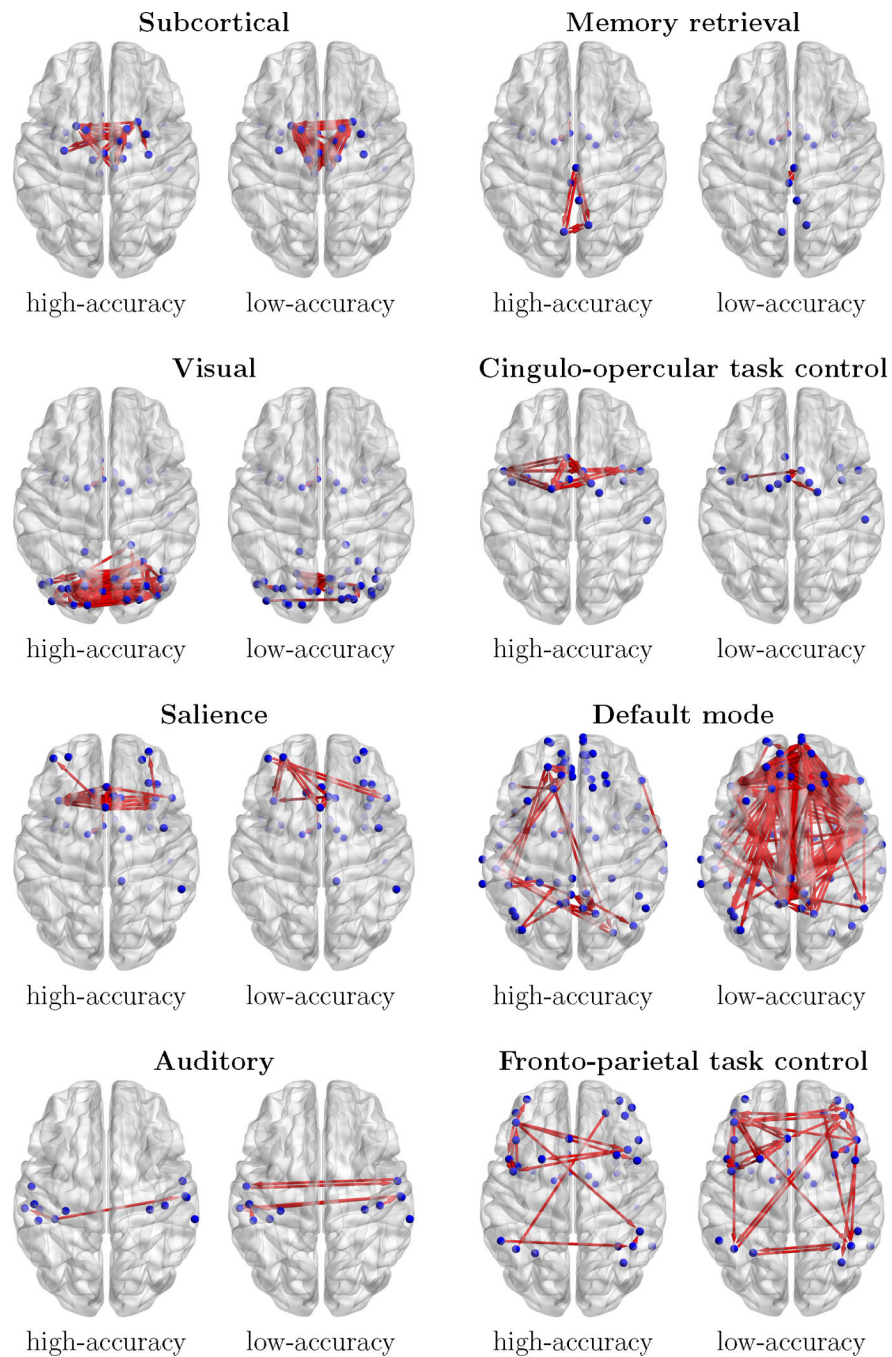


Figure 4: Visualization of the identified brain regions and their within-module connections of the high-accuracy and low-accuracy subjects for the eight functional modules.

Table 1:

Empirical size and power, in percentage, of the global test for four network structures, and the varying (p, T) (left column), the varying signal strength $\|A_*\|_2$ (middle column), and the varying noise level $(\sigma_{e,*}, \sigma_{\eta,*})$ (right column). The standard errors are reported in the parentheses.

	(p, T)	Size	Power	$\ A_*\ _2$	Size	Power	$(\sigma_{e,*}, \sigma_{\eta,*})$	Size	Power
banded	(30,500)	3.1 (0.17)	100 (0)	0.7	3.4 (0.18)	100 (0)	(0.1,0.1)	4.5 (0.21)	100 (0)
	(50,500)	2.7 (0.16)	100 (0)	0.8	3.1 (0.17)	100 (0)	(0.2,0.2)	2.4 (0.15)	100 (0)
	(50,1000)	2.4 (0.15)	100 (0)	0.9	2.9 (0.17)	100 (0)	(0.3,0.3)	2.4 (0.15)	100 (0)
	(70,1000)	2.7 (0.16)	100 (0)	0.97	2.4 (0.15)	100 (0)	(0.4,0.4)	2.4 (0.15)	100 (0)
Erdős-Rényi	(30,500)	3.1 (0.17)	100 (0)	0.7 (0.18)	3.4 (0)	100 (0)	(0.1,0.1)	5.0 (0.22)	100 (0)
	(50,500)	2.6 (0.16)	100 (0)	0.8	3.5 (0.18)	100 (0)	(0.2,0.2)	2.8 (0.17)	100 (0)
	(50,1000)	2.8 (0.17)	100 (0)	0.9	3.1 (0.17)	100 (0)	(0.3,0.3)	2.7 (0.16)	100 (0)
	(70,1000)	3 (0.17)	100 (0)	0.97	2.8 (0.17)	100 (0)	(0.4,0.4)	2.7 (0.16)	100 (0)
stochastic block	(30,500)	3.2 (0.18)	100 (0)	0.7	3.8 (0.19)	100 (0)	(0.1,0.1)	5.9 (0.24)	100 (0)
	(50,500)	3.1 (0.17)	100 (0)	0.8	3.7 (0.19)	100 (0)	(0.2,0.2)	3.4 (0.18)	100 (0)
	(50,1000)	3.4 (0.18)	100 (0)	0.9	3.5 (0.18)	100 (0)	(0.3,0.3)	3.3 (0.18)	100 (0)
	(70,1000)	2.1 (0.14)	100 (0)	0.97	3.4 (0.18)	100 (0)	(0.4,0.4)	3.3 (0.18)	100 (0)
hub	(30,500)	3.5 (0.18)	100 (0)	0.7	2.9 (0.17)	100 (0)	(0.1,0.1)	6.3 (0.24)	100 (0)
	(50,500)	2 (0.14)	100 (0)	0.8	2.9 (0.17)	100 (0)	(0.2,0.2)	2.6 (0.16)	100 (0)
	(50,1000)	2.6 (0.16)	100 (0)	0.9	2.5 (0.16)	100 (0)	(0.3,0.3)	2.6 (0.16)	100 (0)
	(70,1000)	3.7 (0.19)	100 (0)	0.97	2.6 (0.16)	100 (0)	(0.4,0.4)	2.5 (0.16)	100 (0)

Table 2:

Average false discovery proportion (FDP) and true positive rate (TPR), in percentage, of the simultaneous test for four network structures, and the varying (p, T) (left column), the varying signal strength $\|A_*\|_2$ (middle column), and the varying noise level $(\sigma_{e,*}, \sigma_{\eta,*})$ (right column). The standard errors are reported in the parentheses.

	(p, T)	FDR	TPR	$\ A_*\ _2$	FDR	TPR	$(\sigma_{e,*}, \sigma_{\eta,*})$	FDR	TPR
banded	(30,500)	4.34 (0.03)	73.65 (0.05)	0.7	4.59 (0.02)	71.42 (0.04)	(0.1,0.1)	5.39 (0.02)	93.59 (0.02)
	(50,500)	3.91 (0.02)	67.19 (0.05)	0.8	4.45 (0.02)	82.55 (0.03)	(0.2,0.2)	3.91 (0.02)	92.3 (0.02)
	(50,1000)	3.91 (0.02)	92.3 (0.02)	0.9	4.18 (0.02)	89.24 (0.03)	(0.3,0.3)	3.91 (0.02)	92.27 (0.02)
	(70,1000)	3.73 (0.02)	88.44 (0.02)	0.97	3.91 (0.02)	92.3 (0.02)	(0.4,0.4)	3.91 (0.02)	92.27 (0.02)
Erdős-Rényi	(30,500)	3.95 (0.03)	75.45 (0.06)	0.7	4.63 (0.02)	70.53 (0.05)	(0.1,0.1)	5.61 (0.02)	98.27 (0.01)
	(50,500)	3.93 (0.02)	65.68 (0.05)	0.8	4.59 (0.02)	86.83 (0.03)	(0.2,0.2)	3.98 (0.02)	97.4 (0.02)
	(50,1000)	3.98 (0.02)	97.4 (0.02)	0.9	4.21 (0.02)	94.57 (0.02)	(0.3,0.3)	3.97 (0.02)	97.35 (0.02)
	(70,1000)	4.07 (0.02)	91.36 (0.02)	0.97	3.98 (0.02)	97.4 (0.02)	(0.4,0.4)	3.97 (0.02)	97.34 (0.02)
stochastic block	(30,500)	4.14 (0.03)	73.96 (0.05)	0.7	4.7 (0.03)	66.79 (0.05)	(0.1,0.1)	5.62 (0.02)	90.98 (0.02)
	(50,500)	3.75 (0.02)	61.12 (0.05)	0.8	4.65 (0.02)	78.86 (0.04)	(0.2,0.2)	4.18 (0.02)	89.6 (0.03)
	(50,1000)	4.18 (0.02)	89.6 (0.03)	0.9	4.42 (0.02)	86.15 (0.03)	(0.3,0.3)	4.17 (0.02)	89.55 (0.03)
	(70,1000)	3.97 (0.02)	84.63 (0.02)	0.97	4.18 (0.02)	89.6 (0.03)	(0.4,0.4)	4.18 (0.02)	89.54 (0.03)
hub	(30,500)	4.33 (0.03)	81.02 (0.05)	0.7	4.75 (0.03)	65.07 (0.05)	(0.1,0.1)	6.37 (0.02)	96.96 (0.02)
	(50,500)	3.81 (0.02)	58.74 (0.06)	0.8	4.7 (0.02)	82.08 (0.04)	(0.2,0.2)	4.25 (0.02)	95.28 (0.02)
	(50,1000)	4.25 (0.02)	95.28 (0.02)	0.9	4.45 (0.02)	91.76 (0.03)	(0.3,0.3)	4.26 (0.02)	95.28 (0.02)
	(70,1000)	4.39 (0.02)	77.21 (0.03)	0.97	4.25 (0.02)	95.28 (0.02)	(0.4,0.4)	4.19 (0.02)	95.26 (0.02)