Mixed Membership Estimation in Bayesian Network Autoregression*

Endong Wang† Siao Xu‡

Abstract

We propose a generative network autoregressive model in which the network coefficient is a random matrix following a joint distribution specified by a Bayesian belief network. Rather than treating the coefficient as fixed and observed, we model it as the output of a latent generative process that simultaneously governs the network structure and the resulting panel data. A central feature is that units may exhibit mixed membership across latent groups rather than belonging to a single group. We develop an estimator that recovers the latent membership structure from panel observations, establish identification and consistency as both panel dimensions grow, and assess finite-sample performance through simulations. An application to the U.S. financial system documents substantial mixed exposure across latent groups, showing that mixed membership materially shapes the dependence structure among financial firms. *Keywords:* Network autoregression, mixed membership, latent group, Bayesian belief network, cross-sectional dependence.

JEL Classification: C14, C32, C38, G10, G17.

*We are very grateful to Carsten Trenkler, Jean-Marie Dufour, Matthew O. Jackson, Mengshan Xu, Evan Munro, Christopher Rothe, Frank Schorfheide, Weining Wang, Amilcar Velez, Bin Peng, for their very insightful suggestions and discussions. Seminar participants at Universität Mannheim have provided invaluable feedback. All faults are our own. This work was supported by the research funds provided by the University of Mannheim.

†Department of Economics, University of Mannheim, L7, 3–5, Room 124, Mannheim, Germany, 68161. TEL: +49 621-181-1879; Email: endong.wang@uni-mannheim.de.

‡Department of Economics, University of Mannheim, L7, 3–5, Room 104, Mannheim, Germany, 68161. Email: siao.xu@students.uni-mannheim.de.

1 Introduction

Learning latent group structure from panel data is a central problem in econometrics. A large literature studies this problem through grouped panel models, where heterogeneity is represented by a finite number of latent groups and each unit belongs to one group only, see Bonhomme and Manresa (2015); Su, Shi and Phillips (2016); Chetverikov and Manresa (2022); Mugnier (2025). Related ideas have entered the analysis of cross-sectional dependence in dynamic systems since the development of network autoregression (Zhu et al., 2017). In that literature, the network coefficient governing cross-sectional dynamic dependence is often assumed to exhibit latent group structure, and the econometric task is to recover group memberships from the panel induced by that matrix, see Zhu and Pan (2020); Chen, Fan and Zhu (2023); Guðmundsson and Brownlees (2021); Guðmundsson (2026). Whether in the group panel model or in the network autoregression with group structure mentioned above, a common maintained restriction is that the latent structure takes the form of pure membership. That is, one unit is 100% affiliated with one specific group and 0% affiliated with the others. For example, a U.S. voter exclusively belongs to either the Democrat or Republican. A firm exclusively belongs to one specific industry.

That restriction can be too strong in applications. For example, U.S. voters interchangeably vote for the Republican and Democrat. A fintech company may belong to both the financial and the high-tech industries. People may have mixed ancestry from different nations. The mixed membership structure exists everywhere. Furthermore, in many economic settings, units are linked to several latent groups at once, and their dynamic dependence reflects that mixed membership structure. Financial institutions provide a natural example. Large firms often combine commercial banking, securities intermediation, asset management, insurance, payments, and real-estate-related activities within the same organization. More generally, economic units are often shaped by several latent dimensions of interaction at the same time. In such settings, a pure-membership representation may be systematically misleading. The issue is not simply that units differ. The issue is that the heterogeneity is organized through mixed affiliation rather than through pure affiliation.

This paper develops a generative mixed-membership network autoregression for panel data. The observed panel enters through a large-dimensional network autoregression, while the coefficient matrix, assumed to be a network governing cross-sectional dynamic dependence, is treated as a latent random object. Rather than taking that matrix as fixed and directly observed, we model it as generated by an underlying probabilistic mechanism encoded through a Bayesian belief network. This yields a joint description of latent memberships, coefficient formation, and panel dynamics. It also changes the econometric problem in a basic way. The matrix carrying the mixed-membership structure is not observed and must be recovered from the panel data. The matrix denotes the group structure of interest in the present model.

Our contribution is threefold. At the modeling level, we introduce mixed membership and generative structure into a dynamic environment in which the dependence matrix is latent. Existing grouped specifications in panel data and related large network autoregressive formulations are organized around hard assignment. By contrast, our framework allows each unit to load on several latent groups simultaneously. This extension is economically important in settings where interaction patterns reflect overlapping affiliations rather than a sharp partition.

At the econometric level, we develop a feasible estimator for this setting by combining ordinary

least squares with an adaptation of SCORE (Jin, 2015; Jin, Ke and Luo, 2024). The procedure first estimates the network coefficient of a large network autoregression from panel data, then uses its spectral structure to recover mixed memberships. The theoretical analysis establishes identification of the latent membership structure and derives consistency and error bounds for the estimator. The argument links the leading eigenspace of the population coefficient matrix to the underlying mixed memberships and then controls the estimation error.

At the empirical level, the paper shows why the mixed-membership extension matters substantively. We apply the method to daily volatility data for U.S. financial firms and ask whether the cross section is better described by overlapping latent dependence groups than by an exclusive partition. This is a natural environment for the question, since large financial institutions rarely operate along a single business line and the dependence structure in firm-level volatility need not coincide with conventional sectoral boundaries. The estimates point to substantial overlap. In the application, many institutions lie well inside the simplex, the average largest membership weight is only 0.632, and an exclusive classification would mask an important part of the estimated dependence structure. The empirical finding is therefore not merely that firms are heterogeneous. It is that the heterogeneity takes the form of overlapping latent group exposure.

The paper is related to several strands of the literature. First, it contributes to the literature on latent group structure in panel data, including Brownlees, Guðmundsson and Lugosi (2022); Su, Shi and Phillips (2016); Bonhomme and Manresa (2015); Zhang, Wang and Zhu (2019). Relative to that work, our framework replaces pure membership with mixed membership in a dynamic setting with cross-sectional dependence. Second, it relates to the literature on mixed-membership models and their estimation in networks, including Airoldi et al. (2008); Mao, Sarkar and Chakrabarti (2017, 2021); Jin, Ke and Luo (2024); Zhang, Levina and Zhu (2020). The key difference is that those papers recover mixed memberships from an observed network or adjacency object, whereas in our setting the relevant dependence matrix is latent and must be inferred from panel dynamics. Third, it connects to the literature on network autoregressions and large VARs, including Zhu et al. (2017, 2019a,b); Guðmundsson and Brownlees (2021). Our contribution there is to allow the latent group structure of the dependence matrix to be mixed rather than exclusive. Finally, the paper draws on probabilistic graphical models and Bayesian belief networks as a way to encode the latent generative structure, see Blei, Ng and Jordan (2003); Munro and Ng (2022); Koller and Friedman (2009).

The remainder of the paper is organized as follows. Section 2 introduces the model, identification, estimation, and asymptotic theory. Section 3 reports the finite-sample evidence. Section 4 presents the empirical application to the U.S. financial system. Section 5 concludes. Technical proofs are collected in the Appendix.

Notation. Let N denote the number of time series and T the number of observations. Of the N series, N_0 are pure-membership units and $N - N_0$ are mixed-membership units. Lowercase letters denote scalars, and uppercase letters denote matrices. For a matrix Y, $Y_{i\cdot}$, $Y_{\cdot j}$, and Y_{ij} denote its i-th row, j-th column, and (i, j) entry, respectively. For a square matrix Y, let $\lambda_i(Y)$ denote its i-th eigenvalue, ordered by magnitude, and let $\rho(Y)$ denote its spectral radius. For a symmetric matrix Y, $\lambda_{\max}(Y)$ and $\lambda_{\min}(Y)$ denote its largest and smallest eigenvalues. We write $\|Y\|$ for the spectral norm and $\|Y\|_F$ for the Frobenius norm. A directed weighted graph is denoted by $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$, where $\mathcal{V} = \{1, \ldots, N\}$, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$, and $\mathcal{W}$ assigns edge weights. Its adjacency matrix is

$A \in \mathbb{R}^{N\times N}$, with $A_{ij} \neq 0$ if there is an edge from node j to node i, and $A_{ij} = 0$ otherwise.

2 Model, identification, and estimation

2.1 The mixed membership network autoregressive model

We study a mixed membership network autoregressive model in which a latent weighted network enters the law of motion of a large-dimensional vector autoregression. The model has a hierarchical structure. Mixed-membership parameters determine link probabilities, link realizations determine which coefficients are active, and the resulting coefficient matrix governs dynamic dependence across units.

Let K denote the number of latent groups. For each unit $i = 1, \dots, N$, let $\pi_i \in \mathbb{R}^K$ denote an unknown membership vector, with nonnegative entries summing to one, and let $\Pi \in \mathbb{R}^{N\times K}$ collect the row vectors $\pi_1', \dots, \pi_N'$. Let $A \in \{0,1\}^{N\times N}$ denote the adjacency matrix, $B \in \mathbb{R}^{K\times K}$ the group interaction matrix, $\Theta = \operatorname{diag}(\theta_1, \dots, \theta_N)$ the diagonal matrix of activity parameters, $\Phi \in \mathbb{R}^{N\times N}$ the latent coefficient matrix, and $\mu \in \mathbb{R}^N$ and $\sigma > 0$ the parameters in the conditional Gaussian law of the observed process.

The model is defined as follows.

Definition 2.1 (Mixed membership network autoregression)**.** *Given parameters* Π*,* B*,* Θ*,* $(\underline{u}, \overline{u})$*,* μ*, and* σ*, the process* $\{Y_t\}_{t=1}^T$ *is generated according to*

$$
\begin{aligned}
A_{ij} \mid \Pi, B, \Theta &\sim \text{Bernoulli}\big(\theta_i \pi_i' B \pi_j \theta_j\big), \qquad i, j = 1, \dots, N, \\
\Phi_{ij} \mid A_{ij}, (\underline{u}, \overline{u}) &\sim \begin{cases} \text{Uniform}(\underline{u}, \overline{u}), & \text{if } A_{ij} = 1, \\ 0, & \text{if } A_{ij} = 0, \end{cases} \qquad i, j = 1, \dots, N, \\
Y_t \mid Y_{t-1}, \mu, \sigma, \Phi &\sim \mathcal{N}\big(\mu + \Phi Y_{t-1}, \sigma^2 I_N\big), \qquad t = 1, \dots, T.
\end{aligned}
$$

The central object in Definition 2.1 is the latent network coefficient matrix Φ. Conditional on Φ, the observed process follows a first-order large-dimensional vector autoregression with intercept μ and homoskedastic innovations of variance σ^2. The entry Φ_{ij} measures the effect of the lagged outcome of unit j on the current outcome of unit i, so Φ governs the pattern of dynamic dependence across units. The model does not treat this coefficient matrix as unrestricted. Instead, it imposes a latent mixed-membership structure on the cross-sectional pattern of nonzero coefficients.

For each unit i, the vector π_i records how its latent exposure is allocated across the K groups. Mixed membership allows a unit to load on several groups simultaneously, rather than assigning each unit to one group only. The matrix B governs how group memberships interact in link formation. Its (k, ℓ) entry determines how exposure of the receiving unit to group k and exposure of the sending unit to group ℓ affect the probability of a directed link. The diagonal matrix Θ captures unit-specific variation in overall network activity, with larger θ_i corresponding to a more active role in cross-sectional dependence.

Given Π, B, and Θ, the Bernoulli variable A_{ij} determines whether there is an active directed

link from unit j to unit i. The matrix A therefore determines which entries of Φ are nonzero. Conditional on $A_{ij} = 1$, the coefficient Φ_{ij} measures the strength of that link, while $\underline{u}$ and $\overline{u}$ determine the support of active link weights. The model thus separates link formation from link strength. The objects Π, B, and Θ organize the latent cross-sectional structure, and Φ maps that structure into the dynamic law of motion of the system.

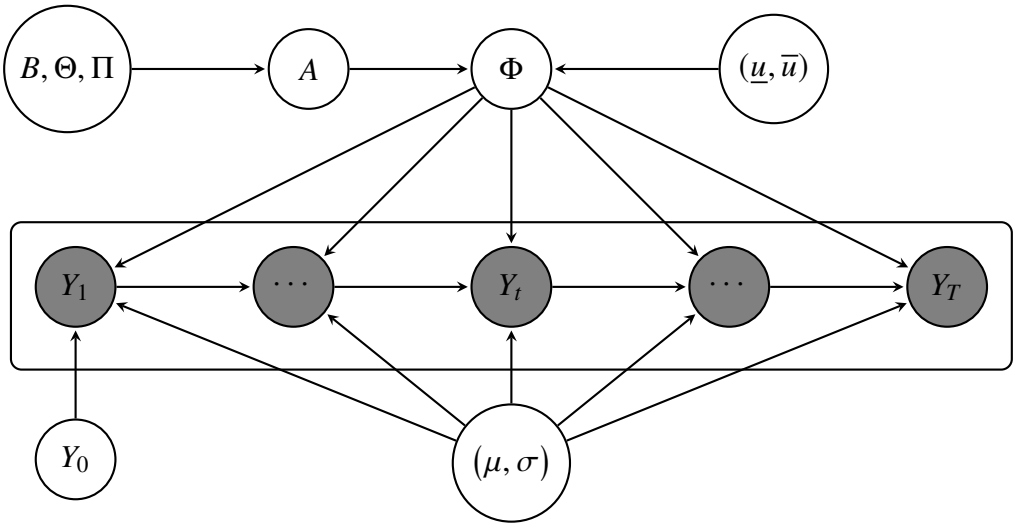

Figure 1: Graphical representation of the MMNAR model. The graph highlights the hierarchical dependence structure in Definition 2.1. The latent objects Π, B, and Θ determine the link-indicator matrix A, which in turn determines the sparsity pattern of the network coefficient matrix Φ. Conditional on Φ, μ, and σ, the observed process $\{Y_t\}_{t=1}^T$ follows a first-order large-dimensional vector autoregression. Shaded nodes are observed.

We maintain the restriction $\theta_i \pi_i' B \pi_j \theta_j \in [0, 1]$ for all i and j, so that the implied link probabilities are well defined. A sufficient condition for stability is $\rho(\Phi) < 1$, which is the standard stationarity requirement for a first-order autoregressive system. The SB-VAR(1) model of Guðmundsson and Brownlees (2021) is nested as a special case of Definition 2.1. Their specification normalizes active links by the degree matrix and scales them by a common autoregressive parameter. Our formulation instead allows the network coefficient matrix to be latent, weighted, and heterogeneous across links.

Figure 1 summarizes the hierarchical structure of the model. At the top level, the latent objects Π, B, and Θ determine the link-indicator matrix A. The matrix A in turn determines which entries of Φ are active, while $(\underline{u}, \overline{u})$ determine the support of the corresponding nonzero coefficients. Conditional on Φ, μ, and σ, the observed process $\{Y_t\}_{t=1}^T$ follows a first-order large-dimensional vector autoregression. The figure therefore makes clear the ordering of the model: latent-group structure determines network topology, network topology determines the coefficient matrix, and the coefficient matrix governs the evolution of the observed data.

Conditional on the initial value Y_0, and holding Π, μ, σ, $\underline{u}$, $\overline{u}$, Θ, and B fixed, the recursive

structure of the model implies the factorization

$$\mathbb{P}\big(\{Y_t\}_{t=1}^T, \Phi, A\big) = \left\{\prod_{t=1}^{T} \mathbb{P}\big(Y_t \mid Y_{t-1}, \Phi, \mu, \sigma\big)\right\} \mathbb{P}\big(\Phi \mid A, \underline{u}, \overline{u}\big) \times \left\{\prod_{i=1}^{N}\prod_{j=1}^{N} \mathbb{P}\big(A_{ij} \mid B, \pi_i, \pi_j, \Theta\big)\right\}. \tag{2.1}$$

This factorization makes clear that the population object relevant for estimation is the first moment of the latent coefficient matrix.

Theorem 2.1. *Under Definition 2.1,*

$$\mathbb{E}(\Phi) = \frac{\underline{u} + \overline{u}}{2} \Theta \Pi B \Pi' \Theta.$$

Theorem 2.1 provides the key link between the latent membership structure and the coefficient matrix recovered from the data. In particular, the mean coefficient matrix inherits a low-rank mixed-membership form. The matrix Π encodes the latent-group composition of each unit, B governs interactions across latent groups, and Θ captures unit-specific variation in network activity. The scalar $(\underline{u} + \overline{u})/2$ is the mean weight of an active link. Thus, $\mathbb{E}(\Phi)$ summarizes the systematic component of cross-sectional dependence implied by the latent structure, while the realized matrix Φ also reflects variation from random link formation and random link weights.

This representation is central for identification. Once the first moment of Φ is characterized, the problem reduces to recovering the latent mixed-membership structure encoded in $\Theta\Pi B\Pi'\Theta$. The Appendix proves this result by integrating out the latent adjacency matrix and then taking expectations over the remaining sources of randomness.

2.2 Identification

Our objective is to recover the mixed membership matrix Π. By Theorem 2.1, this problem reduces to an inverse problem defined by the population moment $\mathbb{E}(\Phi) = \frac{\underline{u}+\overline{u}}{2} \Theta\Pi B\Pi'\Theta$. Throughout this subsection, the objects Π, B, Θ, and $(\underline{u}, \overline{u})$ are treated as fixed unknown parameters.

The following conditions are sufficient for identification.

Assumption 2.1. *Under Definition 2.1,*

1. *B is full rank and positive definite, with $B_{kk} > 0$ for $k = 1, \ldots, K$.*
2. *For each community $k = 1, \ldots, K$, there exists at least one pure node, namely an i such that $\pi_{ik} = 1$.*
3. *The degree parameters satisfy $\theta_i > 0$ for all i, and $N^{-1} \sum_{i=1}^{N} \theta_i = 1$.*
4. *$(\underline{u}, \overline{u})$ is fixed.*

These conditions exclude observationally equivalent representations of the same population moment. Full rank and positive definiteness of B rule out redundant community interactions and ensure that the community structure enters the mean matrix in a nondegenerate way. The pure-node condition anchors the simplex generated by the membership vectors and thereby fixes the geometric support of mixed membership. The normalization on Θ removes the remaining scale indeterminacy between node activity and the other latent components. Taken together, these restrictions ensure that variation in $\mathbb{E}(\Phi)$ can be attributed uniquely to membership composition, community interaction, and node heterogeneity, apart from a relabeling of the communities.

The next theorem states the identification result.

Theorem 2.2. *Suppose that a MMNAR satisfying Definition 2.1 also satisfies Assumption 2.1. Then any parameter tuple* $\left((\underline{u}, \overline{u}), \Theta, \Pi, B\right)$ *that generates the population matrix* $\mathbb{E}(\Phi)$ *is unique up to a permutation of the community labels.*

Theorem 2.2 formalizes the identifying content of the model. In the population mean matrix, a large entry may reflect strong overlap in membership profiles, strong interaction across communities through B, high node activity through Θ, or some combination of these channels. Assumption 2.1 separates these components. As a result, the mixed membership structure is pinned down by the way the low-rank matrix $\Theta\Pi B\Pi'\Theta$ organizes cross-sectional dependence. Related identification arguments for overlapping and mixed-membership models appear in Latouche, Birmelé and Ambroise (2011) and Zhang, Levina and Zhu (2020).

The next lemma links the latent membership structure to the leading eigenspace of $\mathbb{E}(\Phi)$. This provides the population basis for the estimator developed below.

Lemma 2.3. *Let the MMNAR in Definition 2.1 satisfy Assumption 2.1. Let* $\lambda_1, \ldots, \lambda_K$ *denote the* K *largest eigenvalues of* $\mathbb{E}(\Phi)$*, ordered from largest to smallest, and let* $x_1, \ldots, x_K$ *denote the associated eigenvectors. Define* $X = [x_1, \ldots, x_K]$*. Then there exists a unique nonsingular matrix* $P \in \mathbb{R}^{K\times K}$ *such that*

$$X = \left(\sqrt{\frac{\underline{u} + \overline{u}}{2}}\,\Theta\Pi\right)P.$$

Lemma 2.3 shows that the leading eigenspace of $\mathbb{E}(\Phi)$ spans the same column space as $\Theta\Pi$, up to a nonsingular linear transformation. This is the key population implication underlying the spectral estimator. Once the leading eigenspace is recovered from a sample analogue of the coefficient matrix, estimation of Π reduces to recovering the simplex geometry induced by mixed membership, in the spirit of SCORE and Mixed-SCORE, see Jin (2015); Jin, Ke and Luo (2024).

2.3 Estimation

We estimate Π in two steps. The first step recovers the latent coefficient matrix from the panel by ordinary least squares. The second step treats a symmetrized version of that estimate as a weighted network and applies a Mixed-SCORE decomposition to recover the membership profiles. Throughout, the number of communities K is taken as known.

The first step recovers the coefficient matrix of the large VAR from the panel. To remove the intercept, the data are first demeaned over time. Let

$$\mathbf{Y} = [Y_2, \ldots, Y_T], \qquad \mathbf{X} = [Y_1, \ldots, Y_{T-1}].$$

Provided $\mathbf{X}\mathbf{X}'$ is invertible, the OLS estimator is

$$\hat{\Phi} = \mathbf{Y}\mathbf{X}'(\mathbf{X}\mathbf{X}')^{-1}.$$

This estimator serves as a sample analogue of the latent network coefficient matrix Φ.

The symmetrization in Step 2 is a finite-sample regularization device. It reduces the effect of directional estimation noise and stabilizes the subsequent eigen-analysis. In this sense, $\hat{\Phi}^S$ serves as a weighted network representation of the estimated structure, and its leading spectral components provide the empirical counterpart to the low-rank object characterized in the identification analysis. This step follows Guðmundsson and Brownlees (2021).

The remaining steps adapt Mixed-SCORE to the estimated coefficient matrix. By Lemma 2.3, the leading eigenspace of the population matrix spans the same column space as $\Theta\Pi$, up to a nonsingular transformation. This implication is central for estimation, since it converts recovery of the membership matrix into a geometric problem in a low-dimensional spectral space.

The ratio transformation in Step 4 removes the nuisance scale induced by node heterogeneity and maps the rows of the eigenvector matrix into a simplex structure. Under the mixed-membership model, pure-community units are located at the vertices of this simplex, while mixed-membership units lie in its interior as convex combinations of the vertices. The corner-search step therefore recovers the empirical analogues of the pure-community directions. The barycentric projection in Step 6 then expresses each unit as a convex combination of those directions. Step 7 rescales these weights to undo the eigenvalue normalization embedded in the spectral representation, and Step 8 imposes the unit-sum restriction required of membership vectors.

The estimator differs from standard applications of SCORE and Mixed-SCORE in one important respect. In those settings, the relevant network matrix is directly observed. Here, by contrast, the network object of interest is latent and must first be recovered from panel observations through the large VAR. The algorithm therefore combines time-series estimation with spectral recovery of the mixed-membership structure.

2.4 Large-sample theory

We now establish consistency of the MMNAR-Mixed-SCORE estimator. The argument has two parts. We first characterize the geometry of the population matrix $\mathbb{E}(\Phi)$. We then control the discrepancy between the estimated coefficient matrix and its population counterpart. This separation between population geometry and sample perturbation is standard in spectral analysis, see Rohe, Chatterjee and Yu (2011).

Let $c = (\underline{u} + \overline{u})/2$. By Theorem 2.1,

$$\mathbb{E}(\Phi) = c\,\Theta\Pi B \Pi'\Theta. \tag{2.2}$$

Algorithm 1 MMNAR-Mixed-SCORE

Input: panel data $\{Y_t\}_{t=1}^T$, number of communities K.

1: Estimate the coefficient matrix $\hat{\Phi}$ by OLS.

2: Symmetrize the estimate, $\hat{\Phi}^S = (\hat{\Phi} + \hat{\Phi}')/2$.

3: Let $\hat{X} \in \mathbb{R}^{N\times K}$ collect the eigenvectors associated with the K largest eigenvalues of $\hat{\Phi}^S$.

4: Set the threshold $S = \log N$, and construct $\hat{R} \in \mathbb{R}^{N\times(K-1)}$ by

$$\hat{R}_{ik} = \text{sign}\left(\hat{X}_{i,k+1}/\hat{X}_{i1}\right) \min\left\{\left|\hat{X}_{i,k+1}/\hat{X}_{i1}\right|, S\right\}, \qquad 1 \le i \le N, \;\; 1 \le k \le K-1.$$

5: Find K corner points $\{\hat{v}_1, \ldots, \hat{v}_K\}$ in $\hat{R}$ using a corner-search routine, for example successive projection or sketched vertex search.

6: For each $i = 1, \ldots, N$, compute barycentric weights by solving

$$\hat{w}_i = \arg\min_{w\in\mathbb{R}^K} \left\|\hat{R}_i - \sum_{k=1}^K w_k \hat{v}_k\right\|^2 \quad \text{subject to} \quad w_k \ge 0, \;\; \sum_{k=1}^K w_k = 1.$$

7: For $1 \le i \le N$ and $1 \le k \le K$, define

$$\hat{\pi}^*_{ik} = \max\left\{0, \hat{w}_{ik}/\hat{b}_{1k}\right\}, \qquad \hat{b}_{1k} = \left[\hat{\lambda}_1 + \hat{v}_k' \text{diag}(\hat{\lambda}_2, \ldots, \hat{\lambda}_K)\hat{v}_k\right]^{-1/2}.$$

8: Normalize,

$$\hat{\pi}_i = \hat{\pi}^*_i / \|\hat{\pi}^*_i\|_1, \qquad i = 1, \ldots, N.$$

Output: $(N \times K)$ matrix of estimated mixed memberships $\hat{\Pi}$.

Defining $B^* = cB$, the population matrix takes the form $\Theta\Pi B^*\Pi'\Theta$, which coincides with the algebraic structure studied in degree-corrected mixed-membership models, see Fan et al. (2022); Jin, Ke and Luo (2024). The geometric arguments developed there therefore apply here as well. In substantive terms, $\mathbb{E}(\Phi)$ summarizes the systematic component of dynamic dependence across units. The matrix Π collects the latent exposure profiles of the units, B^* governs interaction across latent groups, and Θ scales the overall intensity with which each unit enters the dependence structure. The low-rank structure in $\mathbb{E}(\Phi)$ therefore reflects a small number of latent groups that organize dynamic dependence in the cross section.

The required regularity conditions are collected next.

Assumption 2.2. *Under Definition 2.1, define*

$$G = K\left\|\sqrt{\frac{\underline{u}+\overline{u}}{2}}\,\Theta\right\|^{-2}\left(\frac{\underline{u}+\overline{u}}{2}\,\Pi'\Theta^2\Pi\right). \tag{2.3}$$

There exists a constant $z > 0$ such that $\Theta_{\max} \le z$, $\|B\|_{\max} \le z$, $\|G\| \le z$, $\|G^{-1}\| \le z$, $\Theta_{\max} \le z\Theta_{\min}$, and $\sqrt{\frac{\log N}{\left(\sum_{i=1}^N \theta_i\right)^2/N}} \to 0$.

Assumption 2.2 controls the magnitude of degree heterogeneity and interaction strength across latent groups, and requires the effective average degree to dominate $\log N$. These restrictions ensure that the latent group structure is neither too weak to detect nor too unevenly distributed across units. Under these conditions, the rows of Π admit the same simplex representation as in the degree-corrected mixed-membership model. Pure nodes determine the vertices, mixed-membership nodes lie in the interior, and the SCORE ratio map removes the nuisance scale induced by Θ. In substantive terms, the assumption ensures that the latent group structure leaves a sufficiently strong imprint on the coefficient matrix to be recoverable from the data.

We next turn to the sample problem. The perturbation of interest is $\hat{\Phi} - \mathbb{E}(\Phi)$. It has two components, the network realization error, $\Phi - \mathbb{E}(\Phi)$, and the time-series estimation error, $\hat{\Phi} - \Phi$. The first component reflects the fact that the realized coefficient matrix is only one draw from the latent network formation mechanism. The second reflects the sampling error that arises because the coefficient matrix must be estimated from a finite time series. The large-sample argument must therefore control both cross-sectional randomness in the latent network and estimation uncertainty in the large VAR.

Write the MMNAR as

$$Y_t = \mu + \Phi Y_{t-1} + \varepsilon_t. \tag{2.4}$$

The concentration argument is carried out conditional on the realized latent graph and the corresponding coefficient matrix Φ. Let $\mathcal{G}$ denote the realized latent graph and let $\Sigma_Y = \mathbb{E}(Y_t Y_t' \mid \mathcal{G}, \Phi)$. For $r \in \mathbb{Z}$ and $m \geq 1$, let $\mathcal{B}^r_{-\infty}$ and $\mathcal{B}^\infty_{r+m}$ denote the σ-fields generated by $\{\Sigma_Y^{-1/2} Y_t : t \leq r\}$ and $\{\Sigma_Y^{-1/2} Y_t : t \geq r + m\}$, respectively. The associated α-mixing coefficient is

$$\alpha(m) = \sup_r \sup_{\mathcal{A} \in \mathcal{B}^r_{-\infty},\, \mathcal{B} \in \mathcal{B}^\infty_{r+m}} \left| \mathbb{P}(\mathcal{A} \cap \mathcal{B} \mid \mathcal{G}, \Phi) - \mathbb{P}(\mathcal{A} \mid \mathcal{G}, \Phi)\mathbb{P}(\mathcal{B} \mid \mathcal{G}, \Phi) \right|. \tag{2.5}$$

We impose the following sampling conditions.

Assumption 2.3. *Let $\{Y_t\}_{t=1}^T$ be generated by the MMNAR in Definition 2.1. Then:*

1. *$\{\varepsilon_t\}$ is covariance stationary, serially uncorrelated, and $\mathbb{E}(Y_{t-1}\varepsilon_t' \mid \mathcal{G}, \Phi) = 0$ for all t.*
2. *$\|\mathbb{E}(\varepsilon_t \varepsilon_t')\| < \infty$ and $\|\mathbb{E}(\varepsilon_t \varepsilon_t')^{-1}\| < \infty$.*
3. *$\{\Sigma_Y^{-1/2} Y_t\}$ is strongly mixing, with $\alpha(m) \leq \exp(-C_1 m^{\gamma_1})$ for some constants $C_1 > 0$ and $\gamma_1 > 0$.*
4. *For any unit vector x and any $s > 0$,*

$$\mathbb{P}\Big(|x' \Sigma_Y^{-1/2} Y_t| > s \mid \mathcal{G}, \Phi\Big) \leq \exp\Big(1 - (s/C_2)^{\gamma_2}\Big),$$

   for some constants $C_2 > 0$ and $\gamma_2 > 0$.
5. *$T = \Omega\big(N^{2/\gamma - 1}\big)$, where $1/\gamma = 1/\gamma_1 + 1/\gamma_2$ and $\gamma < 1$.*

Assumption 2.3 is standard in high-dimensional time-series analysis, see Fan, Liao and Mincheva (2013). It delivers exponential mixing, exponential-type tails, and a growth condition under which Bernstein-type concentration inequalities for dependent data apply, in particular the result of Merlevède, Peligrad and Rio (2011). In substantive terms, these conditions rule out persistence and tail behavior that are too strong relative to the sample size. This ensures that the time-series information in the panel is sufficiently informative to separate systematic dynamic dependence from estimation noise.

The next sparsity restriction enters the perturbation argument.

Assumption 2.4. *Define* $B_N = \min_{i,j} B_{ij}$. *Assume* $B_N = \Omega\big(\log N/N\big)$.

Assumption 2.4 places the model in a sparse regime in which recovery of the latent group structure remains feasible, see Abbe, Bandeira and Hall (2016). In substantive terms, the condition allows the latent network to be sparse, but not so sparse that the group structure disappears from the coefficient matrix. Under Assumptions 2.2–2.4, the OLS estimator of the coefficient matrix is uniformly close to its population counterpart in spectral norm.

Lemma 2.4. *Suppose that the MMNAR in Definition 2.1 satisfies Assumptions 2.1, 2.2, 2.3, and 2.4. Let* $\hat{\Phi}$ *denote the OLS estimator of the coefficient matrix. Then, with high probability,*

$$\|\hat{\Phi} - \mathbb{E}(\Phi)\| \leq O_p\left(\sqrt{\frac{\log N}{NB_N}} + \sqrt{\frac{N}{T}}\right). \tag{2.6}$$

Lemma 2.4 provides the basic perturbation bound. The first term reflects the discrepancy between the realized latent coefficient matrix and its population mean. The second term is the time-series estimation error arising from recovery of Φ from the panel by OLS. The decomposition shows that consistent recovery requires both a sufficiently informative cross section and a sufficiently long time dimension. The cross section identifies the latent group structure, while the time series pins down the realized pattern of dynamic dependence.

The next result transfers this bound to the leading eigenspace.

Lemma 2.5. *Suppose that the MMNAR in Definition 2.1 satisfies Assumptions 2.1, 2.2, 2.3, and 2.4. Let* $\hat{X}$ *denote the matrix of estimated eigenvectors associated with the* K *largest eigenvalues of* $\hat{\Phi}^S$, *and let* X *denote the corresponding population eigenvector matrix. Then there exists an orthonormal matrix* $O \in \mathbb{R}^{K\times K}$ *such that, with high probability,*

$$\|\hat{X} - XO\| \leq O_p\left(\sqrt{\frac{\log N}{B_N}} + \frac{N}{\sqrt{T}}\right). \tag{2.7}$$

Lemma 2.5 shows that the estimated leading eigenspace is close to its population counterpart, up to the usual orthogonal alignment. This is the step that transfers the population simplex geometry to the sample problem. In substantive terms, the result implies that the dominant directions of dynamic dependence recovered from the panel converge to those induced by the latent group structure.

We now state the corresponding bound for the SCORE ratio matrix.

Theorem 2.6. *Suppose that the MMNAR in Definition 2.1 satisfies Assumptions 2.1, 2.2, 2.3, and 2.4. Let R and $\hat{R}$ denote the SCORE ratio matrices constructed from X and $\hat{X}$, respectively. Then*

$$\|\hat{R} - R\| = O_p\left(\sqrt{\frac{\log N}{B_N}} + \frac{N}{\sqrt{T}}\right). \tag{2.8}$$

Theorem 2.6 shows that the feasible ratio matrix is uniformly close to its population counterpart. In particular, if $\sqrt{\frac{\log N}{B_N}} + \frac{N}{\sqrt{T}} \to 0$, then $\|\hat{R} - R\| = o_p(1)$. Combined with the population simplex representation and continuity of the vertex-search, barycentric recovery, and normalization steps, this yields consistency of the MMNAR-Mixed-SCORE estimator $\hat{\Pi}$ for Π. In substantive terms, the theorem implies that, as both the cross-sectional dimension and the time dimension increase, the estimated membership profiles recover each unit's latent pattern of mixed exposure across latent groups.

3 Finite-sample performance

3.1 Simulation design

This section studies the finite-sample performance of MMNAR-Mixed-SCORE under the data-generating process in Definition 2.1. The design varies three features that are central for recovery of the membership structure: the time dimension, the separation of the latent communities, and the number of pure nodes. Throughout, the number of latent communities is fixed at $K = 3$, and the cross-sectional dimension is fixed at $N = 120$.

The community interaction matrix B is taken to be symmetric. Its diagonal entries govern within-community interaction strength, whereas its off-diagonal entries govern between-community interaction strength. We consider three interaction designs, with $(b_{ii}, b_{ij}) \in \{(0.25, 0.01), (0.25, 0.05), (0.5, 0.01)\}$. These choices vary the sharpness of the latent block structure. The design $(0.25, 0.05)$ yields the weakest separation between within- and between-community interactions, whereas $(0.5, 0.01)$ yields the strongest.

We also vary the number of pure nodes. Let N_0 denote the number of pure nodes, so that $N - N_0$ is the number of mixed-membership nodes. In the main design reported in Table 3.2, we consider $N_0 \in \{30, 50, 80\}$. For pure nodes, π_i is one of the canonical basis vectors $[1, 0, 0]$, $[0, 1, 0]$, or $[0, 0, 1]$. For mixed-membership nodes, we use the two profiles $[0.5, 0.5, 0]$ and $[0, 0.5, 0.5]$. This construction preserves a transparent geometry while allowing for nontrivial overlap across communities.

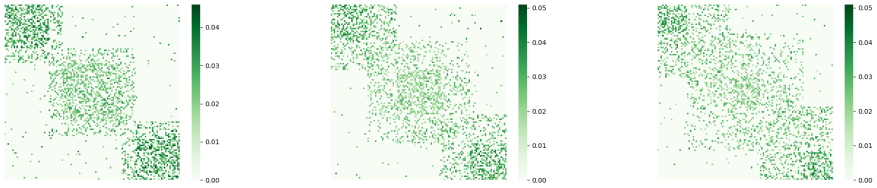

(a) True Φ, $\frac{N_0}{N} = \frac{100}{120}$ (b) True Φ, $\frac{N_0}{N} = \frac{80}{120}$ (c) True Φ, $\frac{N_0}{N} = \frac{60}{120}$

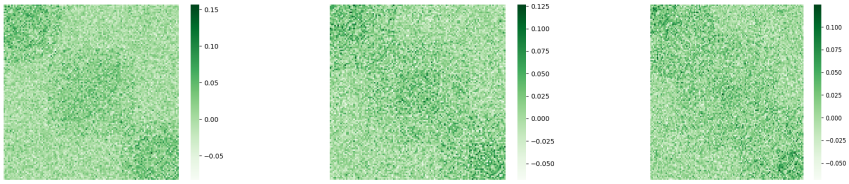

(d) Estimated $\hat{\Phi}$, $\frac{N_0}{N} = \frac{100}{120}$ (e) Estimated $\hat{\Phi}$, $\frac{N_0}{N} = \frac{80}{120}$ (f) Estimated $\hat{\Phi}$, $\frac{N_0}{N} = \frac{60}{120}$

Figure 2: Heatmaps of the true and estimated coefficient matrices for three values of the pure-node fraction, holding the interaction design fixed. Moving from left to right, the share of pure nodes falls from $100/120$ to $60/120$. As the share of mixed-membership nodes rises, the block structure becomes less distinct, and the recovery of Φ becomes less accurate.

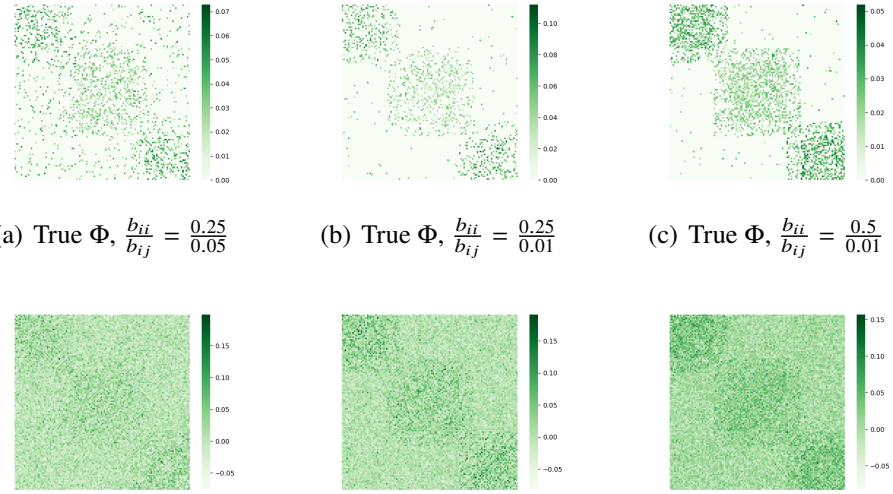

(a) True Φ, $\frac{b_{ii}}{b_{ij}} = \frac{0.25}{0.05}$ (b) True Φ, $\frac{b_{ii}}{b_{ij}} = \frac{0.25}{0.01}$ (c) True Φ, $\frac{b_{ii}}{b_{ij}} = \frac{0.5}{0.01}$

(d) Estimated $\hat{\Phi}$, and $T = 5000$ (e) Estimated $\hat{\Phi}$ and $T = 5000$ (f) Estimated $\hat{\Phi}$ and $T = 5000$

Figure 3: Heatmaps of the true and estimated coefficient matrices under alternative interaction designs and time dimensions. A larger contrast between within-community and between-community interaction yields a sharper block pattern in the population matrix. For a fixed interaction design, the estimated matrix tracks that pattern more closely as T increases.

The remaining parameters are held fixed across designs. The support of active link weights is set to $[\underline{u}, \overline{u}] = [0.3, 1.0]$, and the heterogeneity parameters $\{\theta_i\}_{i=1}^N$ are evenly spaced over $[0.9, 1.0]$. For each specification, we generate data from the MMNAR model and repeat the experiment 500 times. Estimation accuracy is measured by the average squared ℓ_2 loss, $N^{-1}\sum_{i=1}^N \|\hat{\pi}_i - \pi_i\|^2$, and Table 3.2 reports the mean of this loss across replications. The figures reported below provide heatmap illustrations of the associated coefficient matrices.

3.2 Finite-sample results

Table 3.2 reports the main Monte Carlo results. The most stable pattern concerns the time dimension. For every interaction design and every value of N_0, the estimation error declines as T increases. This pattern is consistent with the structure of the estimator. A longer time series sharpens first-step recovery of the coefficient matrix Φ, and this improvement carries over to the spectral recovery of the membership simplex.

	$N = 120$				
	$T = 1000$	2000	3000	4000	5000
$N_0 = 30$					
0.25/0.01	0.14	0.12	0.109	0.099	0.092
0.50/0.01	0.13	0.11	0.094	0.076	0.064
0.25/0.05	0.146	0.142	0.138	0.135	0.130
$N_0 = 50$					
0.25/0.01	0.16	0.12	0.096	0.073	0.060
0.50/0.01	0.15	0.11	0.067	0.048	0.036
0.25/0.05	0.17	0.15	0.14	0.12	0.11
$N_0 = 80$					
0.25/0.01	0.17	0.12	0.085	0.052	0.040
0.50/0.01	0.16	0.10	0.066	0.040	0.031
0.25/0.05	0.19	0.16	0.13	0.114	0.094

Table 1: **Estimation error of MMNAR-Mixed-SCORE.** Each entry reports the average over 500 replications of $N^{-1}\sum_{i=1}^N \|\hat{\pi}_i - \pi_i\|^2$, where $\hat{\pi}_i$ denotes the estimated membership vector and π_i denotes its population counterpart.

A second pattern concerns the separation of the latent block structure. Holding N_0 and T fixed, estimation error is generally smaller when within-community interaction is stronger relative to between-community interaction. In particular, the design $(0.5, 0.01)$ typically outperforms $(0.25, 0.05)$, with $(0.25, 0.01)$ lying between them. This ordering is natural. Stronger within-community interaction yields a clearer separation across latent groups, so the leading eigenspace

of the coefficient matrix more sharply captures the underlying mixed-membership structure.

The role of pure nodes also emerges in the table, although more clearly once the time dimension is sufficiently large. For moderate and large values of T, designs with more pure nodes tend to produce lower estimation error. For example, under $(0.5, 0.01)$ and $T = 5000$, the loss falls from 0.064 at $N_0 = 30$ to 0.036 at $N_0 = 50$ and 0.031 at $N_0 = 80$. The ordering is less regular when T is small. This is consistent with the geometry of the problem. Pure nodes sharpen the simplex only to the extent that the first-step estimate of Φ is already precise enough for that geometry to be visible in the sample. When the time series is short, estimation noise in $\hat{\Phi}$ can obscure part of the gain from additional anchors.

Figure 3 complements Table 3.2 by illustrating the role of block separation and the time dimension. The top row reports the true coefficient matrices under the three interaction designs. As the gap between b_{ii} and b_{ij} widens, the latent block structure becomes more pronounced. The remaining rows report the corresponding estimates at $T = 2000$ and $T = 5000$. For a given interaction design, the estimated coefficient matrix tracks the population structure more closely as T increases.

Figure 2 isolates the role of pure nodes while holding the interaction design fixed. As the share of mixed-membership nodes rises, the block pattern becomes less distinct, and recovery of Φ becomes less accurate. This is the finite-sample manifestation of the pure-node condition underlying identification. Pure nodes determine the vertices of the simplex, whereas a larger mass of mixed nodes compresses the geometry of the membership space and makes recovery more demanding.

The simulation evidence points to a clear conclusion. MMNAR-Mixed-SCORE performs better when the time dimension is longer, the latent block structure is more sharply separated, and the design contains a larger set of pure nodes. These features strengthen the spectral signal associated with the mixed-membership structure and improve finite-sample recovery of Π.

4 Empirical application: mixed-membership estimation in the U.S. financial system

We apply the MMNAR framework to daily volatility data for U.S. financial firms to assess whether cross-sectional dependence is better described by mixed exposure across latent groups than by an exclusive partition. The distinction is economically relevant: large financial institutions routinely combine commercial banking, securities intermediation, asset management, insurance, payments, and real-estate activities, so a pure-membership specification may impose a sharper classification than the data support.

The sample consists of U.S. financial firms with SIC codes between 6000 and 6799, covering depository and non-depository institutions, broker-dealers, insurers, real-estate-related firms, and holding and investment offices. Daily prices are obtained from Yahoo Finance for the period from 3 January 2013 to 31 December 2024. After calendar alignment and removal of incomplete histories, the balanced panel contains $N = 98$ institutions and $T = 3018$ trading days. For firm i on day t, we

use the Parkinson range estimator

$$\tilde{\sigma}_{i,t}^2 = 0.361\big(p_{i,t}^{high} - p_{i,t}^{low}\big)^2, \tag{4.1}$$

where high and low prices are split-adjusted using the close-to-adjusted-close ratio. The empirical input is the panel of log volatilities $\log(\tilde{\sigma}_t^2)$.

We estimate the VAR

$$\log\big(\tilde{\sigma}_t^2\big) = \Phi_0 + \Phi_1 \log\big(\tilde{\sigma}_{t-1}^2\big) + \epsilon_t, \tag{4.2}$$

with lag order one, in line with Guðmundsson and Brownlees (2021). The estimated system is dynamically stable, with a largest eigenvalue 0.865. We fix $K = 3$; formal selection of the number of groups lies outside the scope of this application. We estimate Φ_1 by OLS, symmetrize, and apply MMNAR-Mixed-SCORE to obtain the 98×3 membership matrix $\hat{\Pi}$, whose rows are nonnegative and sum to one.

Table 4 reports the estimated memberships. The central finding is that mixed exposure, rather than exclusive assignment, dominates the cross section. The average largest membership weight is only 0.632: 83 of 98 firms have a largest weight below 0.8, 77 below 0.7, only 12 are estimated as exact pure members, and 16 load positively on all three groups. Many large institutions lie well inside the simplex. GS, JPM, MMC, NDAQ, SPGI, and FITB are close to an even split between groups 2 and 3, while BAC, BRK-B, ICE, MA, CME, TFC, NTRS, and CCI display nontrivial exposure to all three. Pure members are the exception, confined to firms such as AFL, AXP, BRO, IVZ, KKR, MAA (group 2), BX, V, HST (group 3), and FIS, SBAC, UDR (group 1).

The estimated structure is also markedly uneven across groups. Assigning each firm to its largest weight allocates only 6 firms to group 1, against 42 and 50 to groups 2 and 3, with average shares of 0.079, 0.453, and 0.469. Although the latent groups are unobserved, the estimated pattern admits a restrained reading: group 2 concentrates institutions associated with broad intermediation activities, including insurers and brokers (AFL, BRO) and asset managers (IVZ, KKR); group 3 is more heavily populated by market-based and asset-price-sensitive institutions (BX, V, IBKR, MET, STT, and several REITs); group 1 is a smaller, more specialized component (FIS, SBAC, UDR, with nontrivial weights for BK, PFG, CCI, JKHY, FDS). These descriptions are not literal industry labels, and the recovered groups do not reproduce SIC categories: firms from similar segments need not load on the same group, and firms from different segments may exhibit closely related membership vectors.

The application yields two conclusions. First, cross-sectional heterogeneity in the U.S. financial system is well captured by a small number of latent groups. Second, this heterogeneity does not take the form of an exclusive partition; the prevalence of interior membership vectors implies that hard classification would discard economically consequential variation that the mixed-membership specification preserves.

Firm	G1	G2	G3	Firm	G1	G2	G3	Firm	G1	G2	G3
AFL	0	1	0	GS	0	0.532	0.468	PSA	0.264	0.736	0
ALL	0	0.483	0.517	HIG	0	0.492	0.508	O	0	0.531	0.469
AXP	0	1	0	HBAN	0	0.414	0.586	REG	0.146	0.356	0.498
AIG	0	0.462	0.538	IBKR	0	0.207	0.793	SBAC	1	0	0
AMP	0	0.510	0.490	ICE	0.083	0.488	0.429	SPG	0	0.367	0.633
AON	0	0.511	0.489	IVZ	0	1	0	UDR	1	0	0
APO	0	0.691	0.309	JKHY	0.357	0.086	0.557	VTR	0	0.289	0.711
ACGL	0	0.539	0.461	JPM	0	0.533	0.467	WELL	0	0.149	0.851
AJG	0	0.695	0.305	KEY	0	0.361	0.639	WY	0	0.626	0.374
AIZ	0	0.406	0.594	KKR	0	1	0	WRB	0	0.532	0.468
BAC	0.135	0.237	0.628	L	0	0.482	0.518	WFC	0	0.327	0.673
BRK-B	0.289	0.327	0.383	MTB	0	0.220	0.780	WTW	0	0.433	0.567
BLK	0	0.559	0.441	MMC	0	0.500	0.500	ARE	0	0.431	0.569
BX	0	0	1	MA	0.067	0.509	0.424	AMT	0	0.514	0.486
BK	0.571	0.327	0.102	MET	0	0.238	0.762	AVB	0	0.598	0.402
BRO	0	1	0	MCO	0.087	0.533	0.381	BXP	0	0.577	0.423
COF	0	0.503	0.497	MS	0	0.574	0.426	CPT	0.071	0.463	0.466
CBOE	0	0.456	0.544	MSCI	0	0.514	0.486	CBRE	0	0.755	0.245
SCHW	0	0.428	0.572	NDAQ	0	0.487	0.513	CSGP	0.059	0.508	0.433
CB	0	0.003	0.997	NTRS	0.031	0.149	0.820	CCI	0.471	0.292	0.238
CINF	0	0.483	0.517	PNC	0	0.494	0.506	DLR	0	0.464	0.536
C	0.400	0	0.600	PFG	0.632	0.368	0	EQIX	0	0.451	0.549
CME	0.019	0.441	0.540	PGR	0	0.565	0.435	EQR	0	0.612	0.388
CPAY	0.090	0.300	0.611	PRU	0	0.391	0.609	ESS	0	0.434	0.566
ERIE	0	0.514	0.486	RJF	0	0.529	0.471	EXR	0	0.418	0.582
EG	0	0.484	0.516	RF	0	0.571	0.429	FRT	0	0.682	0.318
FDS	0.379	0	0.621	SPGI	0	0.493	0.507	DOC	0	0.611	0.389
FIS	1	0	0	STT	0	0.379	0.621	HST	0	0	1
FITB	0	0.496	0.504	TROW	0.307	0.390	0.303	IRM	0	0.571	0.429
FISV	0	0.471	0.529	TRV	0	0.454	0.546	KIM	0	0.414	0.586
BEN	0	0.509	0.491	TFC	0.270	0.319	0.411	MAA	0	1	0
GPN	0	0.334	0.663	USB	0	0.651	0.349	PLD	0	0.645	0.355
GL	0	0.470	0.530	V	0	0	1				

Table 2: **Estimated three-group mixed memberships of U.S. financial firms.** Each row reports the estimated membership vector for one institution. The columns G1, G2, and G3 denote the three latent groups discussed in the text. Entries are nonnegative and sum to one. Larger entries indicate stronger exposure to the corresponding group.

5 Conclusion

This paper studies a mixed membership network autoregressive model for panel data in which the network coefficient matrix is latent and generated by a Bayesian belief network. The main departure from existing grouped formulations is that units are not restricted to an exclusive partition. Instead, each unit may exhibit mixed exposure across latent groups. This structure allows the model to capture overlapping forms of cross-sectional heterogeneity in dynamic dependence.

We develop an econometric approach for recovering the latent membership structure from panel observations. We establish identification of the mixed-membership component and show that the proposed estimator is consistent as both dimensions of the panel grow. The empirical application to the U.S. financial system indicates that mixed membership is substantively important. Many institutions exhibit non-negligible exposure across more than one latent group, so a hard classification misses an important feature of the underlying dependence structure. More broadly, the results suggest that allowing for mixed exposure across latent groups can improve empirical analysis in large network autoregressive systems when cross-sectional dependence is not well described by disjoint latent types.

References

Abbe, Emmanuel, Afonso S. Bandeira and Georgina Hall. 2016. "Exact recovery in the stochastic block model." *IEEE Transactions on Information Theory* 62(1):471–487.

Airoldi, Edoardo M., David M. Blei, Stephen E. Fienberg and Eric P. Xing. 2008. "Mixed membership stochastic blockmodels." *Journal of Machine Learning Research* 9(65):1981–2014.

Blei, David M., Andrew Y. Ng and Michael I. Jordan. 2003. "Latent Dirichlet allocation." *Journal of Machine Learning Research* 3:993–1022.

Bonhomme, Stéphane and Elena Manresa. 2015. "Grouped patterns of heterogeneity in panel data." *Econometrica* 83(3):1147–1184.

Brownlees, Christian, Guðmundur Stefán Guðmundsson and Gábor Lugosi. 2022. "Community detection in partial correlation network models." *Journal of Business & Economic Statistics* 40(1):216–226.

Chen, Elynn Y., Jianqing Fan and Xuening Zhu. 2023. "Community network auto-regression for high-dimensional time series." *Journal of Econometrics* 235(2):1239–1256.

Chetverikov, Denis and Elena Manresa. 2022. "Spectral and post-spectral estimators for grouped panel data models." *arXiv e-prints* pp. arXiv–2212.

Fan, Jianqing, Yingying Fan, Xiao Han and Jinchi Lv. 2022. "SIMPLE: Statistical inference on membership profiles in large networks." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 84(2):630–653.

Fan, Jianqing, Yuan Liao and Martina Mincheva. 2013. "Large covariance estimation by thresholding principal orthogonal complements." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 75(4):603–680.

Guðmundsson, Guðmundur Stefán. 2026. "Detecting giver and receiver spillover groups in large vector autoregressions." *Journal of Business & Economic Statistics* 44(1):297–308.

Guðmundsson, Guðmundur Stefán and Christian Brownlees. 2021. "Detecting groups in large vector autoregressions." *Journal of Econometrics* 225(1):2–26.

Jin, Jiashun. 2015. "Fast community detection by SCORE." *The Annals of Statistics* 43(1):57–89.

Jin, Jiashun, Zheng Tracy Ke and Shengming Luo. 2024. "Mixed membership estimation for social networks." *Journal of Econometrics* 239(2):105369.

Koller, Daphne and Nir Friedman. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.

Latouche, Pierre, Etienne Birmelé and Christophe Ambroise. 2011. "Overlapping stochastic block models with application to the French political blogosphere." *The Annals of Applied Statistics* 5(1):309–336.

Mao, Xueyu, Purnamrita Sarkar and Deepayan Chakrabarti. 2017. On mixed memberships and symmetric nonnegative matrix factorizations. In *International Conference on Machine Learning*. PMLR pp. 2324–2333.

Mao, Xueyu, Purnamrita Sarkar and Deepayan Chakrabarti. 2021. "Estimating mixed memberships with sharp eigenvector deviations." *Journal of the American Statistical Association* 116(536):1928–1940.

Merlevède, Florence, Magda Peligrad and Emmanuel Rio. 2011. "A Bernstein type inequality and moderate deviations for weakly dependent sequences." *Probability Theory and Related Fields* 151:435–474.

Mugnier, Martin. 2025. "A simple and computationally trivial estimator for grouped fixed effects models." *Journal of Econometrics* 250:106011.

Munro, Evan and Serena Ng. 2022. "Latent Dirichlet analysis of categorical survey responses." *Journal of Business & Economic Statistics* 40(1):256–271.

Rohe, Karl, Sourav Chatterjee and Bin Yu. 2011. "Spectral clustering and the high-dimensional stochastic blockmodel." *The Annals of Statistics* 39(4):1878–1915.

Su, Liangjun, Zhentao Shi and Peter C. B. Phillips. 2016. "Identifying latent structures in panel data." *Econometrica* 84(6):2215–2264.

Tang, Minh, Daniel L. Sussman and Carey E. Priebe. 2013. "Universally consistent vertex classification for latent positions graphs." *The Annals of Statistics* 41(3):1406–1430.

Zhang, Yingying, Huixia Judy Wang and Zhongyi Zhu. 2019. "Quantile-regression-based clustering for panel data." *Journal of Econometrics* 213(1):54–67.

Zhang, Yuan, Elizaveta Levina and Ji Zhu. 2020. "Detecting overlapping communities in networks using spectral methods." *SIAM Journal on Mathematics of Data Science* 2(2):265–283.

Zhu, Xuening and Rui Pan. 2020. "Grouped network vector autoregression." *Statistica Sinica* 30(3):1437–1462.

Zhu, Xuening, Rui Pan, Guodong Li, Yuewen Liu and Hansheng Wang. 2017. "Network vector autoregression." *The Annals of Statistics* 45(3):1096–1123.

Zhu, Xuening, Weining Wang, Hansheng Wang and Wolfgang Karl Härdle. 2019a. "Network quantile autoregression." *Journal of Econometrics* 212(1):345–358.

Zhu, Xuening, Xiangyu Chang, Runze Li and Hansheng Wang. 2019b. "Portal nodes screening for large scale social networks." *Journal of Econometrics* 209(2):145–157.

Appendix: Proofs

Proof of Theorem 2.1. Treat Π as fixed throughout. By construction of the model, $\mathbb{P}(\Phi, A) = \mathbb{P}(\Phi \mid A, \underline{u}, \overline{u})\mathbb{P}(A \mid \Theta, B, \Pi)$. Hence, by iterated expectations,

$$\mathbb{E}(\Phi) = \mathbb{E}\big[\mathbb{E}(\Phi \mid A)\big]. \tag{5.1}$$

Now fix (i, j). If $A_{ij} = 0$, then $\Phi_{ij} = 0$. If $A_{ij} = 1$, then $\Phi_{ij} \mid A_{ij} = 1 \sim \text{Uniform}(\underline{u}, \overline{u})$, so $\mathbb{E}(\Phi_{ij} \mid A) = \frac{\underline{u}+\overline{u}}{2} A_{ij}$. Therefore, $\mathbb{E}(\Phi_{ij}) = \frac{\underline{u}+\overline{u}}{2}\mathbb{E}(A_{ij})$. Under Definition 2.1, $A_{ij} \sim \text{Bernoulli}(\theta_i \pi_i' B \pi_j \theta_j)$, and hence $\mathbb{E}(A_{ij}) = \theta_i \pi_i' B \pi_j \theta_j$. Collecting entries yields $\mathbb{E}(\Phi) = \frac{\underline{u}+\overline{u}}{2}\Theta\Pi B\Pi'\Theta$, which establishes the claim. □

Proof of Theorem 2.2. By Theorem 2.1, the relevant population object is $\mathbb{E}(\Phi) = c\,\Theta\Pi B\Pi'\Theta$, where $c = (\underline{u} + \overline{u})/2$. Since Assumption 2.1(4) fixes $(\underline{u}, \overline{u})$, the scalar c is fixed. Identification therefore reduces to recovering (Θ, Π, B) from $\Theta\Pi B\Pi'\Theta$.

Suppose two parameter tuples, (Θ_1, Π_1, B_1) and (Θ_2, Π_2, B_2), generate the same population matrix. Then

$$\Theta_1\Pi_1 B_1 \Pi_1'\Theta_1 = \Theta_2\Pi_2 B_2 \Pi_2'\Theta_2. \tag{5.2}$$

Assumption 2.1(1) ensures that B is full rank and positive definite. Assumption 2.1(2) requires at least one pure node in each community. Assumption 2.1(3) removes the remaining scale indeterminacy through the normalization $N^{-1}\sum_{i=1}^{N}\theta_i = 1$. Under these conditions, the identification argument in Theorem 2.1 of Zhang, Levina and Zhu (2020) applies directly to (5.2). Hence there exists a permutation matrix P such that $\Pi_2 = \Pi_1 P$, $B_2 = P' B_1 P$, and $\Theta_2 = \Theta_1$. Therefore the parameter tuple $\big((\underline{u}, \overline{u}), \Theta, \Pi, B\big)$ is unique up to a permutation of the community labels. □

Proof of Lemma 2.3. It suffices to show that $\mathrm{Span}\big(\sqrt{(\underline{u}+\overline{u})/2}\,\Theta\Pi\big) = \mathrm{Span}(X)$. Let $c = (\underline{u}+\overline{u})/2$ and define $C = \sqrt{c}\,\Theta\Pi$. By Theorem 2.1, $\mathbb{E}(\Phi) = CBC'$.

By Assumption 2.1(2), there exists at least one pure node in each community, so Π has rank K. Since Θ is diagonal with strictly positive diagonal entries by Assumption 2.1(3), C also has rank K. Because B is full rank by Assumption 2.1(1), $\mathbb{E}(\Phi) = CBC'$ has rank K as well.

Now $\mathrm{Col}\big(\mathbb{E}(\Phi)\big) \subseteq \mathrm{Col}(C)$, since $\mathbb{E}(\Phi) = CBC'$. But both spaces have dimension K, and therefore they coincide. Let $X = [x_1, \ldots, x_K]$ collect the eigenvectors associated with the K largest eigenvalues of $\mathbb{E}(\Phi)$. Then the columns of X form a basis for $\mathrm{Col}\big(\mathbb{E}(\Phi)\big)$, and hence also for $\mathrm{Col}(C)$. Since C has full column rank, there exists a unique nonsingular matrix $P \in \mathbb{R}^{K\times K}$ such that $X = CP = \big(\sqrt{(\underline{u}+\overline{u})/2}\,\Theta\Pi\big)P$. This proves the lemma. □

Proof of Lemma 2.4. We decompose the estimation error as $\hat{\Phi} - \mathbb{E}(\Phi) = (\hat{\Phi} - \Phi) + (\Phi - \mathbb{E}(\Phi))$. Hence, by the triangle inequality,

$$\|\hat{\Phi} - \mathbb{E}(\Phi)\| \le \|\hat{\Phi} - \Phi\| + \|\Phi - \mathbb{E}(\Phi)\|. \tag{5.3}$$

Consider first the second term on the right-hand side. By Theorem 2.1, $\mathbb{E}(\Phi) = c\,\Theta\Pi B\Pi'\Theta$, where $c = (\underline{u}+\overline{u})/2$. This has the same factor structure as the corresponding population coefficient matrix in Guðmundsson and Brownlees (2021), after replacing the pure-membership matrix by the mixed-membership matrix Π. Under Assumptions 2.2 and 2.4, the same concentration argument yields $\|\Phi - \mathbb{E}(\Phi)\| = O_p\big(\sqrt{\log N/(NB_N)}\big)$.

Next consider the first term in (5.3). Conditional on the realized latent graph and the corresponding coefficient matrix Φ, the MMNAR satisfies (2.4). Since the first-step regression is run on demeaned data, the intercept plays no role in the OLS representation. The estimation error then admits the usual score decomposition, and the argument in Lemma 3 of Guðmundsson and Brownlees (2021) applies under Assumption 2.3. Hence $\|\hat{\Phi} - \Phi\| = O_p\big(\sqrt{N/T}\big)$.

Combining the two bounds with (5.3) gives $\|\hat{\Phi} - \mathbb{E}(\Phi)\| = O_p\big(\sqrt{\log N/(NB_N)} + \sqrt{N/T}\big)$, which proves the result. □

Proof of Lemma 2.5. Let $M = \mathbb{E}(\Phi)$ and let $\hat{M} = \hat{\Phi}^S = (\hat{\Phi} + \hat{\Phi}')/2$. Since M is symmetric, $\|\hat{M} - M\| \le \|\hat{\Phi} - M\|$. By Lemma 2.4, $\|\hat{M} - M\| = O_p\big(\sqrt{\log N/(NB_N)} + \sqrt{N/T}\big)$.

Let X collect the K leading eigenvectors of M, and let $\hat{X}$ collect the K leading eigenvectors of $\hat{M}$. By Assumption 2.2, the nonzero eigenvalues of M are of order NB_N. Applying Lemma A.1 of Tang, Sussman and Priebe (2013), there exists an orthonormal matrix $O \in \mathbb{R}^{K\times K}$ such that

$$\|\hat{X} - XO\|_F \le \frac{\sqrt{K}\,\|\hat{M} - M\|\big(\sqrt{\|\hat{M}\|} + \sqrt{\|M\|}\big)}{\lambda_K(M)}. \tag{5.4}$$

Since $\|M\| \asymp NB_N$, $\lambda_K(M) \asymp NB_N$, and $\|\hat{M}\| = O_p(NB_N)$ by Weyl's inequality, the right-hand side of (5.4) is $O_p\big(\sqrt{\log N/B_N} + N/\sqrt{T}\big)$. Since the operator norm is bounded by the Frobenius norm, the same rate holds for $\|\hat{X} - XO\|$. This proves the lemma. □

Proof of Theorem 2.6. The argument extends the proof of Theorem 3.1 in Jin, Ke and Luo (2024), using Lemma 2.5. Let O be the orthonormal matrix in Lemma 2.5, and define $\tilde{X} = \hat{X}O$. Write $\tilde{x}_{ik}$ and x_{ik} for the (i,k) entries of $\tilde{X}$ and X, respectively. Since the truncation map $a \mapsto \mathrm{sign}(a)\min\{|a|, S\}$ is Lipschitz with constant one, it suffices to control the untruncated ratio map.

For a fixed i, let $r_i = x_{i1}^{-1} X_{i,0}$ and $\hat{r}_i = \tilde{x}_{i1}^{-1} \tilde{X}_{i,0}$, where $X_{i,0}$ and $\tilde{X}_{i,0}$ denote the ith rows after removing the first column. Then

$$\hat{r}_i - r_i = \frac{1}{\tilde{x}_{i1}}(\tilde{X}_{i,0} - X_{i,0}) - \frac{\tilde{x}_{i1} - x_{i1}}{\tilde{x}_{i1}} r_i. \tag{5.5}$$

By Lemma C.4 of Jin, Ke and Luo (2024), $\|r_i\|$ is uniformly bounded. By Lemma C.3 of Jin, Ke and Luo (2024), the leading eigenvector is bounded away from zero after scaling by degree heterogeneity, so there exists a constant $c_0 > 0$ such that $x_{i1} \geq c_0 \Theta_{\min}/\|\Theta\|$ uniformly in i. Lemma 2.5 then implies that, with probability approaching one, $\tilde{x}_{i1} \geq c_1 \Theta_{\min}/\|\Theta\|$ for some constant $c_1 > 0$. Therefore the denominator in (5.5) is uniformly bounded away from zero.

Using (5.5) and Lemma 2.5, we obtain $\max_{1\leq i\leq N} \|\hat{r}_i - r_i\| = O_p\big(\sqrt{\log N/B_N} + N/\sqrt{T}\big)$. Since truncation does not enlarge the error, the same rate holds for the ratio matrices, namely $\|\hat{R} - R\| = O_p\big(\sqrt{\log N/B_N} + N/\sqrt{T}\big)$. This proves the theorem. □