

NETWORK ROBUST INFERENCE FOR FIXED-EFFECT REGRESSIONS

KENSUKE SAKAMOTO

SEPTEMBER 25, 2025

PRELIMINARY: Please click [HERE](#) for the most recent version

Department of Economics, University of Wisconsin-Madison

ABSTRACT. Fixed effects in network data are central to many economic applications, yet most existing methods assume conditional independence of networked observations—an assumption whose plausibility is often application-dependent. This paper studies fixed-effect regressions on network data under a conditional dependence structure in which errors arise from both node- and edge-level shocks that are not fully captured by the fixed effects. We show that the least-squares estimator of the fixed effects can be inconsistent due to a persistent noise term induced by the dependence structure. By leveraging information from the regression residuals, we propose new inference methods for fixed effects that explicitly account for dependence. We also introduce a bias-correction procedure for variance estimation under dependence. An empirical application using worker–firm matched data demonstrates the practical relevance of our methods and highlights the substantial impact of dependence on inference and variance estimation.

E-mail address: `ksakamoto2@wisc.edu`.

The author thanks Jack Porter, Bruce Hansen, Harold Chiang, and Yong Cai for their advice and support. The author also appreciates comments from Xiaoxia Shi, Kohei Yata, and participants at the New York Camp Econometrics XIX, the Bristol Econometrics Study Group Annual Conference 2025, and Celebrating James MacKinnon Conference. The author gratefully acknowledges financial support from the Douglas W. and Sherry A. Caves Dissertation Fellowship.

1. INTRODUCTION

Network data—describing interactions among economic agents—have become increasingly prevalent in applied economics, providing additional dimensions that are not available in individual-level data. One important application of network data, given this extra dimensionality, is the identification and estimation of agent-level fixed effects to capture unobserved heterogeneity that is crucial for understanding economic phenomena. For example, Card et al. (2013) estimate the contribution of firm-level heterogeneity to wage inequality by leveraging worker mobility, which creates a network connecting firms and generates wage differences, from which firm-level fixed effects can be extracted.¹ Such network structures naturally induce cross-sectional dependence in outcomes, such as firm-to-firm wage differences, because agent-level shocks are shared across networked observations. This paper examines how dependence in network data affects the estimation and inference of fixed effects and proposes a novel inference method that explicitly accounts for the dependence structure.

Despite the networked structure of the data, most state-of-the-art econometric methods for fixed-effects estimation rely on the assumption that outcomes are independent, or at most weakly dependent, conditional on fixed effects. For example, Engbom and Moser (2022) estimate changes in workers’ and firms’ contributions to wage determination using the methodology of Kline et al. (2020), which assumes conditional independence of wages. However, the plausibility of this assumption is application-dependent. For example, in worker-firm matched data, the wage a worker receives from firm j may be correlated with the wage received by another worker at the same firm, perhaps due to a firm-level shock such as the introduction of new technology that is not fully captured by time-invariant firm fixed effects. Such shocks can induce correlation in wage differences between firm j and other firms, which is ruled out by the conventional independence assumption. The presence of network dependence can alter the behavior of fixed-effect estimators and invalidate current econometric approaches.

In this paper, we study fixed-effect regressions on network data with dependent errors. The data are represented by a graph whose nodes correspond to economic agents and whose edges represent interactions; the unit of observation is an edge, and each edge has

¹Other examples include studies in labor: Abowd et al. (1999); Engbom and Moser (2022); Bonhomme et al. (2023), education: Jackson (2013); Bacher-Hicks and Koedel (2023), health: Finkelstein et al. (2016, 2021), and trade: Bernard et al. (2022).

an associated scalar outcome. We consider a linear model in which an edge outcome is driven by the difference between the incident nodes’ fixed effects, edge-level covariates, and an edge-level error term that may be correlated across edges. For example, in labor economics applications, the graph’s nodes are firms and an edge links a pair of firms when a worker moves between them; the edge outcome is the wage change for that mover, determined in part by the difference in the two firms’ fixed effects.

In this high-dimensional setting, the dependence across outcomes poses significant challenges for inference on fixed effects. To deal with these challenges and for analytical tractability, we assume that errors are generated by both node-level and edge-level shocks, which are not controlled by the fixed effects, introducing conditional dependence among the errors. This dependence structure allows for independence as a special case.

When evaluating the least-squares estimator for fixed effects, we highlight the fundamental trade-off between the connectivity of the graph and the dependence structure. The underlying graph must be well-connected to consistently estimate the fixed effects, as otherwise, there is insufficient information to identify fixed effects. This issue, termed “limited mobility” (e.g., Andrews et al., 2008; Bonhomme et al., 2023), is well known in the literature. However, by allowing for dependence, we find a new countervailing effect. If the graph is too densely connected, the resulting dependence may become too strong. This, in turn, hinders precise estimation of the fixed effects.

Our main results are as follows. First, we derive a first-order approximation for the fixed-effect estimator that features a persistent noise term stemming from the dependence structure and hindering consistent estimation. We also show that the distribution of this noise term can be identified and consistently estimated by exploiting information contained in residuals from the regression. Using this result, we propose a non-standard inference method for each fixed effect *à la* Conley and Taber (2011), who provide an inference method for treatment effects under small treatment group asymptotics. Specifically, we propose confidence intervals for each fixed effect and a hypothesis test for homogeneity of fixed effects across nodes within a specific group. We show that these procedures are asymptotically valid. Our simulation exercises show that our inference methods perform well in finite samples provided that the graph is relatively well-connected.

Second, we introduce a new bias-correction method for moment estimation of fixed effects under the dependence structure. Focusing specifically on the variance of fixed

effects, we characterize the bias inherent in the plug-in estimator and propose a bias-correction method that is consistent under the dependence structure. Additionally, we show that the bias stemming from network dependence can dominate the bias addressed by existing methods, which typically assume independence in errors (e.g., Andrews et al., 2008; Kline et al., 2020; Bonhomme et al., 2023).

In our empirical application, we implement our inference and bias-correction methods using worker-firm matched data from the Veneto Worker History (VWH) file. For inference, we construct a novel confidence interval for the fixed effect of the most central firm in the mobility network. For bias correction, we show that our method substantially reduces the bias in the plug-in estimator of the variance of firm fixed effects. Our results highlight that compared to conventional methods that assume independence, accounting for dependence can significantly affect both inference for fixed effects and variance estimation.

Our paper contributes to the literature on fixed-effects estimation in network data (Abowd et al., 1999; Andrews et al., 2008; Graham, 2017; Jochmans and Weidner, 2019; Kline et al., 2020; Bonhomme et al., 2023). This body of work generally assumes conditional independence of outcomes. The most closely related work is Jochmans and Weidner (2019), which develops finite-sample theory and establishes asymptotic normality for fixed effects estimators under the assumption of conditional independence. Jochmans and Weidner (2019) also covers first-order approximations under a weak dependence structure. In contrast, our paper introduces a strong dependence structure and provides a finite-sample theory and inference method for fixed effects under this dependence structure.

The estimation of moments of fixed effects is also a key topic in this literature. Andrews et al. (2008) provide a bias-correction approach for the estimation of variance components under independence and homoskedasticity. Kline et al. (2020) extend this method to accommodate heteroskedasticity, maintaining the independence assumption or allowing at most weak cluster dependence. Our paper is the first to introduce a bias-correction approach for variance components estimation under the strong dependence structure, showing that the bias stemming from network dependence is non-negligible and can dominate the bias addressed by the existing methods.

We are also connected to the literature on two-way/dyadic cluster-robust inference (Cameron et al., 2011; Cameron and Miller, 2014; Tabord-Meehan, 2019; Verdier, 2020; Davezies et al., 2021; Menzel, 2021; Chiang et al., 2024). While our dependence structure

shares elements with these papers, our focus differs by centering on estimating fixed effects, which are typically treated as nuisance parameters. We show that there is a fundamental difficulty in conducting inference on fixed effects under the dependence structure and develop new methods to address this issue.

Finally, this paper connects to the literature on inference with few treated units (Conley and Taber, 2011; MacKinnon and Webb, 2018, 2020); see Alvarez et al. (2025) for a recent review. We are the first to adapt the approach of Conley and Taber (2011) to construct confidence intervals for fixed effects and joint hypothesis tests for groups of fixed effects under the network dependence structure considered here. Intuitively, a node together with its incident edges can be viewed as “few treated units” and the remainder of the graph as “control units” in the sense of Conley and Taber (2011); we exploit this analogy to estimate the distribution of the persistent noise term in the first-order approximation to conduct inference.

2. MODEL

In this section, we introduce the data structure and the linear model we consider in this paper. We then introduce the dependence structure of error terms, along with a useful decomposition to understand the behavior of the estimator.

2.1. Setup. We consider network data based on a directed multigraph $\mathcal{G} = (V, E, s, t)$, consisting of a node set V , an edge set E , and source and target functions $s, t : E \rightarrow V$. The node set $V = \{1, \dots, n\}$ contains n nodes, and the edge set $E = \{1, \dots, m\}$ contains m edges. We use e to denote a typical edge in E . For each edge $e \in E$, the source function $s(e)$ specifies its origin node, while the target function $t(e)$ specifies its destination node.

Because multiple edges can exist between the same pair of nodes, we define an edge subset $E_{(i,j)} = E_{(j,i)} \subset E$ as the set of edges whose origin or destination is either node i or j :

$$E_{(i,j)} = \{e \in E : s(e) = i, t(e) = j \text{ or } s(e) = j, t(e) = i\} \text{ for each } i, j \in V.$$

We allow $E_{(i,j)}$ to be empty when there are no edges between nodes i and j . By convention, we do not allow self-loops, i.e., $E_{(i,i)} = \emptyset$ for all $i \in V$. Additionally, for each $i \in V$, let $E_i^s = \{e \in E : s(e) = i\}$ be the set of edges that originate from node i and $E_i^t = \{e \in E : t(e) = i\}$ be the set of edges that terminate at node i . Then, $E_i = E_i^s \cup E_i^t$ is the set of edges incident to node i . We define the out-degree of node i as $d_i^s = |E_i^s|$, the number of

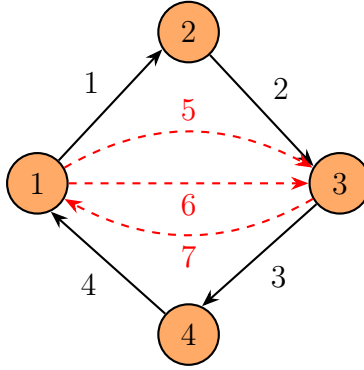
edges flowing out of node i , and the in-degree as $d_i^t = |E_i^t|$, the number of edges flowing into node i . We refer to $d_i = d_i^s + d_i^t$ as the (total) degree of node i , which is the number of edges incident to node i .

To illustrate the notation, consider the directed multigraph \mathcal{G} shown in Figure 2.1, with nodes $V = \{1, 2, 3, 4\}$ and edges $E = \{1, \dots, 7\}$. Focus on node 1: this node is the source of edges $1 \in E_{(1,2)} = \{1\}$, $5, 6 \in E_{(1,3)} = \{5, 6, 7\}$, and the target of edges $4 \in E_{(1,4)} = \{4\}$ and $7 \in E_{(1,3)}$. Consequently, we have

$$s(1) = s(5) = s(6) = t(4) = t(7) = 1,$$

and so $E_1^s = \{1, 5, 6\}$ and $E_1^t = \{4, 7\}$. The out-degree of node 1 is $d_1^s = 3$, the in-degree is $d_1^t = 2$, and thus the degree is $d_1 = d_1^s + d_1^t = 5$.

FIGURE 2.1. A directed multigraph \mathcal{G} with $n = 4$ nodes and $m = 7$ edges.



The following matrices associated with graph \mathcal{G} are crucial for our analysis. First, the $n \times n$ (symmetric) adjacency matrix \mathbf{A} is defined as

$$\mathbf{A}_{i,j} = |E_{(i,j)}|,$$

for $i, j \in V$. Note that \mathbf{A} is the adjacency matrix of the undirected graph derived from the directed graph \mathcal{G} by ignoring the direction of edges and counting the number of edges between each pair of nodes. The degree of node i is also defined as $d_i = \sum_{j \neq i} \mathbf{A}_{i,j}$. Let \mathbf{D} be the $n \times n$ diagonal matrix with degrees d_i placed along its diagonal.

Next, define the $m \times n$ incidence matrix \mathbf{B} as follows:

$$\mathbf{B}_{e,i} = \begin{cases} 1 & \text{if } t(e) = i \\ -1 & \text{if } s(e) = i \\ 0 & \text{otherwise} \end{cases}$$

for $e \in E$ and $i \in V$. The incidence matrix \mathbf{B} maps a node-level vector to an edge-level vector of differences. For example, for a vector $\mathbf{v} \in \mathbb{R}^n$, the e -th element of $\mathbf{B}\mathbf{v}$ equals the difference $v_{t(e)} - v_{s(e)}$. Finally, these matrices are connected through the graph Laplacian \mathbf{L} , defined as

$$\mathbf{L} = \mathbf{B}'\mathbf{B} = \mathbf{D} - \mathbf{A}.$$

As we will see later, the graph Laplacian plays a crucial role in our analysis.

Remark 1. *We focus on unweighted directed multigraphs for simplicity, but our analysis can be extended to weighted graphs as in Jochmans and Weidner (2019). In a weighted graph, each edge $e \in E$ is associated with a positive weight $w_e > 0$. The adjacency matrix \mathbf{A} is then defined as $\mathbf{A}_{i,j} = \sum_{e \in E(i,j)} w_e$ for $i, j \in V$, and the incidence matrix \mathbf{B} is defined as $\mathbf{B}_{e,i} = \sqrt{w_e}$ if $t(e) = i$, $\mathbf{B}_{e,i} = -\sqrt{w_e}$ if $s(e) = i$, and $\mathbf{B}_{e,i} = 0$ otherwise. Focusing on unweighted graphs simplifies the notation and exposition, and does not affect the generality of our results when applied to Example 1 and our empirical application.*

2.2. Linear Model. Suppose that each edge $e \in E$ is equipped with a scalar outcome y_e and a p -dimensional covariate X_e . These outcomes and covariates are stacked into the m -dimensional vector $\mathbf{y} = (y_e)_{e=1}^m$ and the $m \times p$ matrix $\mathbf{X} = (X_e)_{e=1}^m$, respectively. We consider the following linear model:

$$\mathbf{y} = \mathbf{B}\boldsymbol{\alpha} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\alpha} = (\alpha_i)_{i=1}^n$ is the n -dimensional vector of node-level fixed effects, $\boldsymbol{\beta}$ is the p -dimensional vector of coefficients for the covariates, and $\boldsymbol{\epsilon} = (\epsilon_e)_{e=1}^m$ is the m -dimensional vector of edge-level errors. For now, we focus on estimation of $\boldsymbol{\alpha}$ and leave out covariates or equivalently treat $\boldsymbol{\beta}$ as known and redefine $\mathbf{y} - \mathbf{X}\boldsymbol{\beta}$ as \mathbf{y}^2 . Then, the model can be simplified to

$$\mathbf{y} = \mathbf{B}\boldsymbol{\alpha} + \boldsymbol{\epsilon}, \tag{2.1}$$

which we refer to as our regression model.

²See Appendix A for a discussion on the case where $\boldsymbol{\beta}$ is unknown and needs to be estimated.

Element-wise, the regression model (2.1) is written as

$$y_e = \alpha_{t(e)} - \alpha_{s(e)} + \epsilon_e,$$

for $e \in E$. Thus, each outcome is driven by the difference between the destination node fixed effect $\alpha_{t(e)}$ and the origin node effect $\alpha_{s(e)}$. The following example illustrates how such differences in fixed effects naturally arise in economic applications.

Example 1. (*Two-Period AKM model*): An important example of the model (2.1) arises from the so-called AKM model (Abowd et al., 1999), which links worker and firm effects to wages. Leaving out covariates, in this model, worker g is employed by firm $J(g, t)$ at time t , and the (log) wage of worker g at time t is given by

$$w_{g,t} = \beta_g + \alpha_{J(g,t)} + u_{g,t}, \quad (2.2)$$

where β_g is the worker fixed effect, $\alpha_{J(g,t)}$ is the firm fixed effect corresponding to the worker's employer at time t , and $u_{g,t}$ is the idiosyncratic error. Suppose that the economy lasts for two periods, as considered in Section 2 of Kline et al., 2020. Then, 'movers,' i.e., workers g with $J(g, 1) \neq J(g, 2)$, form an edge e by transitioning from firm $s(e) = J(g, 1)$ to firm $t(e) = J(g, 2)$. Then, we can write the wage difference as

$$\underbrace{w_{g,2} - w_{g,1}}_{\equiv y_e} = \underbrace{\alpha_{J(g,2)} - \alpha_{J(g,1)}}_{=\alpha_{t(e)} - \alpha_{s(e)}} + \underbrace{u_{g,2} - u_{g,1}}_{\equiv \epsilon_e},$$

which corresponds to the model (2.1).

Theorem 1 of Jochmans and Weidner (2019) shows that when the graph \mathcal{G} is connected and the fixed effects α are normalized such that $\sum_{i \in V} \alpha_i d_i = 0$, the least-squares estimator for α is uniquely given by:

$$\hat{\alpha} = (\mathbf{B}'\mathbf{B})^* \mathbf{B}'\mathbf{y} = \mathbf{L}^* \mathbf{B}'\mathbf{y}, \quad (2.3)$$

where, for any $n \times n$ matrix \mathbf{C} , \mathbf{C}^* denotes the pseudo-inverse defined as:

$$\mathbf{C}^* = \mathbf{D}^{-1/2} (\mathbf{D}^{-1/2} \mathbf{C} \mathbf{D}^{-1/2})^+ \mathbf{D}^{-1/2}$$

with $^+$ representing the Moore-Penrose inverse. Throughout the paper, we maintain such assumptions on the graph structure and the normalization of fixed effects that validate the estimator (2.3):

Assumption 1. *The graph \mathcal{G} is connected in the following sense: for any $i, j \in V$, there exists a path between i and j on the undirected graph derived from \mathcal{G} . Furthermore, the fixed effects $\boldsymbol{\alpha}$ are normalized such that $\sum_{i \in V} \alpha_i d_i = 0$.*

If the graph is not connected, the normalization condition must be applied separately to each connected component of the graph, and the subsequent analysis can then be applied to each component individually. In practice, researchers typically focus on the largest connected component if it constitutes a significant portion of the entire graph (e.g., Engbom and Moser, 2022).

The specific normalization employed in Assumption 1 is not crucial for our analysis, and alternative normalizations can be used to obtain similar results. For example, we could normalize the fixed effects by imposing $\sum_{i \in V} \alpha_i = 0$, which leads to a least-squares estimator with a different pseudo-inverse:

$$\hat{\boldsymbol{\alpha}} = (\mathbf{B}'\mathbf{B})^+\mathbf{B}'\mathbf{y}.$$

Another common normalization is to set $\alpha_i = 0$ for a specific node $i \in V$. In this case, the least-squares estimator can be written as

$$\hat{\boldsymbol{\alpha}} = (\mathbf{B}'_{-i}\mathbf{B}_{-i})^{-1}\mathbf{B}'_{-i}\mathbf{y},$$

where \mathbf{B}_{-i} is the $m \times (n-1)$ matrix obtained by removing the i -th column of \mathbf{B} . Although this normalization is often used in practice, it complicates theoretical analysis by breaking the direct link to graph-related objects such as the graph Laplacian \mathbf{L} . As discussed in Kline (2024), one can nonetheless relate $(\mathbf{B}'_{-i}\mathbf{B}_{-i})^{-1}\mathbf{B}'_{-i}$ to \mathbf{L}^* using the results of Bozzo (2013). We leave the exploration of that connection to future work, focusing here on the normalization in Assumption 1.

Remark 2. *Note that randomness in our setup arises solely from the error vector $\boldsymbol{\epsilon}$; the fixed effects $\boldsymbol{\alpha}$ and the graph \mathcal{G} , and hence the matrices $\mathbf{A}, \mathbf{B}, \mathbf{D}, \mathbf{L}$, are treated as fixed. Accordingly, all probabilistic statements are understood as conditional on $\boldsymbol{\alpha}$ and \mathcal{G} . This conditional approach is standard in the literature on fixed-effect estimation with network data (e.g., Jochmans and Weidner, 2019; Kline et al., 2020; Bonhomme et al., 2023).*

2.3. Dependence Structure. Next, we introduce a model of dependence for the edge-level errors ϵ_e . Previous literature has typically imposed independence of the error terms

across edges, or at most weak dependence, meaning that the covariance matrix of ϵ is nearly diagonal. However, since outcomes and errors are defined at the edge level, more realistic dependence structures would allow for dependence between edges that share a node, which is ruled out by the independence or weak dependence assumptions.

Here, we focus on a first-order strong dependence structure: we allow for sources of dependence including (i) node-level shocks, which induce dependence between edges sharing a node; and (ii) edge-level shocks, which induce dependence between edges in the same edge subset. These shocks are shared across edges that either share a node or belong to the same edge subset, respectively, and can generate strong dependence among the errors at the first order.

Specifically, for each edge $e \in E$, we assume the following structure for ϵ_e :

$$\epsilon_e = f(U_{s(e)}, U_{t(e)}, V_e); \quad \mathbb{E}[\epsilon_e] = 0, \quad (2.4)$$

where f is a measurable function unknown to the researcher and $U_{s(e)}, U_{t(e)}$ and V_e are random vectors. Specifically, $U_{s(e)}$ and $U_{t(e)}$ represent node-level shocks associated with the source and target nodes of edge e , while V_e is an edge-level shock associated with edge e . Since our analysis does not depend on their dimension, in the following, we treat each U_i and V_e as one-dimensional random variables unless otherwise specified. Importantly, we do not require prior knowledge of the form of f : it may be additive, interactive, or of any other form.

We impose the following structure on (2.4) to facilitate the analysis:

Assumption 2. *The error term $\epsilon = (\epsilon_e)_{e=1}^m$ satisfies the following conditions:*

- (i) *For each $e \in E$, $\mathbb{E}[\epsilon_e] = 0$ and $\max_{e \in E} \mathbb{E}[\epsilon_e^4] \leq C < \infty$ for some absolute constant $C > 0$.*
- (ii) *The node-level shocks $U_i, i \in V$ are independently and identically distributed. The edge-level shocks V_e satisfy the following independence structure:*

$$V_e \perp\!\!\!\perp V_{e'} \text{ if } \{s(e), t(e)\} \neq \{s(e'), t(e')\}.$$

- (iii) *$(U_i)_{i=1}^n$ is independent of $(V_e)_{e=1}^m$*

Assumption 2 imposes a specific dependence structure on the error terms ϵ . Part (i) ensures that the errors are centered at zero and possess uniformly bounded fourth

moments. Part (ii) requires that the node-level shocks U_i be homogeneously distributed across nodes, which is a strong restriction but simplifies the analysis. In contrast, the edge-level shocks V_e are allowed to have heterogeneous distributions and may exhibit correlation among them if two edges belong to the same edge subset $E_{(i,j)}$. For example, one might write the covariance structure of V_e as follows:

$$\text{Cov}(V_e, V_{e'}) = \begin{cases} v_e^2 & \text{if } e = e'; \\ v_{e,e'} & \text{if } e, e' \in E_{(i,j)} \text{ for some } i, j \in V; \\ 0 & \text{otherwise,} \end{cases}$$

where v_e^2 and $v_{e,e'}$ can be different across $e, e' \in E$. Finally, part (iii), the mutual independence between the node-level shocks $(U_i)_{i \in V}$ and the edge-level shocks $(V_e)_{e \in E}$, is a common assumption in the network and two-way clustering literature (e.g., Graham, 2017; Menzel, 2021; Chiang et al., 2024). This assumption is not as restrictive as it might seem because the function f in (2.4) can incorporate both U_i and V_e in a nonlinear fashion, for example, by allowing them to interact with each other.

An important implication of the dependence structure in (2.4) under Assumption 2 is that the error terms ϵ are dependent across edges either through (i) shared node-level shocks U_i when two edges share a common node, or (ii) correlation induced by edge-level shocks when two edges belong to the same edge subset $E_{(i,j)}$. Consequently, we can write the covariance structure of ϵ as

$$\text{Cov}(\epsilon_e, \epsilon_{e'}) = \begin{cases} \sigma_e^2 & \text{if } e = e'; \\ \sigma_{e,e'} & \text{if } e \neq e' \text{ and } e, e' \in E_i \text{ for some } i \in V; \\ 0 & \text{otherwise} \end{cases}, \quad (2.5)$$

for $e, e' \in E$. Note that we are allowing for heteroskedasticity in the error terms, as the variance σ_e^2 and covariance $\sigma_{e,e'}$ can differ across edges due to the heterogeneity of the edge-level shocks V_e .

Note that the dependence structure in (2.4) does not capture higher-order correlations among edges that are indirectly connected through common neighbors. While this restriction simplifies the analysis, it is flexible enough to accommodate a broad range of first-order correlations, and it subsumes the traditional assumption of independence as a special case.

We can connect the dependence structure in (2.4) to the AKM model in Example 1:

Example 2. (*Two-Period AKM model, continued*). Consider the setting described in Example 1. Recall that the error term in this model is given by $\epsilon_e = u_{g,2} - u_{g,1}$ for a mover g . This error term can absorb misspecification in the linear model (2.2), such as time-varying firm effects (Engbom et al., 2023; Lachowska et al., 2023) and match-specific heterogeneity (Bonhomme et al., 2019) that depends on both destination and origin firms, as captured in search models with on-the-job search (e.g., Postel-Vinay and Robin, 2002; Bagger et al., 2014; Di Addario et al., 2023). For instance, if firm effects vary over time as $\alpha_{J(g,t),t} = \alpha_{J(g,t)} + U_{J(g,t),t}$ for $t = 1, 2$, where $U_{J(g,t),t}$ is a firm-level shock at time t , then the error term contains $U_{J(g,2),2} - U_{J(g,1),1}$ and can be expressed in the form of (2.4) by setting $U_i = (U_{i,1}, U_{i,2})$ for each firm $i \in V$ and $V_e = u_{g,2} - u_{g,1}$. Similarly, if there is match-specific heterogeneity that depends on interactions among the worker, origin firm, and destination firm, we can express ϵ_e in the form of (2.4) by viewing U_i as a vector that includes private information of firm i affecting bargaining with workers and competition with other firms, and V_e as a vector that includes any worker-specific information that affects bargaining with firms.

2.4. Decomposition. To isolate the impact of the dependence structure on the fixed-effect estimator $\hat{\alpha}$, we propose the following decomposition. For each $e \in E$, define the functions:

$$\tau_e^s(\cdot) \equiv \mathbb{E}[\epsilon_e | U_{s(e)} = \cdot], \quad \tau_e^t(\cdot) \equiv \mathbb{E}[\epsilon_e | U_{t(e)} = \cdot].$$

These functions capture, respectively, the effects of the origin and destination node-level shocks $U_{s(e)}$ and $U_{t(e)}$ on the error term ϵ_e . Note that they are not necessarily equal; the function f in (2.4) can be asymmetric in its arguments $U_{s(e)}$ and $U_{t(e)}$.

For example, suppose we have

$$\epsilon_e = U_{t(e)}U_{s(e)} + U_{t(e)} - 2U_{s(e)} + 1,$$

and assume that $\mathbb{E}[U_i] = 1$. Then we can show that

$$\tau_e^s(u) = -u + 2 \text{ and } \tau_e^t(u) = 2u - 1,$$

for any $u \in \mathcal{U}$, the support of U_i . This example illustrates that the influence of the node-level shocks on the error term can differ depending on whether the shock comes from the origin or destination node.

To reduce dimensionality, we impose the following homogeneity on these functions:

Assumption 3. *For any $e, e' \in E$, we have*

$$\tau_e^s = \tau_{e'}^s \equiv \tau^s, \quad \tau_e^t = \tau_{e'}^t \equiv \tau^t.$$

This assumption ensures that τ_e^t and τ_e^s do not depend on the particular edge e . In practice, under Assumption 2, this condition is satisfied in a broad class of models where there is a certain separability between the node-level shocks U_i and the edge-level shocks V_e .

A straightforward case is the additive separable model:

$$\epsilon_e = U_{s(e)} + U_{t(e)} + V_e,$$

which implies that $\tau^s(u) = \tau^t(u) = u$ for any $u \in \mathcal{U}$.

We can also allow for interactions between U_i and V_e , provided V_e is not nonlinearly transformed. For instance, if we assume that $\mathbb{E}[V_e] = 1$ and consider the model

$$\epsilon_e = (U_{s(e)} + U_{t(e)}) \times (1 - V_e) + V_e^2 - v_e^2,$$

then it follows that $\tau^s(u) = \tau^t(u) = u$ for any $u \in \mathcal{U}$.

In contrast, Assumption 3 is violated when V_e is nonlinearly transformed in the interaction. For example, if we have

$$\epsilon_e = (U_{s(e)} + U_{t(e)})V_e^2$$

then

$$\mathbb{E}[\epsilon_e | U_{s(e)} = u] = \mathbb{E}[\epsilon_e | U_{t(e)} = u] = v_e^2,$$

even though $\mathbb{E}[\epsilon_e] = 0$ when $\mathbb{E}[U_i] = 0$.

Notice that we can decompose the error term ϵ_e as follows:

$$\epsilon_e = \tau^s(U_{s(e)}) + \tau^t(U_{t(e)}) + \epsilon_e - \tau^s(U_{s(e)}) - \tau^t(U_{t(e)}).$$

This decomposition is useful because it expresses the error term ϵ_e as a sum of two orthogonal components: the first component is the sum of the node-level shocks $\tau^s(U_{s(e)})$ and $\tau^t(U_{t(e)})$, while the second component is the residual error term $\epsilon_e - \tau^s(U_{s(e)}) - \tau^t(U_{t(e)})$.

In matrix form, we can write this decomposition as:

$$\boldsymbol{\epsilon} = \mathbf{F}^s \boldsymbol{\tau}^s + \mathbf{F}^t \boldsymbol{\tau}^t + (\boldsymbol{\epsilon} - \mathbf{F}^s \boldsymbol{\tau}^s - \mathbf{F}^t \boldsymbol{\tau}^t),$$

where \mathbf{F}^s and \mathbf{F}^t are the $m \times n$ matrices defined as:

$$\mathbf{F}_{e,i}^s = \begin{cases} 1 & \text{if } s(e) = i \\ 0 & \text{otherwise} \end{cases}, \quad \mathbf{F}_{e,i}^t = \begin{cases} 1 & \text{if } t(e) = i \\ 0 & \text{otherwise} \end{cases},$$

and $\boldsymbol{\tau}^s = (\tau^s(U_i))_{i \in V}$ and $\boldsymbol{\tau}^t = (\tau^t(U_i))_{i \in V}$ are the n -dimensional vectors of node-level origin and destination shocks, respectively.

Now we can decompose $\hat{\boldsymbol{\alpha}}$ as follows:³

$$\hat{\boldsymbol{\alpha}} = \mathbf{L}^* \mathbf{B}' \mathbf{y} = \boldsymbol{\alpha} + \mathbf{L}^* \mathbf{B}' (\mathbf{F}^s \boldsymbol{\tau}^s + \mathbf{F}^t \boldsymbol{\tau}^t) + \mathbf{L}^* \mathbf{B}' (\boldsymbol{\epsilon} - \mathbf{F}^s \boldsymbol{\tau}^s - \mathbf{F}^t \boldsymbol{\tau}^t). \quad (2.6)$$

The first term in (2.6) represents the true fixed effects, the second term is the bias term stemming from the dependence structure, and the third term is the remaining noise term.

2.5. Variance Bound. Since the second term and third term in (2.6) are orthogonal to each other, i.e., uncorrelated element-wise, the variance of $\hat{\boldsymbol{\alpha}}$ is given by:

$$\text{Var}(\hat{\boldsymbol{\alpha}}) = \mathbf{L}^* \mathbf{B}' \boldsymbol{\Omega}_1 \mathbf{B} \mathbf{L}^* + \mathbf{L}^* \mathbf{B}' \boldsymbol{\Omega}_2 \mathbf{B} \mathbf{L}^*, \quad (2.7)$$

where

$$\boldsymbol{\Omega}_1 = \mathbb{E}[\tau^s(U_i)^2] \mathbf{F}^s (\mathbf{F}^s)' + \mathbb{E}[\tau^s(U_i) \tau^t(U_i)] (\mathbf{F}^s (\mathbf{F}^t)' + \mathbf{F}^t (\mathbf{F}^s)') + \mathbb{E}[\tau^t(U_i)^2] \mathbf{F}^t (\mathbf{F}^t)',$$

$$\boldsymbol{\Omega}_2 = \text{Var}(\boldsymbol{\epsilon} - \mathbf{F}^s \boldsymbol{\tau}^s - \mathbf{F}^t \boldsymbol{\tau}^t).$$

The first term in (2.7) captures the variance of the bias term, while the second term captures the variance of the remaining noise term. Up to some constant, the first term is bounded by

$$\mathbf{L}^* \mathbf{B}' \boldsymbol{\Omega}_1 \mathbf{B} \mathbf{L}^* \preceq \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}' \mathbf{B} \mathbf{L}^*,$$

³Note that $\mathbf{L}^* \mathbf{L} \boldsymbol{\alpha} = \boldsymbol{\alpha}$ as $\mathbf{L}^* \mathbf{L}$ works as a projection matrix on the null space of \mathbf{d} and $\boldsymbol{\alpha}$ is in the null space of \mathbf{d} : $\mathbf{d}' \boldsymbol{\alpha} = 0$ under Assumption 1

where $\mathbf{F} = \text{abs}(\mathbf{B})$ is the signless incidence matrix, and \lesssim indicates that the right-hand side multiplied by some constant upper bounds the left-hand side in the positive semidefinite ordering. The matrix $\mathbf{\Omega}_2$ in the second term is block-diagonal:

$$\mathbf{\Omega}_2 = \begin{pmatrix} \ddots & & \\ & \mathbf{\Omega}_{(i,j)} & \\ & & \ddots \end{pmatrix}, \quad \mathbf{\Omega}_{(i,j)} = \begin{pmatrix} \sigma_{e_k}^2 & \sigma_{e_k, e_l} & \cdots \\ \sigma_{e_k, e_l} & \sigma_{e_l}^2 & \cdots \\ \vdots & \vdots & \ddots \end{pmatrix} - (\mathbb{E}[\tau^s(U_i)^2 + \tau^t(U_i)^2]) \mathbf{u}_{|E_{(i,j)}|} \mathbf{u}_{|E_{(i,j)}|}' ,$$

for each (i, j) such that $E_{(i,j)} \neq \emptyset$. If the number of multiple edges between each pair of nodes is bounded, then the second term in (2.7) is also bounded by

$$\mathbf{L}^* \mathbf{B}' \mathbf{\Omega}_2 \mathbf{B} \mathbf{L}^* \lesssim \mathbf{L}^* \mathbf{B}' \mathbf{B} \mathbf{L}^* = \mathbf{L}^*.$$

Hence, the upper bound on the variance of the fixed-effect estimator is proportional to

$$\text{Var}(\hat{\alpha}) \lesssim \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}' \mathbf{B} \mathbf{L}^* + \mathbf{L}^*.$$

The second term is proportional to the variance in the absence of the dependence structure and is analyzed in Jochmans and Weidner (2019): it is inversely proportional to the connectivity of the graph. On the other hand, the additional variance in the first term is proportional to the density of the line graph represented by

$$\mathbf{H} \equiv \mathbf{F} \mathbf{F}' - 2\mathbf{I}_m,$$

which is the adjacency matrix of the line graph derived from \mathcal{G} : \mathbf{H} has a non-zero entry if two edges share a common node. Thus, the decomposition (2.7) highlights a fundamental trade-off: while a well-connected graph provides sufficient information for accurately estimating node-level effects, the strong dependence resulting from dense connections hampers consistent estimation.

To measure the connectivity, as in Jochmans and Weidner (2019), let $\lambda_{2,L}$ be the second smallest eigenvalue of the normalized Laplacian

$$\mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}.$$

Additionally, for each $i \in V$, define

$$h_i = \left(\frac{1}{d_i} \sum_{j \neq i} \frac{\mathbf{A}_{i,j}^2}{d_j} \right)^{-1},$$

which is a harmonic mean of the degrees of the neighbors of node i . These two objects reflect the graph's connectivity: $\lambda_{2,L}$ captures the global connectivity, while h_i characterizes the local connectivity around node i .

For the dependence structure, let $\lambda_{n,F}$ be the largest eigenvalue of

$$\mathbf{F}\mathbf{F}' = \mathbf{H} + 2\mathbf{I}_m,$$

which quantifies the density of the graph by summarizing how edges are connected to each other through common nodes.

Finally, let $\sigma_\tau^2 \equiv \max_{o \in \{s,t\}} \mathbb{E}[\tau^o(U_i)^2]$ and

$$\tilde{\sigma}_n^2 \equiv \max_{i,j} \mathbf{A}_{i,j} \max_{e \in E(i,j)} (\sigma_e^2 - \mathbb{E}[\tau^s(U_i)^2 + \tau^t(U_i)^2]),$$

which essentially bounds the largest eigenvalue of $\mathbf{\Omega}_2$ in the second term of (2.7).

The following proposition provides a worst-case variance bound for the fixed-effect estimator $\hat{\alpha}_i$:

Proposition 1. *Under Assumptions 1-3, the variance of each fixed-effect estimator $\hat{\alpha}_i$ is bounded as follows:*

$$\text{Var}(\hat{\alpha}_i) \leq (\tilde{\sigma}_n^2 + \lambda_{n,F}\sigma_\tau^2) \times \left\{ \frac{1}{d_i} \left(1 + \frac{1}{\lambda_{2,L}h_i} \right) - \frac{2}{m} \right\}.$$

Proposition 1 highlights that the worst-case variance of the fixed-effect estimator depends critically on the interplay between graph connectivity (as captured by $\lambda_{2,L}$) and the density of the line graph (as measured by $\lambda_{n,F}$). In particular, when the node-level shocks U_i are degenerate (i.e., $\mathbb{E}[\tau^2(U_i)] = 0$), this bound recovers the variance bound found in Jochmans and Weidner (2019) if we further assume that V_e is independently and identically distributed across edges.

It is well known in the spectral graph literature that (Cvetković et al., 2007)

$$2 \min_{k \in V} d_k \leq \lambda_{n,F} \leq 2 \max_{k \in V} d_k,$$

Thus, the upper bound in Proposition 1 is of order

$$O(\lambda_{n,F}/d_i + \lambda_{n,F}/(d_i\lambda_{2,L}h_i)),$$

which is $O(1)$ if $d_i \propto \lambda_{n,F}$ and $\lambda_{2,L}h_i$ is bounded away from zero. This implies that even when the graph is well-connected, the consistency of the fixed-effect estimator is not guaranteed in the worst-case scenario.

3. ASYMPTOTIC THEORY

The variance decomposition in (2.7) and the finite-sample variance bound in Proposition 1 help clarify the behavior of the fixed-effect estimator under the dependence structure. Notably, the estimator can be inconsistent even when the graph is well-connected—contrasting with the consistency results under independence (see Jochmans and Weidner, 2019). However, the finite-sample theory above provides only a worst-case bound and does not sharply characterize the estimator’s behavior. Therefore, in this section, we develop a first-order approximation for the fixed-effect estimator and propose a new inference method for fixed effects under an asymptotic framework, where the sequence of graphs $\mathcal{G}_1, \mathcal{G}_2, \dots$ grows both locally around each node of interest and globally in terms of the number of nodes and edges.

3.1. First-Order Approximation. We begin with the following first-order approximation for the fixed-effect estimator. The idea is to decompose the estimator as follows:

$$\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha} = \mathbf{D}^{-1}\mathbf{B}'\boldsymbol{\epsilon} + \underbrace{\mathbf{D}^{-1}\mathbf{A}(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha})}_{\equiv \mathbf{r}}.$$

This decomposition is algebraically derived by leveraging the restriction $\mathbf{d}'\boldsymbol{\alpha} = 0$.⁴ The remainder term \mathbf{r} aggregates the deviation of the fixed-effect estimator from the true effects, and it will be shown to be negligible, depending on both the connectivity of the graph and the dependence structure as seen in the variance bound in Proposition 1.

Element-wise, the first term of the decomposition for node i is given by

$$\frac{1}{d_i} \sum_{e \in E_i} \mathbf{B}_{e,i} \epsilon_e.$$

Thus, each fixed-effect estimator $\hat{\alpha}_i$ is locally driven by the average of the error terms ϵ_{e_i} over the set E_i of edges incident to node i .

We can further decompose the estimator as follows:

$$\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha} = \mathbf{D}^{-1}\mathbf{B}'(\mathbf{F}^s\boldsymbol{\tau}^s + \mathbf{F}^t\boldsymbol{\tau}^t) + \mathbf{D}^{-1}\mathbf{B}'(\boldsymbol{\epsilon} - \mathbf{F}^s\boldsymbol{\tau}^s - \mathbf{F}^t\boldsymbol{\tau}^t) + \mathbf{r}. \quad (3.1)$$

⁴See the proof for Theorem 4 in Jochmans and Weidner (2019)

The first term on the right-hand side of (3.1) can be expressed as

$$\left(\frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i) + \frac{1}{d_i} \sum_{e \in E_i^s} \tau^t(U_{t(e)}) - \frac{1}{d_i} \sum_{e \in E_i^t} \tau^s(U_{s(e)}) \right)_{i \in V}.$$

Note that the term

$$\frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i)$$

is generally persistent as d_i increases as at least one of d_i^t/d_i or d_i^s/d_i converges to a non-zero constant. On the other hand, since $\tau^s(U_j)$ and $\tau^t(U_j)$ are mean-zero random variables for any $j \in V$, the averages

$$\frac{1}{d_i} \sum_{e \in E_i^s} \tau^t(U_{t(e)}), \quad \frac{1}{d_i} \sum_{e \in E_i^t} \tau^s(U_{s(e)})$$

are of smaller order than the persistent term. Furthermore, the second and third terms in (3.1) are also of smaller order than the persistent term if the graph is well-connected. The following assumption ensures that the remaining terms are negligible and simplifies the asymptotic analysis:

Assumption 4. *The number of edges between each pair of nodes is absolutely bounded, i.e., there exists an absolute constant $C > 0$ such that $\max_{i,j \in V} |E_{(i,j)}| \leq C$ for all $n \in \mathbb{N}$.*

Assumption 4 can be relaxed to allow for a growing number of edges between each pair of nodes, provided that the growth is sufficiently slow. For example, if node i is of interest and $d_i \rightarrow \infty$, we can allow for $\max_{i,j \in V} |E_{(i,j)}| \rightarrow \infty$ as long as the growth rate is slower than d_i . Also, we can force this assumption to hold by randomly removing a subset of multiple edges between each pair of nodes such that the number of remaining edges is bounded.

The following result summarizes the above discussion and provides a first-order approximation for the fixed-effect estimator $\hat{\alpha}$:

Theorem 1. *Under Assumptions 1-4, for each $i \in V$, we have*

$$\hat{\alpha}_i - \alpha_i = \frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i) + O_p \left(\sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}} \right),$$

as $d_i \rightarrow \infty$.

Theorem 1 shows that the fixed-effect estimator is approximated by the true fixed effect plus a mean-zero noise term, expressed as a weighted difference in the origin and destination effects, $\tau^s(U_i)$ and $\tau^t(U_i)$. Thus, the estimator is unbiased but inconsistent. The inconsistency arises because the fixed-effect estimator is, approximately, a local average of the error terms in which the node-level random effects are not fully averaged out. That is why the noise term is closely related to the proportions of edges where node i appears as an origin (d_i^s/d_i) versus as a destination (d_i^t/d_i). Since the noise does not vanish, this result is consistent with the intuition that the fixed-effect estimator is potentially inconsistent even in a well-connected graph due to the dependence structure.

Also note that Theorem 1 generalizes Theorem 4 in Jochmans and Weidner (2019), which establishes that the fixed-effect estimator is consistent for α_i under a weak dependence structure. In Jochmans and Weidner (2019), this weak dependence structure is characterized by requiring that the largest eigenvalue of $\mathbb{E}[\epsilon\epsilon']$ remains bounded by a positive absolute constant. In our framework, however, the dependence structure is sufficiently strong that the largest eigenvalue of $\mathbb{E}[\epsilon\epsilon']$ diverges with d_i . This divergence leads directly to the inconsistency of the fixed-effect estimator in our setting.

The remainder term is negligible if the graph is well-connected, such that $\lambda_{2,L}$ does not decay too quickly and $\lambda_{2,L}h_i \rightarrow \infty$, and the graph itself grows sufficiently slowly, ensuring that $\lambda_{n,F}/d_i$ does not grow too rapidly as d_i increases. In that case, we have

$$\hat{\alpha}_i - \alpha_i = \frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i) + o_p(1),$$

so that the fixed-effect estimator is consistent for

$$\alpha_i + \frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i).$$

A sufficient condition for this is that

$$\lambda_{n,F}/(d_i\lambda_{2,L}h_i) \rightarrow 0.$$

This condition is satisfied in the following examples:⁵

Example 3. Suppose that \mathcal{G} is a complete graph. Then, $d_i = n - 1$, $h_i = n - 1$ for each $i \in V$ and $\lambda_{n,F} = 2(n - 1)$, $\lambda_{2,L} = 1$. Thus, $\lambda_{n,F}/(d_i\lambda_{2,L}h_i) = 1/(n - 1) \rightarrow 0$ as $n \rightarrow \infty$.

⁵The dependence measure $\lambda_{n,F}$ can be conservative in highly heterogeneous graphs where a few nodes have a large number of links. Improving this bound will be an important direction for future work.

Example 4. Suppose that \mathcal{G} is an Erdős–Rényi random graph with edge probability $p_n = c \log(n)/n$ for some constant $c > 1$. Then, $d_i/\log(n) \rightarrow c$, $h_i/\log(n) \rightarrow c$ for each $i \in V$ and $\lambda_{n,F}/\log(n) \rightarrow c$, $\lambda_{2,L} \rightarrow 1$ almost surely as $n \rightarrow \infty$. Thus, $\lambda_{n,F}/(d_i \lambda_{2,L} h_i) \sim 1/\log(n) \rightarrow 0$ almost surely as $n \rightarrow \infty$.

A counterexample that violates the condition $\lambda_{n,F}/(d_i \lambda_{2,L} h_i) \rightarrow 0$ is given below:

Example 5. Suppose that \mathcal{G} is a stochastic-block random graph composed of two blocks of size $n/2$. Let the intra-block edge probability be $p_n = p_0 \log(n)/n$ for some constant $p_0 > 2$, and let the inter-block edge probability be $q_n = q_0/n$ for some constant $q_0 < p_0$. Then, for each $i \in V$, $d_i/\log(n) \rightarrow p_0/2$ and $h_i/\log(n) \rightarrow p_0/2$, and $\lambda_{n,F}/\log(n) \rightarrow p_0/2$, $\lambda_{2,L}^{-1} \times \log(n) = O(1)$ almost surely as $n \rightarrow \infty$. Thus, we have $\lambda_{n,F}/(d_i \lambda_{2,L} h_i) = O(1)$ almost surely as $n \rightarrow \infty$.⁶

In Example 5, the graph is not well-connected: the connections between the two blocks are sparse, making it easier for the graph to break into two disconnected components as n increases. In contrast, within each block the graph is well-connected, and the dependence structure remains non-negligible. Consequently, the first-order approximation in Theorem 1 may fail under such conditions. Indeed, without the dependence structure, $\lambda_{n,F}$ would be replaced by 1, so even with the stochastic-block structure, the fixed-effect estimator would be consistent for α_i as $1/(d_i \lambda_{2,L} h_i) \rightarrow 0$ as $d_i \rightarrow \infty$.

3.2. Inference.

The first-order approximation in Theorem 1 suggests a potential approach for performing inference on the fixed effect α_i by estimating the distribution of

$$\frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i).$$

However, the challenge is that we must either directly estimate the differences across nodes or estimate each $\tau^s(U_i)$ and $\tau^t(U_i)$ separately. This is in general not feasible without additional assumptions as α_i is unknown and only the sum $\tau^t(U_i) + \tau^s(U_i)$ can be consistently estimated.

⁶This result is based on Deng et al. (2021), who characterized the asymptotic behavior of $\lambda_{2,L}$ in the stochastic-block random graphs

To illustrate this point, consider the residual $\hat{\epsilon}$ defined as

$$\hat{\epsilon} \equiv \mathbf{y} - \mathbf{B}\hat{\alpha} = \mathbf{M}_\mathbf{B}\mathbf{y},$$

where $\mathbf{M}_\mathbf{B} = \mathbf{I}_m - \mathbf{B}\mathbf{L}^*\mathbf{B}'$ projects onto the orthogonal complement of the column space of \mathbf{B} . For each $i \in V$, let c_i be defined as

$$c_i = \frac{1}{d_i} \mathbf{f}_i' \mathbf{M}_\mathbf{B} \mathbf{f}_i,$$

where \mathbf{f}_i is the i -th column of \mathbf{F} . Note that $c_i \in [0, 1]$. Heuristically, c_i is a measure of balance between the inflow and outflow of edges incident to node i : if $d_i^t - d_i^s$ is close to zero, then c_i is close to one, while if $|d_i^t - d_i^s|$ is close to d_i , then c_i is close to zero. In fact, we can show that

$$c_i \sim 1 - \left(\frac{d_i^t - d_i^s}{d_i} \right)^2,$$

when the graph is well-connected with large $\lambda_{2,L}h_i$.

By locally averaging the residuals, we can obtain the following:

Proposition 2. *Under Assumptions 1-4, for each $i \in V$, we have*

$$\frac{1}{d_i} \sum_{e \in E_i} \hat{\epsilon}_e = c_i \times \frac{\tau^t(U_i) + \tau^s(U_i)}{2} + O_p \left(\sqrt{\frac{1}{d_i}} \right),$$

as $d_i \rightarrow \infty$.

Proposition 2 shows that if node i has sufficient balance between its inflow and outflow of edges, the local average of the residuals $\hat{\epsilon}_e$ for edges incident to node i contains information about the average $(\tau^t(U_i) + \tau^s(U_i))/2$. Specifically, if $c_i > 0$, we have

$$\hat{\tau}_i \equiv \frac{1}{d_i c_i} \sum_{e \in E_i} \hat{\epsilon}_e = \frac{\tau^t(U_i) + \tau^s(U_i)}{2} + O_p \left(\sqrt{\frac{1}{c_i^2 d_i}} \right). \quad (3.2)$$

Thus, as long as the balance measure c_i does not converge to zero quickly, we can consistently estimate the average $(\tau^t(U_i) + \tau^s(U_i))/2$. An extreme case would be when $c_i = 0$, which occurs when $d_i^t = 0$ or $d_i^s = 0$. In this case, the local average of the residuals is exactly zero and uninformative about the average $(\tau^t(U_i) + \tau^s(U_i))/2$.

Remark 3. *The convergence rate of $\hat{\tau}_i$ in (3.2) is notable because it does not depend on $\lambda_{n,F}$ or $\lambda_{2,L}h_i$. In other words, the convergence rate is determined solely by the balance*

and the number of edges incident to node i , regardless of the overall connectivity of the graph or the strength of the dependence structure. This implies that even if the graph is not well-connected and the dependence structure is strong—so the first-order approximation in Theorem 1 may not hold, as in Example 5—we can still consistently estimate the average $(\tau^t(U_i) + \tau^s(U_i))/2$ as long as the edges incident to node i are sufficiently balanced.

As argued above, the average $(\tau^t(U_i) + \tau^s(U_i))/2$ is not sufficient for performing inference on the fixed effect α_i because essentially we have three unknowns $(\alpha_i, \tau^t(U_i), \tau^s(U_i))$ for two equations:

$$\begin{aligned}\hat{\alpha}_i &\sim \alpha_i + \frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i) \\ \hat{\tau}_i &\sim \frac{\tau^t(U_i) + \tau^s(U_i)}{2}.\end{aligned}$$

To facilitate inference, we impose the following additional structure on the $\tau^s(\cdot)$ and $\tau^t(\cdot)$ functions:

Assumption 5. *For any $u \in \mathcal{U}$, we have*

$$\tau^s(u) = \tau^t(u) \equiv \tau(u).$$

Assumption 5 is satisfied in a broad class of models where the dependence structure is symmetric in $U_{s(e)}$ and $U_{t(e)}$, such as the additive separable model or the interactive model discussed above following Assumption 3. Importantly, this assumption is weaker than requiring symmetry of f in (2.4) with respect to the origin and destination shocks, i.e., $f(u, u', v) = f(u', u, v)$ for all u, u', v in their supports. The following example illustrates that Assumption 5 can hold even when f is not symmetric in $U_{s(e)}$ and $U_{t(e)}$: If

$$\epsilon_e = f(U_{s(e)}, U_{t(e)}, V_e) = U_{s(e)}^2 U_{t(e)} - U_{s(e)}^2 + U_{s(e)} - 1,$$

with $\mathbb{E}[U_i] = \mathbb{E}[U_i^2] = 1$, then

$$\tau^s(u) = \tau^t(u) = u - 1$$

for all $u \in \mathcal{U}$, even though $(U_{s(e)}, U_{t(e)}) = (2, 0)$ and $(0, 2)$ yield $\epsilon_e = -5$ and 1 , respectively, which breaks the symmetry of f . Thus, we can allow for asymmetric responses of the error term to the origin and destination shocks, while still satisfying Assumption 5.

When Assumption 5 is violated, ϵ_e is necessarily asymmetric in $U_{s(e)}$ and $U_{t(e)}$. For example, consider the model where

$$\epsilon_e = U_{s(e)}U_{t(e)} + U_{t(e)} - 2U_{s(e)},$$

and suppose $\mathbb{E}[U_i] = 1$. In this case, one can verify that

$$\tau^s(u) = -u, \quad \tau^t(u) = 2u,$$

so that $\tau^s(u) \neq \tau^t(u)$ for any $u \in \mathcal{U}$ except for $u = 0$.

Remark 4. *Instead of imposing Assumption 5, we may consider alternative normalization on the $\tau^s(\cdot)$ and $\tau^t(\cdot)$ functions, such as $\tau^t(\cdot) = 0$ or $\tau^s(\cdot) = 0$. In such cases, the non-zero function can be denoted by τ and inference can proceed similarly as below. More generally, we can consider the case where there is a linear relationship between $\tau^t(\cdot)$ and $\tau^s(\cdot)$, allowing us to set $\tau(\cdot) = \tau^t(\cdot)$ and $\tau^s = c\tau(\cdot)$ or $\tau(\cdot) = \tau^s(\cdot)$ and $\tau^t = c\tau(\cdot)$ for some constant $c > 0$. See Appendix C for inference under such a setting.*

Under Assumption 5, Proposition 2 implies that we can consistently estimate $\tau(U_i)$ as follows:

$$\hat{\tau}_i \sim \tau(U_i),$$

if $c_i^2 d_i \rightarrow \infty$ as $d_i \rightarrow \infty$. Moreover, under Assumption 5, Theorem 1 implies that

$$\hat{\alpha}_i - \alpha_i \sim \frac{d_i^t - d_i^s}{d_i} \tau(U_i),$$

provided that $\lambda_{n,F}/(d_i \lambda_{2,L} h_i) \rightarrow 0$ as $d_i \rightarrow \infty$. Thus, we can leverage the consistency of $\hat{\tau}_i$ to estimate the distribution of $\tau(U_i)$ and to perform inference on fixed effects.

Remark 5. *Since $\hat{\tau}_i$ is a consistent estimator for $\tau(U_i)$, one might consider correcting the fixed-effect estimator as follows:*

$$\hat{\alpha}_i - \frac{d_i^t - d_i^s}{d_i} \hat{\tau}_i \sim \alpha_i.$$

For point estimation, this correction would yield a consistent estimator for α_i if the remainder term in Theorem 1 is negligible and i has sufficient balance. Relatedly, Appendix B discusses an alternative estimator for α_i that directly incorporates this correction. For

inference, however, we do not pursue this approach for two reasons. First, this correction is not feasible for every node $i \in V$ when $c_i = 0$ or close to zero, as $\tau(U_i)$ cannot be consistently estimated in these cases. This issue is particularly relevant when nodes of interest, or a non-negligible number of nodes, have c_i close to zero, as observed in the empirical application in Section 6. Second, the correction does not guarantee asymptotic normality, since the remainder term \mathbf{r} in (3.1) can be dominant with order $O_p(\sqrt{\lambda_{n,F}/(d_i \lambda_{2,L} h_i)})$, and the central limit theorem does not directly apply to this term. Moreover, estimating the standard error of the corrected estimator is not straightforward without further assumptions on the distribution of $(V_e)_{e \in E}$. These considerations motivate the development of a new inference method based on the empirical distribution of $\hat{\tau}_i$ for $i \in \mathcal{C}_n$, as defined below.

For this purpose, we impose the following assumption on the distribution of $\tau(U_i)$:

Assumption 6. $\tau(U_i)$ has a continuous cumulative distribution function.

Assumption 6 ensures that the distribution of $\tau(U_i)$ is well-behaved. This excludes degenerate cases and guarantees that the dependence structure remains relevant as the graph size n increases. Since $\tau(U_i)$ is mean-zero, degeneracy would imply $\tau(U_i) = 0$ almost surely for all $i \in V$, in which case the standard inference results from Jochmans and Weidner (2019) would apply.

We define the set of nodes useful for estimating $\boldsymbol{\tau}$ as $\mathcal{C}_n \equiv \{i \in V : c_i > c\}$, where $c > 0$ is an absolute constant. Let F_τ be the cumulative distribution function of $\tau(U_i)$ and $\hat{F}_{n,\tau}$ be the empirical distribution function of $\hat{\tau}_i$ for $i \in \mathcal{C}_n$, defined as

$$\hat{F}_{n,\tau}(t) \equiv \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{I}(\hat{\tau}_i \leq t),$$

where $\mathbb{I}(\cdot)$ is the indicator function. Also, define the following object to globally control the estimation errors for $|\mathcal{C}_n|$ nodes:

$$\eta_n \equiv \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \frac{1}{d_i},$$

which tends to zero as $n \rightarrow \infty$ provided that nodes in \mathcal{C}_n have growing d_i on average.

Then, we have the following result:

Theorem 2. *Under Assumptions 1-6, if $|\mathcal{C}_n| \rightarrow \infty$ and $\eta_n \rightarrow 0$ as $n \rightarrow \infty$, we have*

$$\sup_{t \in \mathbb{R}} |\hat{F}_{n,\tau}(t) - F_\tau(t)| \rightarrow_p 0,$$

as $n \rightarrow \infty$.

The additional conditions $|\mathcal{C}_n| \rightarrow \infty$ and $\eta_n \rightarrow 0$ ensure that there are enough informative nodes to reliably estimate the distribution of $\tau(U_i)$. These conditions are satisfied in many graphs with growing degrees, where nodes have sufficient inflow and outflow of edges. For example, in Examples 3-5, for each edge, if source and target nodes are assigned independently at random with probability $q \in (0, 1)$ and $1 - q$, respectively, then the conditions are likely satisfied when q is an absolute constant.

Theorem 2 is useful for conducting inference on fixed effects in a manner similar to that of Conley and Taber (2011), who developed a non-standard inference under small treatment group asymptotics. Noting that

$$\left(\frac{d_i^t - d_i^s}{d_i} \right)^{-1} (\hat{\alpha}_i - \alpha_i) \sim \tau(U_i),$$

for $i \in V$ such that $|d_i^t - d_i^s| > d_i c$ for some absolute constant $c > 0$, we can construct a confidence interval for α_i by inverting the test based on $\hat{F}_{n,\tau}$. Specifically, for a given confidence level $1 - \alpha$, let $\hat{c}_{\alpha/2}$ and $\hat{c}_{1-\alpha/2}$ be the $(\alpha/2)$ -th and $(1 - \alpha/2)$ -th quantiles of $\hat{F}_{n,\tau}$, respectively. Then, the $1 - \alpha$ confidence interval for α_i is given by

$$CI_{i,1-\alpha} \equiv \begin{cases} \left[\hat{\alpha}_i - \left(\frac{d_i^t - d_i^s}{d_i} \right) \hat{c}_{1-\alpha/2}, \hat{\alpha}_i - \left(\frac{d_i^t - d_i^s}{d_i} \right) \hat{c}_{\alpha/2} \right] & \text{if } d_i^t - d_i^s > 0; \\ \left[\hat{\alpha}_i - \left(\frac{d_i^t - d_i^s}{d_i} \right) \hat{c}_{\alpha/2}, \hat{\alpha}_i - \left(\frac{d_i^t - d_i^s}{d_i} \right) \hat{c}_{1-\alpha/2} \right] & \text{if } d_i^t - d_i^s < 0. \end{cases} \quad (3.3)$$

Another useful application of Theorem 2 is joint hypothesis testing for fixed effects α . For example, consider testing the null hypothesis that a subset V_0 of nodes share the same fixed effect:

$$H_0 : \alpha_i = \alpha_j \text{ for all } i, j \in V_0 \quad \text{versus} \quad H_1 : \alpha_i \neq \alpha_j \text{ for some } i, j \in V_0.$$

We can test this hypothesis using the following test statistic:

$$T = \hat{\alpha}'_{V_0} \mathbf{M}_{n_0} \hat{\alpha}_{V_0} \quad (3.4)$$

where $\alpha_{V_0} = (\alpha_i)_{i \in V_0}$, $n_0 = |V_0|$, and \mathbf{M}_{n_0} is the demeaning matrix of size n_0 . Under the null hypothesis,

$$\alpha'_{V_0} \mathbf{M}_{n_0} \alpha_{V_0} = 0,$$

and the distribution of T can be approximated by simulating $\hat{\alpha}_{V_0}$ as $((d_i^t - d_i^s)/d_i \times \tau_i^{(m)})_{i \in V_0}$ for M repetitions, where each $\tau_i^{(m)}$ is drawn from the empirical distribution $\hat{F}_{n,\tau}$. Let $\hat{c}_{1-\alpha}^T$ denote the $(1 - \alpha)$ quantile of the simulated distribution of T under the null. The test rejects H_0 if $T > \hat{c}_{1-\alpha}^T$.

More generally, we can test linear hypotheses of the form $H_0 : R\alpha = 0$ for some known matrix R using a similar approach and constructing the test statistic as $g_n(\hat{\alpha}_{V_0})$ for some known function $g_n(\cdot)$ such that the null distribution can be simulated from $\hat{F}_{n,\tau}$.

The following result establishes the asymptotic validity of the confidence interval in (3.3) and the test based on the statistic T in (3.4):⁷

Proposition 3. *Suppose that the conditions of Theorem 2 hold. Suppose also that F_τ is strictly increasing around quantiles of interest. For each $i \in V$ such that $(d_i^t - d_i^s)/d_i = O(1)$ and $\lambda_{n,F}/(d_i \lambda_{2,L} h_i) \rightarrow 0$ as $n \rightarrow \infty$, we have*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\alpha_i \notin CI_{i,1-\alpha}) = \alpha.$$

Also, if for each $i \in V_0$, $(d_i^t - d_i^s)/d_i \rightarrow c_i^{st}$, $\min_{i \in V_0} |c_i^{st}| > 0$, $\max_{i \in V_0} \lambda_{n,F}/(d_i \lambda_{2,L} h_i) \rightarrow 0$ as $n \rightarrow \infty$, and n_0 is fixed, then

$$\lim_{n, M \rightarrow \infty} \mathbb{P}(T > \hat{c}_{1-\alpha}^T) = \alpha,$$

under the null hypothesis that $\alpha_i = \alpha_j$ for all $i, j \in V_0$.

Remark 6. *The proposed inference methods rule out the case where nodes of interest have almost perfect balance between their inflow and outflow of edges, i.e., $d_i^t \sim d_i^s$. In such cases, we have*

$$\hat{\alpha}_i - \alpha_i \sim 0$$

and little information on $\tau(U_i)$ remains in $\hat{\alpha}_i$. The confidence interval in (3.3) and the distribution of the test statistic T will be degenerate, leading to unreliable inference. If

⁷Our result here only establishes point-wise asymptotic validity. Establishing uniform asymptotic validity over a class of DGPs and graph sequences is an important direction for future work.

all nodes of interest have such balance, we can at least acknowledge that the estimates for these nodes are consistent without correcting as done in Remark 5.

4. VARIANCE COMPONENTS ESTIMATION

In the previous sections, we have focused on the estimation and inference of fixed effects α . However, in empirical applications, researchers are often interested in distributional properties of the fixed effects. For example, Card et al. (2013) estimate the sample variance of firm fixed effects in the AKM model to assess the contribution of workplace heterogeneity to rising wage inequality. In this section, we focus on estimating the sample variance of the fixed effects α and address the bias in the sample variance estimator that arises from the dependence structure.

4.1. Estimation.

Our parameter of interest is the sample variance of α , given by

$$\begin{aligned} V_\alpha &\equiv \frac{1}{n} \sum_{i \in V} (\alpha_i - \bar{\alpha})^2 \\ &= \frac{\alpha' M_n \alpha}{n}, \end{aligned}$$

where $M_n = \mathbf{I}_n - 1/n \times \boldsymbol{\iota}_n \boldsymbol{\iota}_n'$ is the demeaning matrix, $\boldsymbol{\iota}_n$ is an n -dimensional vector of ones, and $\bar{\alpha} = \boldsymbol{\iota}_n' \alpha / n$ is the average of α . We can estimate V_α by plugging in $\hat{\alpha}$ for α :

$$\hat{V}_\alpha = \frac{\hat{\alpha}' M_n \hat{\alpha}}{n}.$$

It is well known that \hat{V}_α is biased upward due to the estimation error in $\hat{\alpha}$ even in cases where dependence is ruled out (Andrews et al., 2008). To see this, from (2.3), we have

$$\begin{aligned} \mathbb{E}[\hat{V}_\alpha] &= V_\alpha + \frac{\mathbb{E}[\boldsymbol{\epsilon}' \mathbf{B} \mathbf{L}^* M_n \mathbf{L}^* \mathbf{B}' \boldsymbol{\epsilon}]}{n} \\ &= V_\alpha + \frac{\text{tr}(\mathbf{B} \mathbf{L}^* M_n \mathbf{L}^* \mathbf{B}' \mathbb{E}[\boldsymbol{\epsilon} \boldsymbol{\epsilon}'])}{n}, \end{aligned} \tag{4.1}$$

where $\text{tr}(\cdot)$ denotes the trace of a matrix. Thus, the bias in \hat{V}_α is given by the second term in (4.1).

The bias term in (4.1) is nonzero in general and known as “limited mobility bias” in the literature. As an illustration, suppose that $\boldsymbol{\epsilon}$ is independent and identically distributed.

Then, the bias term simplifies to

$$\mathbb{E}[\epsilon_e^2] \times \frac{\text{tr}(\mathbf{M}_n \mathbf{L}^*)}{n},$$

which is inversely proportional to the connectivity of the graph.

4.2. Bias Correction.

To correct the bias in \hat{V}_α , several bias-correction methods have been proposed in the literature. When ϵ is independent and identically distributed, the bias can be estimated by the following formula (Andrews et al., 2008):

$$\hat{\sigma}^2 \times \frac{\text{tr}(\mathbf{M}_n \mathbf{L}^*)}{n},$$

where $\hat{\sigma}^2$ is an estimator consistent for $\mathbb{E}[\epsilon_e^2]$. Kline et al. (2020) extends this bias-correction method to accommodate the case where ϵ is independent but not identically distributed.

In our setting, however, ϵ is not independent, so the above methods are not directly applicable. To address this issue, note that we can decompose the covariance matrix of ϵ as

$$\mathbb{E}[\epsilon \epsilon'] = \mathbf{\Omega}_1 + \mathbf{\Omega}_2,$$

as in (2.7). Thus, under Assumption 5, the bias term in (4.1) can be rewritten as:⁸

$$\mathbb{E}[\tau(U_i)^2] \times \frac{\text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n} + \frac{\text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{\Omega}_2)}{n}. \quad (4.2)$$

Here, the first term in (4.2) is the new bias term due to the node-level dependence structure, while the second term corresponds to the limited mobility bias in the literature. Note that when there are few multiple edges between any given pair of nodes, the second term is proportional to $\text{tr}(\mathbf{M}_n \mathbf{L}^*)$.

Heuristically, we can assess relative severity of the new bias term compared to the second term by comparing

$$\frac{\text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n} \text{ and } \frac{\text{tr}(\mathbf{M}_n \mathbf{L}^*)}{n}.$$

⁸Since the distributions of $\tau^t(U_i)$ and $\tau^s(U_i)$ are not separately identified as shown above, in this section we focus on the symmetric case and maintain Assumption 5.

These quantities can be computed from the graph structure. In a well-connected graph, the latter will be small in comparison to the former, so the bias arising from the dependence structure will dominate. Conversely, in graphs that are not well-connected, both terms will be significant.

Hypothetically, we can estimate the bias term in (4.2) by estimating $\mathbb{E}[\tau(U_i)^2]$ and $\mathbf{\Omega}_2$ separately. In particular, we can estimate $\mathbb{E}[\tau^2(U_i)]$ via the following estimator:

$$\hat{\sigma}_\tau^2 = \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \left(\frac{1}{d_i c_i} \sum_{e \in E_i} \hat{\epsilon}_e \right)^2,$$

whose consistency is anticipated by Proposition 2. However, estimating $\mathbf{\Omega}_2$, without imposing additional restrictions on the dependence structure or heteroskedasticity is challenging as $\hat{\boldsymbol{\alpha}}$ is not, in general, a consistent estimator of $\boldsymbol{\alpha}$, and the correlation structure within each block of $\mathbf{\Omega}_2$ is too flexible to estimate.

Instead of estimating $\mathbf{\Omega}_2$ directly, we propose the following bias-corrected estimator for V_α by correcting only the first term in (4.2):

$$\hat{V}_\alpha^{bc} = \hat{V}_\alpha - \hat{\sigma}_\tau^2 \times \frac{\text{tr}(\mathbf{B}\mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n}.$$

We need the following additional regularity conditions to establish the consistency of \hat{V}_α^{bc} :

Assumption 7. *The following conditions hold:*

- (i) *The fixed effects $\boldsymbol{\alpha}$ are uniformly bounded, i.e., $\sup_n \sup_{i \in V} |\alpha_i| < C$ for some absolute constant $C > 0$.*
- (ii) *The following quantities are $o(1)$ as $n \rightarrow \infty$:*

$$\frac{\sqrt{\lambda_{n,F}}}{n} \sum_{i \in V} \frac{1}{d_i}, \frac{\sqrt{\eta_n} \lambda_{n,F}}{n} \sum_{i \in V} \frac{1}{d_i}, \frac{\lambda_{n,F}}{n} \sum_{i \in V} \frac{1}{d_i \lambda_{2,L} h_i}.$$

Part (i) of Assumption 7 excludes the case in which some nodes have unbounded fixed effects, a relatively mild requirement. Parts (ii) consists of technical conditions that ensure the bias arising from the dependence structure persists (so our bias correction is essential), while the bias from the remainder of the variance–covariance matrix remains relatively small. These conditions hold, for example, in well-connected, homogeneous graphs such as the Erdős–Rényi random graph of Example 4, but fail if the graph is poorly connected or if edges are concentrated among only a few node pairs.

Then, we have the following consistency result for the bias-corrected estimator:

Theorem 3. *Under Assumptions 1-5 and 7, if $|\mathcal{C}_n| \rightarrow \infty$ and $\eta_n \rightarrow 0$ as $n \rightarrow \infty$, we have*

$$\hat{V}_{\alpha}^{bc} - V_{\alpha} \rightarrow_p 0,$$

as $n \rightarrow \infty$.

Note that although Theorem 3 establishes the consistency of the bias-corrected estimator \hat{V}_{α}^{bc} , in finite sample, it is typically still biased upward as we have not corrected the second term in (4.2). Nonetheless, our finding suggests that ignoring the bias arising from the dependence structure can be more severe than ignoring the second term when the graph is well-connected and homogeneous.

To account for the finite-sample bias, we propose the following rule of thumb modification to the bias-corrected estimator:

$$\hat{V}_{\alpha}^{bc,mod} = \hat{V}_{\alpha} - \hat{\sigma}_{\tau}^2 \times \frac{\text{tr}(\mathbf{B}\mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n} - \max\{0, \hat{\sigma}^2 - 2\hat{\sigma}_{\tau}^2\} \times \frac{\text{tr}(\mathbf{M}_n \mathbf{L}^*)}{n},$$

where

$$\hat{\sigma}^2 = \frac{1}{m} \sum_{e \in E} \hat{\epsilon}_e^2.$$

The second term in $\hat{V}_{\alpha}^{bc,mod}$ serves as the bias-correction term for the second term in (4.2) when the errors are assumed to be homoskedastic and the within-block correlation is neglected. If the block correlation induces positive bias, then this modification will be conservative.

Also, note that $\hat{\sigma}^2$ is not consistent for $\mathbb{E}[\epsilon_e^2]$ in general. However, if the errors are homoskedastic, it can be approximated by

$$\hat{\sigma}^2 \sim \mathbb{E}[\tau^2(U_i)] \times \left[\frac{4}{\sum_{i \in V} d_i^t} \sum_{i \in V} \frac{d_i^t d_i^s}{d_i} - 2 \right] + \mathbb{E}[\epsilon_e^2] \leq \mathbb{E}[\epsilon_e^2]$$

so that $\hat{\sigma}^2$ is downward biased relative to $\mathbb{E}[\epsilon_e^2]$, leading to a conservative modification of the bias-corrected estimator.

Remark 7. *The rule-of-thumb modification may result in over-correction, especially when the errors are heteroskedastic, as observed in the simulation results in the next section. This is because the approximation of $\hat{\sigma}^2$ above assumes homoskedasticity. Using a more robust estimator for $\hat{\sigma}^2$, such as the one proposed by Kline et al. (2020), could potentially*

improve the performance of the modified estimator. We leave a detailed investigation of this approach to future work.

We can compare our bias-corrected estimators with those previously proposed in the literature. The key difference is that our estimator corrects for bias by targeting the first term in (4.2), which arises from the node-level dependence structure, and it does so without requiring the independence assumption on the error term ϵ . By contrast, existing bias-corrected estimators, such as those of Andrews et al. (2008) and Kline et al. (2020), correct only for the second term in (4.2) and are consistent if ϵ is independent.⁹

For example, Andrews et al. (2008)’s bias-corrected estimator is given by

$$\hat{V}_{\alpha}^a = \hat{V}_{\alpha} - \hat{\sigma}^2 \times \frac{\text{tr}(\mathbf{M}_n \mathbf{L}^*)}{n}.$$

Thus, if the graph is not well-connected and the contribution of the node-level shocks (i.e., $\mathbb{E}[\tau^2(U_i)]$) is small relative to $\mathbb{E}[\epsilon_e^2]$, this correction may work well as the first term in (4.2) may be negligible.

However, if the node-level dependence structure is strong with significant $\mathbb{E}[\tau^2(U_i)]$ relative to $\mathbb{E}[\epsilon_e^2]$, then the bias stemming from this dependence structure becomes dominant, as shown in Theorem 3. In this scenario, the Andrews et al. (2008)-type bias correction, which ignores this component, will be inconsistent. Similarly, the bias-correction method proposed by Kline et al. (2020), which also hinges on the independence of ϵ or weak dependence, suffers from the same limitation and will generally be inconsistent when the dependence structure is strong.

5. SIMULATION

In this section, we conduct a simulation study to illustrate our inference procedure and the finite-sample properties of the bias-correction method.

5.1. Design.

We first generate an undirected graph $\mathcal{G} = (V, E)$ with $|V| = n$ nodes from the stochastic block model discussed in Example 5. Recall that the stochastic block model is a random graph model where nodes are partitioned into K blocks, and edges are formed between nodes in the same block with probability p_n and between nodes in different blocks with

⁹While Kline et al. (2020) permits weak dependence in their model, their framework does not accommodate the strong dependence structure we consider here.

probability q_n . We set $p_n = 10 \log(n)/n$ and $q_n = 2/(\log(n)n)$ for each n . We then extract the largest connected component from the generated graph.¹⁰ We vary the number of blocks $K = 1, 2$ and the number of nodes $n = 500, 1000, 2500$, and 5000 to compare the performance of our inference procedure under different graph structures.

Given the connected graph $\mathcal{G} = (V, E)$, we generate and fix the true fixed effects $\boldsymbol{\alpha}$ as follows:

$$v_i \sim U[0, 1] \text{ for } i \in V;$$

$$\boldsymbol{\alpha} = \mathbf{v} - (\mathbf{v}'\mathbf{d}/\mathbf{d}'\mathbf{d}) \times \mathbf{d}$$

which ensures that $\mathbf{d}'\boldsymbol{\alpha} = 0$.

We generate the error terms $\boldsymbol{\epsilon}$ according to

$$\epsilon_e = U_{s(e)} + U_{t(e)} + V_e \text{ for } e \in E,$$

where $U_i \sim N(0, 1)$, and

$$V_e \sim N(0, 1 + |\alpha_{s(e)}| + |\alpha_{t(e)}|)$$

independently. Note that this structure satisfies Assumptions 3 and 5 with $\tau(U_i) = U_i$ for each $i \in V$ and allows for heteroskedasticity in V_e .

Using $\boldsymbol{\alpha}$, we randomly assign the inflow and outflow of each edge $e \in E_{(i,j)}$ as follows:

$$t(e) = i \text{ with probability } \frac{|\alpha_i|}{|\alpha_i| + |\alpha_j|},$$

independently for each edge $e \in E_{(i,j)}$. We can interpret this assignment as a process where a node with large effect is more likely to attract the inflow of an edge. Then, with (V, E, s, t) , we can construct the incidence matrix \mathbf{B} and generate the outcome vector according to (2.1). For each iteration, we compute the least-squares estimator $\hat{\boldsymbol{\alpha}}$ and the empirical distribution of $\tau(U_i)$ using Theorem 2, by setting $\mathcal{C}_n = \{i \in V : c_i > 0.2\}$.

Table 1 reports the degree distribution of the generated graphs. For each n , the generated graphs exhibit a balanced degree distribution, avoiding extremes of sparsity and density, regardless of the underlying generating process.

Table 2 reports the global measures of the generated graphs. The connectivity measure $\lambda_{2,L}$ is well bounded away from zero for $K = 1$, while it converges to zero for $K = 2$,

¹⁰In our simulation setting, randomly generated graphs are typically connected, and even if they are not, the largest connected component covers a large proportion of the nodes. Thus, we do not distinguish the original \mathcal{G} from the largest connected component, and we keep denoting it as \mathcal{G} for simplicity.

reflecting that the latter graph is much easier to partition into separate components. In both models, the dependence measure $\lambda_{n,F}$ increases slowly as n increases. Moreover, more than 90% of the nodes are in the set \mathcal{C}_n , suggesting there are enough informative nodes to estimate the distribution of $\tau(U_i)$. The convergence measure η_n approaches zero for both models, indicating that Theorem 2 provides a good approximation of the distribution of $\tau(U_i)$ for large n . The other convergence measure, $H_n \equiv \lambda_{n,F} n^{-1} \sum_{i \in V} d_i^{-1} \lambda_{2,L} h_i^{-1}$, converges to zero for $K = 1$ but not for $K = 2$, reflecting that the stochastic block model with $K = 2$ is not well-connected while the dependence structure remains non-negligible.

TABLE 1. Degree Distributions of the Generated Graphs

	min	Q1	Q3	max	mean
Panel A: $K = 1$					
$n = 500$	87	110	123	148	116.944
$n = 1000$	100	126	141	174	133.868
$n = 2500$	120	146	162	199	154.169
$n = 5000$	124	160	177	217	168.815
Panel B: $K = 2$					
$n = 500$	41	54	62	76	58.208
$n = 1000$	44	62	72	92	67.362
$n = 2500$	43	71	83	108	77.154
$n = 5000$	56	78	91	124	84.68

Panel A reports the degree distributions when the number of blocks is $K = 1$ (Erdős–Rényi model), and Panel B reports the degree distributions when the number of blocks is $K = 2$. The first column reports the number of nodes n , and the second to the fifth columns report the minimum, 25th percentile, 75th percentile, maximum, and mean of the degree.

5.2. Results: Inference.

In this exercise, we construct 95% confidence intervals for α_i based on (3.3) and evaluate the coverage probability of these confidence intervals. We also construct 95% confidence intervals for α_i based on Theorem 5 in Jochmans and Weidner (2019), which shows the asymptotic normality of $\hat{\alpha}_i$ when ϵ are independent. Specifically, the standard error of $\hat{\alpha}_i$ in the independent case is given by

$$\frac{\sqrt{\sum_{e \in E_i} \hat{\epsilon}_e^2}}{d_i}.$$

This procedure is repeated 2000 times to evaluate the coverage probability. See Appendix D for additional simulation results, including results for joint hypothesis tests discussed above.

TABLE 2. Global Measures of the Generated Graphs

	$\lambda_{2,L}$	$\lambda_{n,F}$	$ \mathcal{C}_n $	η_n	H_n
Panel A: $K = 1$					
$n = 500$	0.842	236.914	496	0.009	0.021
$n = 1000$	0.842	271.302	993	0.008	0.018
$n = 2500$	0.846	311.924	2486	0.007	0.016
$n = 5000$	0.849	341.375	4974	0.006	0.014
Panel B: $K = 2$					
$n = 500$	0.013	119.039	495	0.017	2.852
$n = 1000$	0.008	138.749	989	0.015	4.088
$n = 2500$	0.007	158.181	2478	0.013	3.987
$n = 5000$	0.005	173.399	4938	0.012	4.648

Note: Panel A reports the global measures when the number of blocks is $K = 1$ (Erdős-Rényi model), and Panel B reports the global measures when the number of blocks is $K = 2$. The first column reports the number of nodes n , the second column reports the connectivity measure $\lambda_{2,L}$, the third column reports the dependence measure $\lambda_{n,F}$, and the fourth column reports the number of nodes in $\mathcal{C}_n = \{i \in V : c_i > 0.2\}$. The last two columns report the convergence measures η_n and H_n .

Table 3 summarizes the Monte Carlo simulation results. The confidence intervals based on (3.3) achieve coverage probabilities close to the nominal 95% level for both $K = 1$ and $K = 2$. While this is expected for $K = 1$, it is notable for $K = 2$, where the graph is not well-connected and $H_1 = \lambda_{n,F}/(d_1 \lambda_{2,L} h_1)$ does not converge to zero, so the first-order approximation underlying (3.3) is not theoretically guaranteed. The strong performance in this case suggests that the approximation remains accurate in finite samples, even when the graph is not well-connected.

In contrast, confidence intervals based on the asymptotic normality of $\hat{\alpha}_i$ under independence of ϵ (based on Jochmans and Weidner, 2019) show substantial under-coverage in both $K = 1$ and $K = 2$, highlighting the importance of accounting for the dependence structure in the error terms. This under-coverage is more pronounced in a single block ($K = 1$) than in two blocks ($K = 2$), likely because the relative contribution of the independent component of the error term is smaller in this case than in $K = 2$.

5.3. Results: Variance Components.

In this exercise, we evaluate the performance of our proposed bias-correction method for estimating the sample variance of the fixed effects α . We compare the true variance V_α with the plug-in estimator \hat{V}_α , the bias-corrected estimator \hat{V}_α^{bc} , the rule-of-thumb modified bias-corrected estimator $\hat{V}_\alpha^{bc,mod}$, and the Andrews et al. (2008)-type bias-corrected

TABLE 3. Coverage Probability of the Confidence Intervals

	α_1	d_1	$(d_1^t - d_1^s)/d_1$	$\lambda_{n,F}/(d_1\lambda_{2,L}h_1)$	95% cov	Normal 95% cov
Panel A: $K = 1$						
$n = 500$	-0.596	119	-0.697	0.02	0.946	0.440
$n = 1000$	0.314	143	-0.986	0.016	0.954	0.278
$n = 2500$	0.281	161	-0.652	0.014	0.948	0.370
$n = 5000$	0.539	159	-0.635	0.016	0.944	0.403
Panel B: $K = 2$						
$n = 500$	-0.594	49	-0.878	3.886	0.948	0.516
$n = 1000$	0.312	67	-1.0	3.966	0.954	0.365
$n = 2500$	0.278	72	-0.583	4.407	0.944	0.565
$n = 5000$	0.542	93	-0.613	3.721	0.931	0.508

Note: Panel A reports the results for $K = 1$ (Erdős-Rényi model), and Panel B reports the results for $K = 2$. The first column reports the number of nodes n , the second column reports the true value of α_1 , the third and fourth columns report node 1's degree and the coefficient, respectively. The fifth column reports the convergence measure $H_1 = \lambda_{n,F}/(d_1\lambda_{2,L}h_1)$. The sixth column reports the coverage probability of the confidence intervals based on (3.3), and the seventh column reports the coverage probability of the confidence intervals based on Jochmans and Weidner (2019)'s asymptotic normality.

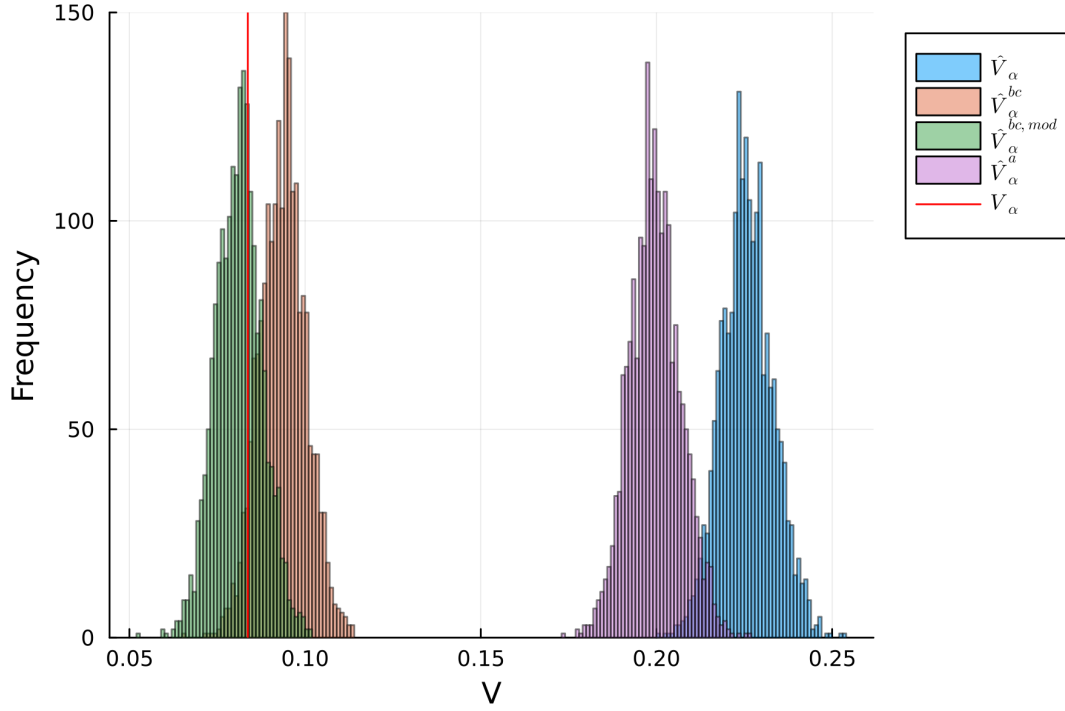
estimator \hat{V}_α^a . We simulate each of these variance estimators 2000 times and compute the mean and standard deviation of the estimated variance.

Table 4 summarizes the results. The plug-in estimator \hat{V}_α exhibits substantial upward bias. The bias-corrected estimator \hat{V}_α^{bc} is much closer to the true value V_α , demonstrating the effectiveness of our bias correction. The rule-of-thumb estimator $\hat{V}_\alpha^{bc,mod}$ performs well for $K = 1$, but tends to overcorrect and display downward bias for $K = 2$, indicating that more adaptive or refined bias-correction methods may be needed for certain graph structures. The Andrews et al. (2008)-type bias-corrected estimator \hat{V}_α^a also reduces bias, but its remaining bias is larger than that of \hat{V}_α^{bc} and $\hat{V}_\alpha^{bc,mod}$ due to unaccounted dependence. The standard deviations of the variance estimators are generally similar (except for $sd(\hat{V}_\alpha^{bc})$ for $K = 2, n = 5000$), suggesting that our bias-corrected estimator is preferable in terms of mean squared error.

6. EMPIRICAL APPLICATION

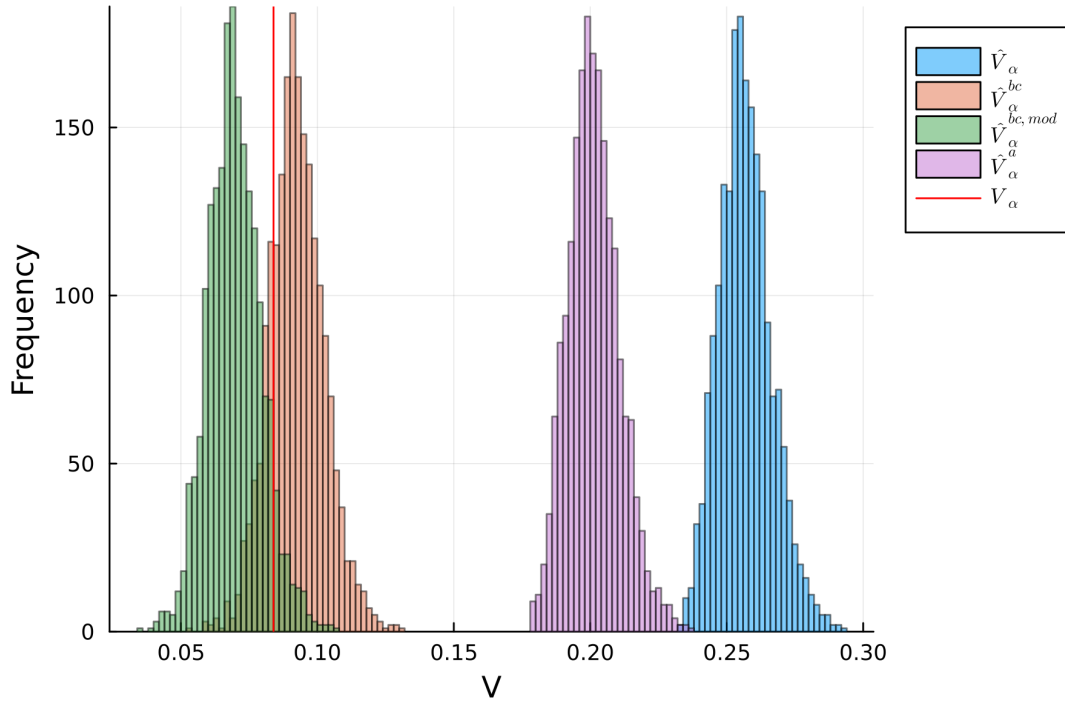
In this section, we apply our methods to an Italian worker-firm matched dataset. First, we construct a confidence interval for the fixed effect of a "central" firm in the firm-firm mobility network. Next, we estimate the variance of the firm fixed effects and report the corresponding bias-corrected variance using our proposed bias-correction method.

FIGURE 5.1. Histogram of the simulated \hat{V}_α : Erdős–Rényi Model



Note: the red line represents the true variance of α . The blue, orange, green, purple bins represent the simulated $\hat{V}_\alpha, \hat{V}_\alpha^{bc}, \hat{V}_\alpha^{bc, mod}, \hat{V}_\alpha^a$, respectively. $n = 2500$.

FIGURE 5.2. Histogram of the simulated \hat{V}_α : Stochastic-Block Model



Note: the red line represents the true variance of α . The blue, orange, green, purple bins represent the simulated $\hat{V}_\alpha, \hat{V}_\alpha^{bc}, \hat{V}_\alpha^{bc, mod}, \hat{V}_\alpha^a$, respectively. $n = 2500$.

TABLE 4. Variance Estimation

	V_{α}	\hat{V}_{α}	\hat{V}_{α}^{bc}	$\hat{V}_{\alpha}^{bc,mod}$	\hat{V}_{α}^a	$sd(\hat{V}_{\alpha})$	$sd(\hat{V}_{\alpha}^{bc})$	$sd(\hat{V}_{\alpha}^{bc,mod})$	$sd(\hat{V}_{\alpha}^a)$
Panel A: $K = 1$									
$n = 500$	0.334	0.544	0.359	0.328	0.494	0.0301	0.0282	0.0282	0.0298
$n = 1000$	0.34	0.549	0.361	0.333	0.505	0.0215	0.0201	0.0201	0.0214
$n = 2500$	0.329	0.53	0.347	0.323	0.492	0.0129	0.0122	0.0122	0.0128
$n = 5000$	0.336	0.528	0.353	0.33	0.494	0.009	0.0085	0.0085	0.0089
Panel B: $K = 2$									
$n = 500$	0.334	0.609	0.388	0.32	0.492	0.04	0.0391	0.0391	0.0396
$n = 1000$	0.34	0.602	0.388	0.329	0.504	0.0302	0.0293	0.0293	0.03
$n = 2500$	0.329	0.574	0.356	0.309	0.495	0.0165	0.0164	0.0163	0.0163
$n = 5000$	0.336	0.568	0.362	0.319	0.496	0.0106	0.0109	0.0107	0.0105

Note: Panel A reports the results for $K = 1$ (Erdős-Rényi model), and Panel B reports the results for $K = 2$. The first column reports the number of nodes n , the second column reports the true variance of α , the third to seventh columns report the mean and standard deviation of \hat{V}_{α} , \hat{V}_{α}^{bc} , $\hat{V}_{\alpha}^{bc,mod}$, and \hat{V}_{α}^a , respectively.

6.1. Data.

The data are drawn from the Veneto Worker History (VWH) file.¹¹ The VWH file contains annual job spells for all workers employed in the Italian region of Veneto. For each spell, annual wages and the number of days worked per year are reported for the period 1975-2001. We follow the same sample selection strategy as in Kline et al. (2020). Specifically, we focus on periods 1999 and 2001, on workers whose ages are between 18 and 64, and on "dominant jobs" (i.e., spells) where the worker earned the most in a given year. We also exclude workers who work in the public sector and whose wages and days worked are outliers.¹²

From this sample, we construct firm-firm network data where nodes represent firms and edges represent links between firms connected by workers who have moved from one firm to another. The construction is similar to the procedure we discussed in Example 1; our handling of the data corresponds to the two-period AKM model. We focus on movers who change firms from i to j between 1999 and 2001. For each mover, we construct a firm pair (i, j) and compute the wage difference between the two firms. We exclude firms with fewer than 15 movers, similarly to the procedure in Bonhomme et al. (2023). The firm-firm network is then defined by the set of firms and the firm pairs connected by

¹¹This dataset was developed by the Economics Department in Università Ca' Foscari Venezia under the supervision of Giuseppe Tattara.

¹²Specifically, "outliers" are defined as workers who (i) report a daily wage less than 5 euros or have zero days worked, (ii) report a log daily wage change one year to the next greater than 1 in absolute value, (iii) have more than 10 jobs in any year or have missing gender.

movers, and we extract the largest connected component from this network. The resulting graph represents our \mathcal{G} , and the wage differences constitute the outcome vector \mathbf{y} in our model.

The resulting graph consists of $|V| = 7123$ firms and $|E| = 66,770$ links. Table 5 presents the degree distribution of the firm-firm network, which is sparse with a mean degree of 18. Table 6 provides global measures of the network. Notably, the connectivity measure $\lambda_{2,L}$ is very small and close to zero, indicating that the graph is not well-connected and can be easily separated into components. The dependence measure $\lambda_{n,F}$ exceeds the maximum degree, and since $\lambda_{2,L}$ is small while $\lambda_{n,F}$ is large, the global measure H_n is also large. This suggests that, on average, the first-order approximation in Theorem 1 may not be valid, and the finite-sample bias in \hat{V}_α^{bc} is likely to be non-negligible.

TABLE 5. Degree Distribution for the Firm-Firm Network

min	Q1	Q3	max	mean
1.0	6.0	17.0	1092.0	18.748

Note: The first and fourth columns report the minimum and maximum degree of the firm-firm network. The second and third columns report the 25th and 75th percentiles of the degree distribution. The fifth column reports the average degree of the firm-firm network.

TABLE 6. Global Measures for the Firm-Firm Network

$\lambda_{2,L}$	$\lambda_{n,F}$	$ C_n $	η_n	H_n
0.006	1779.6	4611.0	0.121	6248.04

Note: The first column reports the second smallest eigenvalue of the normalized Laplacian matrix. The second column reports the largest eigenvalue of the signless Laplacian matrix. The third column reports the number of nodes with $c_i > 0.2$. The fourth and fifth columns report the global measures of convergence.

6.2. Results: Inference for Central Firms.

First, we identify the most central firm in the firm-firm mobility network by computing the PageRank centrality for each firm. PageRank centrality measures the importance of a node based on its connections to other highly central nodes.¹³ In the context of the mobility network, a firm with high PageRank centrality is one that receives many workers from other firms that themselves attract many movers. We are interested in whether the fixed effect of this central firm is significantly different from others, as this may indicate

¹³See Newman (2018) for the precise definition and discussion of PageRank centrality.

that firm fixed effects play an important role in shaping the mobility network. As a related contribution, Sorkin (2018) provides a structural interpretation of PageRank centrality when applied to a mobility network.

TABLE 7. Distribution of the Firm Fixed Effects

min	Q1	Q3	max	mean
-1.41	-1.102	0.096	0.868	-0.017

Note: The first column reports the minimum of the firm fixed effects, the second and third columns report the 25th and 75th percentiles of the firm fixed effects, respectively. The fourth column reports the maximum of the firm fixed effects, and the last column reports the mean of the firm fixed effects.

Applying the PageRank algorithm to our firm-firm mobility network, we find that the most central firm is firm 931 in our dataset. Table 8 summarizes this firm's key characteristics. Its estimated fixed effect $\hat{\alpha}_{931}$ is in approximately the top 34% of the distribution of firm fixed effects. As a central firm, it has a larger number of inflows d_{931}^t than outflows d_{931}^s , which is consistent with the idea that central firms attract more workers. The ratio $H_{931} = \lambda_{n,F}/(d_{931}\lambda_{2,L}h_{931})$ is about 2.90, which is much smaller than the average value $H_n \sim 6248$ and comparable to what we observed in the simulation study when the graph is not well-connected (see Table 3). This suggests that our inference procedure based on the first-order approximation in Theorem 1 may be valid for this central firm, and we can construct a confidence interval for $\hat{\alpha}_{931}$.

TABLE 8. Central Firm Information

$\hat{\alpha}_{931}$	d_{931}	d_{931}^t	d_{931}^s	H_{931}
0.058	455	309	146	2.90

Note: The first column reports the estimated fixed effect of the central firm, the second to fourth columns report the firm's degree, the number of inflows, and the number of outflows, respectively. The last column reports the convergence measure $H_{931} = \lambda_{n,F}/(d_{931}\lambda_{2,L}h_{931})$.

The 95% confidence interval for $\hat{\alpha}_{931}$ is:

$$\begin{aligned}
CI_{931,0.95} &= \left[\underbrace{\hat{\alpha}_{931}}_{=0.058} - \underbrace{\left(\frac{d_{931}^t - d_{931}^s}{d_{931}} \right)}_{=0.358} \underbrace{\hat{c}_{0.975}}_{=0.234}, \quad \hat{\alpha}_{931} - \underbrace{\left(\frac{d_{931}^t - d_{931}^s}{d_{931}} \right)}_{=0.358} \underbrace{\hat{c}_{0.025}}_{=-0.201} \right] \\
&= [-0.026, \quad 0.130].
\end{aligned}$$

Note that the mean of the firm fixed effects is -0.017 (see Table 7), which lies within this confidence interval. This suggests that the central firm’s fixed effect is not significantly different from the average firm. Therefore, we do not find strong evidence that firm fixed effects α are a primary driver of mobility network formation in this data.

For comparison, we can also construct confidence intervals based on the asymptotic normality of $\hat{\alpha}_i$ under the assumption of independent errors, as in the simulation section. The 95% confidence interval for $\hat{\alpha}_{931}$ using this conventional method is

$$CI_{931,0.95}^{\text{Normal}} = [0.031, \quad 0.085].$$

This interval does not include the mean of the firm fixed effects, suggesting that the central firm’s fixed effect is significantly different from the average firm. However, given the under-coverage observed for this method in the simulation study, it likely overstates the significance of the central firm’s fixed effect compared to our proposed approach. This highlights the empirical relevance of dependence in the error terms and the importance of accounting for it in inference.

This exercise demonstrates the practical applicability of our inferential method for fixed effects, which is rarely implemented in empirical work. Another potential application is hypothesis testing to assess whether the fixed effects of firms within a particular group (e.g., firms with similar observable characteristics) are statistically indistinguishable. Such a test could help determine whether grouping fixed effects, as suggested by Bonhomme et al. (2019), is appropriate in empirical analyses.

6.3. Results: Variance Estimation.

We estimate the variance of the firm fixed effects \hat{V}_{α} and several bias-corrected estimators: \hat{V}_{α}^{bc} , $\hat{V}_{\alpha}^{bc,mod}$, and \hat{V}_{α}^a . Table 9 presents the estimated variance of the firm fixed effects α and the bias-corrected estimators.

Our results show that the bias-corrected estimator \hat{V}_{α}^{bc} reduces the plug-in estimator \hat{V}_{α} by approximately 28%, which is a larger reduction than that achieved by the Andrews et al. (2008)-type bias-corrected estimator \hat{V}_{α}^a (about 23%). The modified bias-corrected estimator $\hat{V}_{\alpha}^{bc,mod}$ yields an even lower estimate. These findings indicate that the dependence structure in the data is substantial and that our proposed bias-correction methods more effectively account for this dependence, resulting in more accurate estimation of the variance of the firm fixed effects.

TABLE 9. Variance Component Estimation on the VWH Data

\hat{V}_{α}	\hat{V}_{α}^{bc}	$\hat{V}_{\alpha}^{bc,mod}$	\hat{V}_{α}^a
0.039	0.028	0.025	0.030

Note: This table reports the estimated variances of the firm fixed effects α via the plug-in estimator \hat{V}_{α} , the bias-corrected estimator \hat{V}_{α}^{bc} , the modified bias-corrected estimator $\hat{V}_{\alpha}^{bc,mod}$, and the Andrews et al. (2008)-type bias-corrected estimator \hat{V}_{α}^a .

If we disaggregate each bias correction, we obtain

$$\underbrace{\hat{\sigma}_{\tau}^2}_{=0.015} \times \underbrace{\frac{\text{tr}(\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}')}{n}}_{=0.695}, \quad \underbrace{(\hat{\sigma}^2 - 2\hat{\sigma}_{\tau}^2)}_{=0.022} \times \underbrace{\frac{\text{tr}(\mathbf{M}_n\mathbf{L}^*)}{n}}_{=0.204}$$

Note that $\hat{\sigma}_{\tau}^2$, which is an estimator for $\mathbb{E}[\tau(U_i)^2]$, is about 0.015, which is roughly one-quarter of $\hat{\sigma}^2$, the biased estimator for $\mathbb{E}[\epsilon_e^2]$. This disparity suggests that in our data the dependence structure plays a non-negligible role in generating the error terms, and this effect is transmitted to the bias in the variance estimator. Moreover, although there are uncertainties in estimating $\mathbb{E}[\tau(U_i)^2]$ and $\mathbb{E}[\epsilon_e^2]$, a heuristic comparison of the two quantities characterizing the graph structure

$$\frac{\text{tr}(\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}')}{n} \text{ and } \frac{\text{tr}(\mathbf{M}_n\mathbf{L}^*)}{n}$$

indicates that the first term is more than three times larger than the second term. Thus, even if $\hat{\sigma}_{\tau}^2$ were estimated to be smaller than the current estimate, the bias originating from the dependence structure would remain significant.

This exercise demonstrates that the bias in variance estimation due to dependence, which has been largely overlooked in the literature, can be substantial in empirical applications. It also provides evidence that the dependence structure in the data is non-negligible, highlighting the practical usefulness of our bias-correction method.

What we did not explore in this paper is how the dependence structure affects estimation of covariance between worker and firm fixed effects, which is of interest in empirical applications as it indicates sorting effects between workers and firms. We leave this to future research.

7. CONCLUSION

This paper proposes a new inference method for fixed effects in network regressions with dependent errors. We derive a first-order approximation to the fixed-effect estimator

and show that, under node- and edge-level dependence, the estimator can be inconsistent. Building on this approximation, we develop a non-standard inference procedure that exploits the empirical distribution of appropriately constructed residual averages to obtain valid confidence intervals and hypothesis tests for fixed effects.

We also study estimation of distributional moments of the fixed effects. We characterize the bias of the plug-in variance estimator under the dependence structure and propose a bias correction that is consistent in the presence of node-level dependence.

Monte Carlo simulations and an application to Italian worker–firm matched data illustrate the finite-sample performance of the methods. The results show that network dependence can substantially distort inference for fixed effects and is an important source of bias in variance estimation.

Future research could explore extensions to models with multiple types of fixed effects—such as worker and firm effects in two-way models, and develop methods for estimating covariance between different fixed-effect types. Another avenue is to devise inference procedures for variance components that go beyond the bias-correction proposed here. Finally, it would be useful to extend our methods to high-dimensional, team-based settings in which nodes may belong to multiple teams (e.g., Bonhomme, 2021), thereby accounting for dependence within and across teams.

REFERENCES

- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.
- ALVAREZ, L., B. FERMAN, AND K. WÜTHRICH (2025): “Inference with few treated units,” ArXiv:2504.19841v2 [econ.EM].
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): “High wage workers and low wage firms: negative assortative matching or limited mobility bias?” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171, 673–697.
- BACHER-HICKS, A. AND C. KOEDEL (2023): “Chapter 2 - Estimation and interpretation of teacher value added in research applications,” Elsevier, vol. 6 of *Handbook of the Economics of Education*, 93–134.
- BAGGER, J., F. FONTAINE, F. POSTEL-VINAY, AND J.-M. ROBIN (2014): “Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics,” *American Economic Review*, 104, 1551–96.

- BERNARD, A. B., E. DHYNE, G. MAGERMAN, K. MANOVA, AND A. MOXNES (2022): “The Origins of Firm Heterogeneity: A Production Network Approach,” *Journal of Political Economy*, 130, 1765–1804.
- BILLINGSLEY, P. (1999): *Convergence of Probability Measures*, Hoboken, NJ, USA: John Wiley & Sons, Inc.
- BONHOMME, S. (2021): “Teams: Heterogeneity, Sorting, and Complementarity,” ArXiv:2102.01802v1 [econ.EM].
- BONHOMME, S., K. HOLZHEU, T. LAMADON, E. MANRESA, M. MOGSTAD, AND B. SETZLER (2023): “How Much Should We Trust Estimates of Firm Effects and Worker Sorting?” *Journal of Labor Economics*, 41, 291–322.
- BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): “A Distributional Framework for Matched Employer Employee Data,” *Econometrica*, 87, 699–739.
- BOZZO, E. (2013): “The Moore–Penrose inverse of the normalized graph Laplacian,” *Linear Algebra and its Applications*, 439, 3038–3043.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2011): “Robust Inference With Multiway Clustering,” *Journal of Business & Economic Statistics*, 29, 238–249.
- CAMERON, A. C. AND D. MILLER (2014): “Robust Inference for Dyadic Data,” *Unpublished manuscript*.
- CARD, D., J. HEINING, AND P. KLINE (2013): “Workplace Heterogeneity and the Rise of West German Wage Inequality,” *The Quarterly Journal of Economics*, 128, 967–1015.
- CHIANG, H. D., B. E. HANSEN, AND Y. SASAKI (2024): “Standard Errors for Two-Way Clustering with Serially Correlated Time Effects,” *The Review of Economics and Statistics*, 1–40.
- CONLEY, T. G. AND C. R. TABER (2011): “Inference with “Difference in Differences” with a Small Number of Policy Changes,” *The Review of Economics and Statistics*, 93, 113–125.
- CVETKOVIĆ, D., P. ROWLINSON, AND S. SIMIĆ (2007): “Eigenvalue Bounds for the Signless Laplacian,” *Publications de l’Institut Mathématique*, 81(95), 11–27.
- DAVEZIES, L., X. D’HAULTFŒUILLE, AND Y. GUYONVARCH (2021): “Empirical process results for exchangeable arrays,” *The Annals of Statistics*, 49, 845 – 862.
- DENG, S., S. LING, AND T. STROHMER (2021): “Strong Consistency, Graph Laplacians, and the Stochastic Block Model,” *Journal of Machine Learning Research*, 22, 1–44.

- DI ADDARIO, S., P. KLINE, R. SAGGIO, AND M. SØLVSTEN (2023): “It ain’t where you’re from, it’s where you’re at: Hiring origins, firm heterogeneity, and wages,” *Journal of Econometrics*, 233, 340–374.
- ENGBOM, N. AND C. MOSER (2022): “Earnings Inequality and the Minimum Wage: Evidence from Brazil,” *American Economic Review*, 112, 3803–47.
- ENGBOM, N., C. MOSER, AND J. SAUERMAN (2023): “Firm pay dynamics,” *Journal of Econometrics*, 233, 396–423.
- FINKELSTEIN, A., M. GENTZKOW, AND H. WILLIAMS (2016): “Sources of Geographic Variation in Health Care: Evidence From Patient Migration*,” *The Quarterly Journal of Economics*, 131, 1681–1726.
- (2021): “Place-Based Drivers of Mortality: Evidence from Migration,” *American Economic Review*, 111, 2697–2735.
- GRAHAM, B. S. (2017): “An Econometric Model of Network Formation With Degree Heterogeneity,” *Econometrica*, 85, 1033–1063.
- HORN, R. A. AND C. R. JOHNSON (2012): *Matrix Analysis*, Cambridge University Press, 2 ed.
- JACKSON, C. K. (2013): “Match Quality, Worker Productivity, and Worker Mobility: Direct Evidence From Teachers,” *The Review of Economics and Statistics*, 95, 1096–1116.
- JOCHMANS, K. AND M. WEIDNER (2019): “Fixed-Effect Regressions on Network Data,” *Econometrica*, 87, 1543–1560.
- KALLENBERG, O. (2002): *Foundations of modern probability*, Probability and its Applications (New York), Springer-Verlag, New York, second ed.
- KLINE, P. (2024): “Chapter 2 - Firm wage effects,” Elsevier, vol. 5 of *Handbook of Labor Economics*, 115–181.
- KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2020): “Leave-Out Estimation of Variance Components,” *Econometrica*, 88, 1859–1898.
- LACHOWSKA, M., A. MAS, R. SAGGIO, AND S. A. WOODBURY (2023): “Do firm effects drift? Evidence from Washington administrative data,” *Journal of Econometrics*, 233, 375–395.
- MACKINNON, J. G. AND M. D. WEBB (2018): “The wild bootstrap for few (treated) clusters,” *The Econometrics Journal*, 21, 114–135.

- (2020): “Randomization inference for difference-in-differences with few treated clusters,” *Journal of Econometrics*, 218, 435–450.
- MENZEL, K. (2021): “Bootstrap With Cluster-Dependence in Two or More Dimensions,” *Econometrica*, 89, 2143–2188.
- NEWMAN, M. (2018): *Networks*, Oxford University Press.
- POSTEL-VINAY, F. AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70, 2295–2350.
- SORKIN, I. (2018): “Ranking Firms Using Revealed Preference*,” *The Quarterly Journal of Economics*, 133, 1331–1393.
- TABORD-MEEHAN, M. (2019): “Inference With Dyadic Data: Asymptotic Behavior of the Dyadic-Robust t-Statistic,” *Journal of Business and Economic Statistics*, 37.
- VAN DER VAART, A. W. (1998): *Asymptotic Statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press.
- VERDIER, V. (2020): “Estimation and Inference for Linear Models with Two-Way Fixed Effects and Sparsely Matched Data,” *The Review of Economics and Statistics*, 102, 1–16.

APPENDIX A. COVARIATES

Our model with covariates is given by:

$$\mathbf{y} = \mathbf{B}\boldsymbol{\alpha} + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where \mathbf{X} is a matrix of covariates. Let $\mathbf{M}_\mathbf{B} \equiv \mathbf{I}_m - \mathbf{B}(\mathbf{B}'\mathbf{B})^*\mathbf{B}'$ and $\mathbf{M}_\mathbf{X} \equiv \mathbf{I}_m - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ be the projection matrices onto the nullspace of \mathbf{B} and \mathbf{X} , respectively. Then, the fixed-effect estimators are given by:

$$\check{\boldsymbol{\alpha}} = (\mathbf{B}'\mathbf{M}_\mathbf{X}\mathbf{B})^*\mathbf{B}'\mathbf{M}_\mathbf{X}\mathbf{y}, \quad \hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1}\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{y},$$

where we assume that $\text{rank}(\mathbf{X}) = p$ and $\text{rank}((\mathbf{X}, \mathbf{B})) = p+n-1$. Our goal is to replicate the results in Theorem 1 for $\check{\alpha}_i$.

Algebraically, we have

$$\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha} = \mathbf{D}^{-1}\mathbf{B}'\boldsymbol{\epsilon} + \mathbf{r} + \tilde{\mathbf{r}},$$

where

$$\mathbf{r} = \mathbf{D}^{-1}\mathbf{A}(\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha}), \quad \tilde{\mathbf{r}} = -\mathbf{D}^{-1}\mathbf{B}'\mathbf{X}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}).$$

We already worked out the first term $\mathbf{D}^{-1}\mathbf{B}'\boldsymbol{\epsilon}$ in the proof of Theorem 1. Thus, we need to show that both \mathbf{r} and $\tilde{\mathbf{r}}$ are negligible.

First, we show that $\tilde{\mathbf{r}}$ is negligible. Note that $\mathbb{E}[\tilde{\mathbf{r}}] = 0$. Observe that

$$\hat{\boldsymbol{\beta}} - \boldsymbol{\beta} = (\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1}\mathbf{X}'\mathbf{M}_\mathbf{B}\boldsymbol{\epsilon}$$

with $\mathbb{E}[\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}] = 0$. Then,

$$\begin{aligned} \mathbb{E}[(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'] &= (\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1}\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbb{E}[\boldsymbol{\epsilon}\boldsymbol{\epsilon}']\mathbf{M}_\mathbf{B}\mathbf{X}(\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1} \\ &\leq C(\mathbf{X}\mathbf{M}_\mathbf{B}\mathbf{X})^{-1}\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{F}\mathbf{F}'\mathbf{M}_\mathbf{B}\mathbf{X}(\mathbf{X}\mathbf{M}_\mathbf{B}\mathbf{X})^{-1} \\ &\leq \lambda_{n,F}C(\mathbf{X}\mathbf{M}_\mathbf{B}\mathbf{X})^{-1} \end{aligned}$$

for some absolute constant $C > 0$ under Assumption 1: $\mathbb{E}[\boldsymbol{\epsilon}\boldsymbol{\epsilon}'] \leq C\mathbf{F}\mathbf{F}'$. Let

$$\tilde{\rho} = \|(\mathbf{X}'\mathbf{X})(\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1}\|_2,$$

where $\|\cdot\|_2$ is the spectral norm. By the definition of $\tilde{\rho}$, we have $(\mathbf{X}'\mathbf{M}_\mathbf{B}\mathbf{X})^{-1} \leq \tilde{\rho}(\mathbf{X}'\mathbf{X})^{-1}$. Thus, we have

$$\begin{aligned}\mathbb{E}[\tilde{\mathbf{r}}_i^2] &= e_i' \mathbf{D} \mathbf{B}' \mathbf{X} \mathbb{E}[(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'] \mathbf{X}' \mathbf{B} \mathbf{D}^{-1} e_i \\ &\leq \lambda_{n,F} C e_i' \mathbf{D} \mathbf{B}' \mathbf{X} (\mathbf{X}' \mathbf{M}_\mathbf{B} \mathbf{X})^{-1} \mathbf{X}' \mathbf{B} \mathbf{D}^{-1} e_i \\ &\leq \lambda_{n,F} \tilde{\rho} C e_i' \mathbf{D} \mathbf{B}' \mathbf{X} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{B} \mathbf{D}^{-1} e_i \\ &= \frac{\lambda_{n,F} \tilde{\rho} C}{m} \bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i,\end{aligned}$$

where $\bar{\mathbf{x}}_i = \mathbf{X} \mathbf{D}^{-1} e_i$ and $\boldsymbol{\Omega} = \mathbf{X}' \mathbf{X} / m$. Since $\lambda_{n,F} = O(\max_{i \in V} d_i)$, as long as $\max_{i \in V} d_i / m = o(1)$ and $\tilde{\rho} \times \bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i = O(1)$, we have $\tilde{\mathbf{r}}_i = o_p(1)$.

Second, we show that \mathbf{r} is negligible. Note that $\mathbb{E}[\mathbf{r}] = 0$. Let

$$\rho = \|(\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{M}_\mathbf{B} \mathbf{X}\|_2.$$

We have

$$\begin{aligned}\mathbb{E}[\mathbf{r}_i^2] &= e_i' \mathbf{D}^{-1} \mathbf{A} \mathbb{E}[(\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha})(\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha})'] \mathbf{A}' \mathbf{D}^{-1} e_i \\ &= e_i' \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{M}_\mathbf{X} \mathbf{B})^* \mathbf{B}' \mathbf{M}_\mathbf{X} \mathbb{E}[\boldsymbol{\epsilon} \boldsymbol{\epsilon}'] \mathbf{M}_\mathbf{X} \mathbf{B} (\mathbf{B}' \mathbf{M}_\mathbf{X} \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i \\ &\leq \lambda_{n,F} C \times e_i' \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{M}_\mathbf{X} \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i \\ &\leq \lambda_{n,F} C \times e_i' \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i \\ &\quad + \lambda_{n,F} C \times \rho^{-1} \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{B})^* \mathbf{B}' \mathbf{X} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{B} (\mathbf{B}' \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i\end{aligned}$$

where the last line follows from (S.14) in Jochmans and Weidner (2019). Since $(\mathbf{B}' \mathbf{B})^* \leq \lambda_{2,L}^{-1} \mathbf{D}^{-1}$ and $\mathbf{X} (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \leq I$, we have

$$\begin{aligned}\mathbb{E}[\mathbf{r}_i^2] &\leq \lambda_{n,F} C \times e_i' \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i + \lambda_{n,F} C \times \rho^{-1} \mathbf{D}^{-1} \mathbf{A} (\mathbf{B}' \mathbf{B})^* \mathbf{A}' \mathbf{D}^{-1} e_i \\ &\leq \frac{\lambda_{n,F} C (1 + \rho)}{\rho \lambda_{2,L}} \times e_i' \mathbf{D}^{-1} \mathbf{A} \mathbf{D}^{-1} \mathbf{A}' \mathbf{D}^{-1} e_i \\ &= \frac{\lambda_{n,F} C (1 + \rho)}{\rho \lambda_{2,L} h_i d_i}.\end{aligned}$$

If $\lambda_{n,F} / (d_i \lambda_{2,L} h_i) = o(1)$, $\mathbf{r}_i = o_p(1)$ as long as $\rho^{-1} = O(1)$.

In summary, as long as

- $\max_{i \in V} d_i = o(m)$,
- $\tilde{\rho} \times \bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i = O(1)$

- $\rho^{-1} = O(1)$,

then Theorem 1 holds for $\check{\alpha}$ as well with the same approximation error rates.

APPENDIX B. ALTERNATIVE ESTIMATOR

We can consider an alternative estimator for α based on the following transformation.

B.1. Transformation. Let $V^{st} = \{i \in V : d_i^t > 0 \text{ and } d_i^s > 0\}$ be the set of nodes with both positive in- and out-degrees. Let $V^t = \{i \in V : d_i^t > 0\}$ and $V^s = \{i \in V : d_i^s > 0\}$ be the sets of nodes with positive in-degree and out-degree, respectively. Redefine τ^t and τ^s as $\tau^t = (\mathbb{I}\{i \in V^t\} \cdot \tau^t(U_i))_{i \in V}$ and $\tau^s = (\mathbb{I}\{i \in V^s\} \cdot \tau^s(U_i))_{i \in V}$. Define the following vectors:

$$\begin{aligned}\tilde{\tau}^t &= (\tau^t(U_i))_{i \in V^{st}}, \quad \tilde{\tau}^s = (\tau^s(U_i))_{i \in V^{st}}, \\ \check{\tau}^t &= (\mathbb{I}\{i \in V^t \setminus V^s\} \cdot \tau^t(U_i))_{i \in V}, \quad \check{\tau}^s = (\mathbb{I}\{i \in V^s \setminus V^t\} \cdot \tau^s(U_i))_{i \in V}.\end{aligned}$$

Let $\tilde{\mathbf{F}}$ be the $m \times |V^{st}|$ matrix defined by: for each $e \in E$ and $i \in V^{st}$,

$$\tilde{\mathbf{F}}_{e,i} = \begin{cases} 1 & \text{if } s(e) = i \text{ or } t(e) = i, \\ 0 & \text{otherwise} \end{cases}$$

Finally, let $\tilde{\epsilon} = \epsilon - \mathbf{F}^t \tau^t - \mathbf{F}^s \tau^s$.

Note that the model can be rewritten as

$$\begin{aligned}\mathbf{y} &= \mathbf{B}\alpha + \epsilon \\ &= \mathbf{B}\alpha + \mathbf{F}^t \tau^t + \mathbf{F}^s \tau^s + \tilde{\epsilon} \\ &= \mathbf{B} \left(\alpha + \frac{\tau^t}{2} - \frac{\tau^s}{2} \right) + \mathbf{F} \frac{(\tau^t + \tau^s)}{2} + \tilde{\epsilon},\end{aligned}$$

where we used the facts that $\mathbf{F}^t = (\mathbf{F} + \mathbf{B})/2$ and $\mathbf{F}^s = (\mathbf{F} - \mathbf{B})/2$. For each $e \in E$,

$$\left(\mathbf{F} \frac{(\tau^t + \tau^s)}{2} \right)_e = \begin{cases} \frac{\tau_{t(e)}^t + \tau_{t(e)}^s + \tau_{s(e)}^t + \tau_{s(e)}^s}{2} & \text{if } s(e), t(e) \in V^{st} \\ \frac{\tau_{t(e)}^t + \tau_{t(e)}^s + \tau_{s(e)}^s}{2} & \text{if } t(e) \in V^{st}, s(e) \notin V^{st}, s(e) \in V^s \\ \frac{\tau_{t(e)}^t + \tau_{s(e)}^t + \tau_{s(e)}^s}{2} & \text{if } s(e) \in V^{st}, t(e) \notin V^{st}, t(e) \in V^t \\ \frac{\tau_{t(e)}^t + \tau_{s(e)}^s}{2} & \text{if } t(e), s(e) \notin V^{st}, t(e) \in V^t, s(e) \in V^s \end{cases}$$

This can be decomposed as

$$\left(\mathbf{F} \frac{(\boldsymbol{\tau}^t + \boldsymbol{\tau}^s)}{2} \right)_e = \left(\tilde{\mathbf{F}} \frac{(\tilde{\boldsymbol{\tau}}^t + \tilde{\boldsymbol{\tau}}^s)}{2} \right)_e + \left(\mathbf{B} \frac{(\tilde{\boldsymbol{\tau}}^t - \tilde{\boldsymbol{\tau}}^s)}{2} \right)_e,$$

where

$$\left(\tilde{\mathbf{F}} \frac{(\tilde{\boldsymbol{\tau}}^t + \tilde{\boldsymbol{\tau}}^s)}{2} \right)_e = \begin{cases} \frac{\tau_{t(e)}^t + \tau_{t(e)}^s + \tau_{s(e)}^t + \tau_{s(e)}^s}{2} & \text{if } s(e), t(e) \in V^{st} \\ \frac{\tau_{t(e)}^t + \tau_{t(e)}^s}{2} & \text{if } t(e) \in V^{st}, s(e) \notin V^{st} \\ \frac{\tau_{s(e)}^t + \tau_{s(e)}^s}{2} & \text{if } s(e) \in V^{st}, t(e) \notin V^{st} \\ 0 & \text{if } t(e), s(e) \notin V^{st} \end{cases}$$

$$\left(\mathbf{B} \frac{(\tilde{\boldsymbol{\tau}}^t - \tilde{\boldsymbol{\tau}}^s)}{2} \right)_e = \begin{cases} \frac{\tau_{t(e)}^t}{2} & \text{if } t(e) \in V^t \setminus V^s, s(e) \in V^{st} \\ \frac{\tau_{s(e)}^s}{2} & \text{if } s(e) \in V^s \setminus V^t, t(e) \in V^{st} \\ \frac{\tau_{t(e)}^t + \tau_{s(e)}^s}{2} & \text{if } t(e), s(e) \notin V^{st}, t(e) \in V^t, s(e) \in V^s \\ 0 & \text{if } t(e), s(e) \in V^{st} \end{cases}$$

Therefore,

$$\frac{\mathbf{F}}{2}(\boldsymbol{\tau}^t + \boldsymbol{\tau}^s) = \tilde{\mathbf{F}} \frac{(\tilde{\boldsymbol{\tau}}^t + \tilde{\boldsymbol{\tau}}^s)}{2} + \mathbf{B} \frac{(\tilde{\boldsymbol{\tau}}^t - \tilde{\boldsymbol{\tau}}^s)}{2}.$$

Substituting this into the model, we obtain

$$\mathbf{y} = \mathbf{B} \left(\boldsymbol{\alpha} + \frac{\boldsymbol{\tau}^t + \tilde{\boldsymbol{\tau}}^t}{2} - \frac{\boldsymbol{\tau}^s + \tilde{\boldsymbol{\tau}}^s}{2} \right) + \tilde{\mathbf{F}} \frac{(\tilde{\boldsymbol{\tau}}^t + \tilde{\boldsymbol{\tau}}^s)}{2} + \tilde{\boldsymbol{\epsilon}}. \quad (\text{B.1})$$

This is the transformed model where the effects explained by the incidence matrix \mathbf{B} are separated from those not explained by it.

B.2. Alternative Estimator. Let $\mathbf{M}_{\tilde{\mathbf{F}}} = \mathbf{I}_m - \tilde{\mathbf{F}}(\tilde{\mathbf{F}}'\tilde{\mathbf{F}})^+\tilde{\mathbf{F}}'$, which is the projection matrix onto the nullspace of $\tilde{\mathbf{F}}$. Premultiplying both sides of equation (B.1) by $\mathbf{M}_{\tilde{\mathbf{F}}}$, we have

$$\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{y} = \mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B} \left(\boldsymbol{\alpha} + \frac{\boldsymbol{\tau}^t + \tilde{\boldsymbol{\tau}}^t}{2} - \frac{\boldsymbol{\tau}^s + \tilde{\boldsymbol{\tau}}^s}{2} \right) + \mathbf{M}_{\tilde{\mathbf{F}}}\tilde{\boldsymbol{\epsilon}}.$$

We can consider the following alternative estimator for α as if $\alpha + \frac{\tau^t + \check{\tau}^t}{2} - \frac{\tau^s + \check{\tau}^s}{2}$ were the true fixed effects:

$$\begin{aligned}\hat{\alpha}^{alt} &= (\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{y} \\ &= \alpha + (\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B} \left(\frac{\tau^t + \check{\tau}^t}{2} - \frac{\tau^s + \check{\tau}^s}{2} \right) + (\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\tilde{\epsilon}.\end{aligned}$$

An implicit assumption here is that $\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B}$ has rank $n - 1$ so that $(\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*$ is well-defined. This requires that $\text{rank}((\mathbf{B}, \tilde{\mathbf{F}})) = n - 1 + \text{rank}(\tilde{\mathbf{F}})$, which means that the column spaces of \mathbf{B} and $\tilde{\mathbf{F}}$ do not overlap except for the intercept. This assumption holds if the graph contains enough directed cycles so that a column of $\tilde{\mathbf{F}}$ cannot be represented as a linear combination of columns of \mathbf{B} . This is likely to hold in practice if the graph is dense enough and directions of edges are not too concentrated in one direction.

B.3. Approximation. We can derive a first-order approximation for $\hat{\alpha}^{alt}$ similar to Theorem 1. Note that

$$\begin{aligned}\hat{\alpha}^{alt} - \alpha &= \frac{\tau^t + \check{\tau}^t}{2} - \frac{\tau^s + \check{\tau}^s}{2} \\ &\quad + (\mathbf{I}_n - (\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B}) \left(\frac{\tau^t + \check{\tau}^t}{2} - \frac{\tau^s + \check{\tau}^s}{2} \right) \\ &\quad + (\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\tilde{\epsilon},\end{aligned}$$

where the first term on the right-hand side is the persistent term, the second term is a misspecification error, and the last term is a linear combination of $\tilde{\epsilon}$. Each element of the persistent term can be written as

$$\left(\frac{\tau^t + \check{\tau}^t}{2} - \frac{\tau^s + \check{\tau}^s}{2} \right)_i = \begin{cases} \frac{\tau^t(U_i) - \tau^s(U_i)}{2} & \text{if } i \in V^{st} \\ \tau^t(U_i) & \text{if } i \in V^t \setminus V^s \\ -\tau^s(U_i) & \text{if } i \in V^s \setminus V^t \end{cases}$$

Importantly, in the symmetric case where $\tau^t = \tau^s = \tau$, if $i \in V^{st}$, the persistent term is zero. This means that for nodes with both positive in- and out-degrees, the alternative estimator $\hat{\alpha}_i^{alt}$ can be consistent for α_i , unlike the original estimator $\hat{\alpha}_i$. This is because the transformation effectively removes the strongly dependent component captured by τ for these nodes. Also, we can view this result as an automatic correction for the persistent term, which is similar to $\hat{\alpha}_i - \frac{d_i^t - d_i^s}{d_i} \hat{\tau}_i$ in the original estimator.

For the other two terms, once the difference between $(\mathbf{B}'\mathbf{M}_{\tilde{\mathbf{F}}}\mathbf{B})^*$ and $(\mathbf{B}'\mathbf{B})^*$ is characterized, we can follow similar steps as in the proof of Proposition 2 and Theorem 1 to derive the approximation with different error rates. Noting that $\tilde{\epsilon}$ is weakly dependent and following Appendix A, we can show that the approximation error rate for the last two terms is given by

$$O_p \left(\sqrt{\frac{1}{\rho d_i}} + \sqrt{\frac{1}{\rho \lambda_{2,L} h_i d_i}} + \sqrt{\frac{\tilde{\rho}}{m} \tilde{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \tilde{\mathbf{x}}_i} \right),$$

where $\rho, \tilde{\rho}, \tilde{\mathbf{x}}_i, \boldsymbol{\Omega}$ are defined in Appendix A with \mathbf{X} replaced by $\tilde{\mathbf{F}}$. Since $\mathbf{M}_{\tilde{\mathbf{F}}}$ eliminates the strongly dependent component, we do not have $\lambda_{n,F}$ in this rate unlike the rate in Theorem 1. However, ρ^{-1} and $\tilde{\rho}$ may be large depending on the relationship between \mathbf{B} and $\tilde{\mathbf{F}}$, which may make the approximation error rate larger than that in Theorem 1. A more detailed comparison of the two error rates and the choice between the two estimators is left for future research.

APPENDIX C. INFERENCE WITH ASYMMETRIC τ^t AND τ^s

In the main text, our analysis with regard to inference and variance estimation focused on the symmetric case where $\tau^t = \tau^s$ is imposed by Assumption 5. In this appendix, we discuss potential extensions to the asymmetric case where $\tau^t \neq \tau^s$.

As argued in the main text, the central problem in the asymmetric case is that $\tau^t(U_i)$ and $\tau^s(U_i)$ are not identified separately from the following two approximations:

$$\begin{aligned} \hat{\alpha}_i &\sim \alpha_i + \frac{d_i^t}{d_i} \tau^t(U_i) - \frac{d_i^s}{d_i} \tau^s(U_i); \\ \hat{\tau}_i &\sim \frac{\tau^t(U_i) + \tau^s(U_i)}{2}. \end{aligned}$$

In the following, we first discuss possible restrictions on $\tau^t(U_i)$ and $\tau^s(U_i)$ that can help identify them separately, and then we proceed to show that inference can still be conducted under such restrictions.

C.1. Restrictions on τ^t and τ^s . One possible restriction is to assume the linear relationship between τ^t and τ^s : for some function τ ,

$$\tau^s = \tau, \quad \tau^t = c\tau,$$

where c is an absolute constant. Instead, we can also treat τ^t as a basis function of τ^s . Note also that we do not need to place an intercept restriction on τ^t as $\mathbb{E}[\tau^t(U_i)] = \mathbb{E}[\tau^s(U_i)] = 0$ by Assumption 1. Furthermore, $\tau^t = \tau^s$ is a special case of this restriction with $c = 1$.

Under this restriction, the approximations for $\hat{\alpha}_i$ and $\hat{\tau}_i$ become:

$$\begin{aligned}\hat{\alpha}_i &\sim \alpha_i + \frac{cd_i^t - d_i^s}{d_i} \tau(U_i), \\ \hat{\tau}_i &\sim \frac{(c+1)}{2} \tau(U_i).\end{aligned}$$

Write

$$\begin{aligned}\hat{\alpha}_i &= \alpha_i + \frac{cd_i^t - d_i^s}{d_i} \tau(U_i) + r_{i,1} \\ \hat{\tau}_i &= \frac{c+1}{2} \tau(U_i) + r_{i,2},\end{aligned}$$

where $r_{i,1}$ and $r_{i,2}$ are the approximation errors. By the proof of Theorem 1, $r_{i,2}$ is $O_p(\sqrt{1/d_i})$. For $r_{i,1}$, it is known that $r_{i,1} = O_p(\sqrt{\lambda_{n,F}/(d_i \lambda_{2,L} h_i)})$ by Theorem 1.

We construct the estimator for c in the following manner. First, we regress $\hat{\alpha}_i$ on $\hat{\tau}_i$ and obtain the following regression coefficient:

$$\hat{\beta}_n \equiv \frac{|\mathcal{C}_n|^{-1} \sum_{i \in \mathcal{C}_n} \hat{\tau}_i \hat{\alpha}_i}{|\mathcal{C}_n|^{-1} \sum_{i \in \mathcal{C}_n} \hat{\tau}_i^2}.$$

By Lemma 3, the denominator converges to $((c+1)/2)^2 \mathbb{E}[\tau(U_i)^2]$. The convergence of the numerator is shown as follows:

$$\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \hat{\tau}_i \hat{\alpha}_i = \frac{c+1}{2} \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \frac{cd_i^t - d_i^s}{d_i} \times \tau(U_i)^2 \quad (\text{C.1})$$

$$+ \frac{c+1}{2} \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \alpha_i \tau(U_i) \quad (\text{C.2})$$

$$+ \frac{c+1}{2} \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} r_{i,1} \tau(U_i) \quad (\text{C.3})$$

$$+ \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} r_{i,2} \alpha_i \quad (\text{C.4})$$

$$+ \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} r_{i,2} \frac{cd_i^t - d_i^s}{d_i} \tau(U_i) \quad (\text{C.5})$$

$$+ \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} r_{i,2} r_{i,1}. \quad (\text{C.6})$$

The first term (C.1) converges to $(c+1)/2 \times \mathbb{E}[\tau(U_i)^2] \lim_{n \rightarrow \infty} |\mathcal{C}_n|^{-1} \sum_{i \in \mathcal{C}_n} (cd_i^t - d_i^s)/d_i$ and the second term (C.2) converges to 0 in probability by the law of large numbers as $\mathbb{E}[\tau^2(U_i)] < \infty$, $|(cd_i^t - d_i^s)/d_i| \leq |c| + 1 < \infty$, and $\max_{i \in \mathcal{C}_n} |\alpha_i| = O(1)$. The third term (C.3) converges to 0 in probability since

$$\mathbb{E}[|r_{i,1}\tau(U_i)|] \leq \sqrt{\mathbb{E}[r_{i,1}^2]} \sqrt{\mathbb{E}[\tau(U_i)^2]} = O\left(\sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}}\right)$$

by the Cauchy-Schwarz inequality so that

$$\mathbb{E}[(C.3)] = O_p\left(\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}}\right) = o(1).$$

Thus, (C.3) = $o_p(1)$ by the Markov inequality. By the similar argument, the fourth to sixth terms (C.4), (C.5), and (C.6) also converge to 0 in probability (if $1/|\mathcal{C}_n| \sum_{i \in \mathcal{C}_n} \sqrt{1/d_i} = o(1)$). Therefore, by the continuous mapping theorem, we have

$$\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \hat{\tau}_i \hat{\alpha}_i \rightarrow_p \frac{c+1}{2} \mathbb{E}[\tau(U_i)^2] \lim_{n \rightarrow \infty} \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} (cd_i^t - d_i^s)/d_i.$$

Thus, writing $d^t \equiv \lim_{n \rightarrow \infty} |\mathcal{C}_n|^{-1} \sum_{i \in \mathcal{C}_n} d_i^t/d_i$ and d^s analogously, we have

$$\hat{\beta}_n \rightarrow_p \frac{2(cd^t - d^s)}{c+1}.$$

Therefore, we can construct a consistent estimator for c as follows:

$$\hat{c}_n = \frac{2d^s + \hat{\beta}_n}{2d^t - \hat{\beta}_n} + O_p\left(\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}}\right),$$

as $n \rightarrow \infty$.

Hence, by collecting $2(\hat{c}_n + 1)^{-1} \hat{\tau}_i \sim \tau(U_i)$, we can construct the consistent estimator for F_τ and conduct inference on $\hat{\alpha}_i$ similar to the symmetric case.

One could also consider a more general case where τ^t and τ^s are related by a nonlinear function, such as $\tau^t = g(\tau^s)$ for some function g . This generalizes the linear case, since g could be linear, but additional restrictions on the distribution of $\tau(U_i)$ are needed to obtain useful moment conditions for identifying g . For example, suppose g is invertible and its inverse g^{-1} is odd, and $\tau(U_i)$ is symmetrically distributed around 0. Then, defining $\tilde{\tau}_i = (g(\tau(U_i)) + \tau(U_i))/2$, we have

$$\sum_{i \in \mathcal{C}_n} \mathbb{E}\left[\tilde{\tau}_i^{2m-1} \left(\alpha_i - \tau(U_i) + \frac{2d_i^t}{d_i} g^{-1}(\tilde{\tau}_i)\right)\right] = 0$$

for each integer $m \geq 1$, by g^{-1} being an odd function and the symmetry of $\tilde{\tau}_i$ with respect to $\tau(U_i)$. Estimation of g can proceed by parameterizing g as an odd polynomial function, such as $g(x) = \sum_{k=0}^K a_k x^{2k+1}$, and estimating the coefficients a_k using the method of moments up to some finite order K . We do not pursue this direction in this paper, but it is a potential extension of our method.

C.2. Inference.

Focus on the case where $\tau^s = \tau$ and $\tau^t = c\tau$ for some constant c . We have shown that \hat{c}_n consistently estimates c . Write $\hat{\tau}_i = 2(\hat{c}_n + 1)^{-1}\hat{\tau}_i$, which is a consistent estimator for $\tau(U_i)$. Then, the distribution F_τ can be estimated by the empirical distribution of $\hat{\tau}_i$, denoted by \hat{F}_τ :

$$\hat{F}_\tau(\cdot) \equiv \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{I}(\hat{\tau}_i \leq \cdot),$$

where $\mathbb{I}(\cdot)$ is the indicator function.

The uniform consistency of \hat{F}_τ to F_τ over a compact subset of \mathbb{R} follows from modifying the proof of Theorem 2. First, observe that

$$\frac{2}{\hat{c}_n + 1} = \frac{2}{c + 1} + r_{n,3}, \quad r_{n,3} = O_p \left(\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}} \right),$$

by the delta method so that

$$\begin{aligned} \hat{\tau}_i &= \frac{2}{\hat{c}_n + 1} \hat{\tau}(U_i) \\ &= \tau(U_i) + \frac{2}{c + 1} r_{i,2} + \frac{c + 1}{2} \tau(U_i) r_{n,3} + r_{i,2} r_{n,3}. \end{aligned}$$

We can bound $\mathbb{E}[|\hat{\tau}_i - \tau(U_i)|]$ by the rates of convergence of $r_{i,2}$ and $r_{n,3}$. By replacing $\hat{\tau}_i$ with $\hat{\tau}_i$ in the proof of Theorem 2, we have

$$\sup_{t \in \mathbb{R}} |\hat{F}_\tau(t) - F_\tau(t)| \rightarrow_p 0,$$

as $n \rightarrow \infty$, similarly as before but with slower rates of convergence.

With \hat{F}_τ , we can conduct inference on $\hat{\alpha}_i$ in the same manner as in the symmetric case. For example, a $(1 - \alpha)$ -level confidence interval for α_i is given by

$$CI_{i,1-\alpha} = \begin{cases} \left[\hat{\alpha}_i - \left(\frac{d_i^t \hat{c}_n - d_i^s}{d_i} \right) \hat{c}_{1-\alpha/2}, & \hat{\alpha}_i - \left(\frac{d_i^t \hat{c}_n - d_i^s}{d_i} \right) \hat{c}_{\alpha/2} \right], & \text{if } d_i^t \hat{c}_n - d_i^s > 0, \\ \left[\hat{\alpha}_i - \left(\frac{d_i^t \hat{c}_n - d_i^s}{d_i} \right) \hat{c}_{\alpha/2}, & \hat{\alpha}_i - \left(\frac{d_i^t \hat{c}_n - d_i^s}{d_i} \right) \hat{c}_{1-\alpha/2} \right], & \text{if } d_i^t \hat{c}_n - d_i^s < 0, \end{cases}$$

where \hat{c}_α is the α -quantile of \hat{F}_τ .

APPENDIX D. ADDITIONAL SIMULATION

In this Appendix, we present additional simulation results to complement the exercises in the main text. Unless otherwise specified, we use the same simulation settings as in Section 5.

D.1. Inference with non-normal errors. In Section 5, the distribution of $\tau(U_i)$ was assumed to be standard normal. Here, we consider the case where $\tau(U_i)$ follows a standard logistic distribution, which has heavier tails than the normal distribution, to verify the validity of the inference procedure to the non-normal case. Specifically, we modify the simulation settings by setting

$$U_i \sim \text{Logistic}(0, 1), \quad \tau(U_i) = U_i.$$

The results are reported in Table 10. As in the normal case, the coverage probabilities of the confidence intervals based on (3.3) are close to the nominal level of 95% for both $K = 1$ and $K = 2$. As before, the normal approximation based on Jochmans and Weidner (2019) is not valid in both cases. This result suggests that our inference procedure is not sensitive to the distribution of $\tau(U_i)$ once the regularity conditions are satisfied.

TABLE 10. Coverage Probability of Confidence intervals

	α_1	d_1	$(d_1^+ - d_1^-)/d_1$	$\lambda_{n,F}/(d_1\lambda_{2,L}h_1)$	95% cov	Normal 95% cov
Panel A: $K = 1$						
$n = 500$	-0.596	119	-0.697	0.02	0.949	0.333
$n = 1000$	0.314	143	-0.986	0.016	0.952	0.2015
$n = 2500$	0.281	161	-0.652	0.014	0.944	0.2905
$n = 5000$	0.539	159	-0.635	0.016	0.94	0.3145
Panel B: $K = 2$						
$n = 500$	-0.594	49	-0.878	3.886	0.942	0.3975
$n = 1000$	0.312	67	-1.0	3.966	0.956	0.3
$n = 2500$	0.278	72	-0.583	4.407	0.944	0.4645
$n = 5000$	0.542	93	-0.613	3.721	0.96	0.418

Panel A reports the results for $K = 1$ and Panel B reports the results for $K = 2$. The first column reports the number of nodes n , the second column reports the true value of α_1 , the third and fourth columns report node 1's degree and the coefficient, respectively. The fifth column reports the convergence measure $H_1 = \lambda_{n,F}/(d_1\lambda_{2,L}h_1)$. The sixth column reports the coverage probability of the confidence interval for α_1 at the 95% level based on (3.3). The last column reports the coverage probability of the confidence intervals based on Jochmans and Weidner (2019)'s asymptotic normality.

D.2. Testing the joint hypothesis. In the main text, we proposed a Wald-type test for testing $H_0 : \alpha_{V_0} = \mathbf{a}$ versus $H_1 : \alpha_{V_0} \neq \mathbf{a}$, where \mathbf{a} is a vector of constants. Here, we present some preliminary simulation results to check the size and power of the test.

Let $V_0 = \{1, 2, \dots, 9, 10\}$ and $\mathbf{a} = (0, 0, \dots, 0)$. We then generate the data as before except that we set $\alpha_{V_0 \setminus \{1\}} = 0$ and $\alpha_1 \in \{0, 1, 2, 3\}$, corresponding to the null hypothesis and the alternative hypothesis. When computing the test statistic (3.4) and simulate the critical value, we drop the nodes in V_0 if $(d_i^t + d_i^s)/d_i \leq 0.4$ to avoid the cases where the first-order approximation is not valid. If the threshold is closer to 1, the test is likely to be less powerful, but the size is more likely to be controlled.

Table 11 reports the rejection rates of the joint hypothesis test at the 95% significance level. The results show that the size of the test is controlled at the nominal level of 5% for both $K = 1$ and $K = 2$. As expected, the power of the test increases as α_1 increases. The power tends to be higher for $K = 1$ than for $K = 2$, likely due to the fact that the first-order approximation is more accurate for more well-connected networks. However, there is non-monotonicity in the power of the test as n increases, possibly because the change in the network structure affects the values of d_i^t and d_i^s , which in turn affects which nodes are included in the test and its performance.

TABLE 11. Size and Power of the Joint Hypothesis Test

	$H_0 : \alpha_1 = 0$	$H_1 : \alpha_1 = 1$	$H_1 : \alpha_1 = 2$	$H_1 : \alpha_1 = 3$
Panel A: $K = 1$				
$n = 500$	0.049	0.18	0.626	0.947
$n = 1000$	0.054	0.13	0.429	0.819
$n = 2500$	0.04	0.161	0.58	0.93
$n = 5000$	0.048	0.164	0.57	0.903
Panel B: $K = 2$				
$n = 500$	0.051	0.138	0.471	0.852
$n = 1000$	0.045	0.143	0.429	0.768
$n = 2500$	0.041	0.161	0.551	0.912
$n = 5000$	0.04	0.169	0.617	0.945

Note: Panel A reports the results for $K = 1$ and Panel B reports the results for $K = 2$. Each column reports the rejection rate of the joint hypothesis test at the 95% significance level. In each case, $\alpha_{V_0 \setminus \{1\}} = 0$ but $\alpha_1 = 0$ for the first column, and $\alpha_1 = 1, 2, 3$ for the second to fourth columns.

APPENDIX E. PROOFS

E.1. Proof of Theorem 1.

Proof. We bound each term in the right-hand side of (3.1):

$$\hat{\alpha} - \alpha = \mathbf{D}^{-1} \mathbf{B}'(\mathbf{F}^s \boldsymbol{\tau}^s + \mathbf{F}^t \boldsymbol{\tau}^t) + \mathbf{D}^{-1} \mathbf{B}'(\boldsymbol{\epsilon} - \mathbf{F}^s \boldsymbol{\tau}^s - \mathbf{F}^t \boldsymbol{\tau}^t) + \mathbf{r}$$

We first show the unbiasedness of the fixed-effect estimator:

Lemma 1. *Under the stated assumptions, we have*

$$\mathbb{E}[\hat{\alpha} - \alpha] = 0.$$

Proof. Observe that

$$\hat{\alpha} = \mathbf{L}^* \mathbf{L} \alpha + \mathbf{L}^* \mathbf{B}' \boldsymbol{\epsilon}.$$

Since $\mathbb{E}[\boldsymbol{\epsilon}] = 0$, it suffices to show that $\mathbf{L}^* \mathbf{L} \alpha = \alpha$. By the proof of Theorem 2 and S.1 in Jochmans and Weidner (2019), we have

$$\mathbf{L}^* \mathbf{L} = \mathbf{I}_n - m^{-1} \boldsymbol{\iota}_n \mathbf{d}'.$$

Since $\mathbf{d}' \alpha = 0$ by Assumption 1, we have

$$\mathbf{L}^* \mathbf{L} \alpha = \mathbf{I}_n \alpha - m^{-1} \boldsymbol{\iota}_n \mathbf{d}' \alpha = \alpha.$$

This proves the lemma. □

Another useful lemma is the following:

Lemma 2. *Under the stated assumptions, $C \lambda_{n,F} \geq \tilde{\sigma}_n^2$ for some absolute constant $C > 0$.*

Proof. Note that $\mathbf{F} \mathbf{F}'$ is a non-negative, positive-semidefinite matrix and each block corresponding to edges in $E_{(i,j)}$, $(\mathbf{F} \mathbf{F}')_{(i,j)}$ is a principal submatrix of $\mathbf{F} \mathbf{F}'$. Thus, by Corollary 8.1.20 in Horn and Johnson (2012), we have

$$\lambda_{n,F} \geq \lambda_{\max}((\mathbf{F} \mathbf{F}')_{(i,j)}),$$

for each (i, j) and $\lambda_{\max}(\cdot)$ is the largest eigenvalue of a matrix. Notice that

$$(\mathbf{F} \mathbf{F}')_{(i,j)} = 2 \times \boldsymbol{\iota}_{|E_{(i,j)}|} \boldsymbol{\iota}_{|E_{(i,j)}|}'$$

so that $\lambda_{\max}((\mathbf{F}\mathbf{F}')_{(i,j)}) = 2 \times |E_{(i,j)}|$. Therefore,

$$\lambda_{n,F} \geq 2 \max_{(i,j)} |E_{(i,j)}|.$$

By Assumption 1, we have

$$\begin{aligned} \tilde{\sigma}_n^2 &= \max_{(i,j)} \mathbf{A}_{i,j} \times [\max_{e,e'} E_{(i,j)} \sigma_{e,e'} - \mathbb{E}[\tau^s(U_i)^2 + \tau^t(U_i)^2]] \\ &\leq C \max_{(i,j)} \mathbf{A}_{i,j} = C \max_{(i,j)} |E_{(i,j)}|, \end{aligned}$$

for some absolute constant $C > 0$. Thus, we have

$$\tilde{\sigma}_n^2 \leq C/2 \times 2 \max_{(i,j)} |E_{(i,j)}| \leq C/2 \times \lambda_{n,F}.$$

This proves the lemma. □

Next, we bound the variance of the remaining term $\mathbf{r} = \mathbf{D}^{-1} \mathbf{A}(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha})$. Observe that

$$\text{Var}(\hat{\boldsymbol{\alpha}}) \leq C \lambda_{n,F} \mathbf{L}^*$$

for some constant $C > 0$ by Lemma 2. Thus, we have

$$\begin{aligned} \text{Var}(\mathbf{r}) &= \mathbf{D}^{-1} \mathbf{A} \text{Var}(\hat{\boldsymbol{\alpha}}) \mathbf{A}' \mathbf{D}^{-1} \\ &\leq C \lambda_{n,F} \mathbf{D}^{-1} \mathbf{A} \mathbf{L}^* \mathbf{A}' \mathbf{D}^{-1} \\ &\leq \frac{C \lambda_{n,F}}{\lambda_{2,L}} \times \mathbf{D}^{-1} \mathbf{A} \mathbf{D}^{-1} \mathbf{A}' \mathbf{D}^{-1}, \end{aligned}$$

where the second inequality follows from $\mathbf{L}^* \leq \mathbf{D}^{-1}/\lambda_{2,L}$. Thus, for each element, we have

$$\begin{aligned} \text{Var}(r_i) &= e_i' \text{Var}(\mathbf{r}) e_i \leq \frac{C \lambda_{n,F}}{\lambda_{2,L}} \times e_i' \mathbf{D}^{-1} \mathbf{A} \mathbf{D}^{-1} \mathbf{A}' \mathbf{D}^{-1} e_i \\ &= \frac{C \lambda_{n,F}}{d_i \lambda_{2,L} h_i}. \end{aligned}$$

Since $\mathbb{E}[r_i] = 0$ by Lemma 1, we have

$$r_i = O_p \left(\sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}} \right).$$

by Chebyshev's inequality.

Third, we bound the other error term $\mathbf{o} \equiv \mathbf{D}^{-1}\mathbf{B}'(\boldsymbol{\epsilon} - \mathbf{F}^t\boldsymbol{\tau}^t - \mathbf{F}^s\boldsymbol{\tau}^s)$. This term is mean-zero and its variance is given by:

$$\begin{aligned} \text{Var}(\mathbf{o}) &= \mathbf{D}^{-1}\mathbf{B}'\boldsymbol{\Omega}_2\mathbf{B}\mathbf{D}^{-1} \\ &\leq \tilde{\sigma}_n^2 \times \mathbf{D}^{-1}\mathbf{B}'\mathbf{B}\mathbf{D}^{-1} \end{aligned}$$

Thus, for each element, we have

$$\begin{aligned} \text{Var}(o_i) &= e_i' \text{Var}(\mathbf{o}) e_i = \tilde{\sigma}_n^2 \times e_i' \mathbf{D}^{-1}\mathbf{B}'\mathbf{B}\mathbf{D}^{-1} e_i \\ &\leq \frac{C \max_{(i,j)} |E_{(i,j)}|}{d_i}. \end{aligned}$$

for some constant $C > 0$ under Assumption 1. Thus, we have

$$o_i = O_p \left(\sqrt{\frac{\max_{(i,j)} |E_{(i,j)}|}{d_i}} \right) = o_p(1),$$

by Chebyshev's inequality as $d_i \rightarrow \infty$.

Finally, we work on the main term $\mathbf{T} \equiv \mathbf{D}^{-1}\mathbf{B}'(\mathbf{F}^t\boldsymbol{\tau}^t + \mathbf{F}^s\boldsymbol{\tau}^s)$. Observe that

$$\mathbf{T} = \left(\frac{d_i^t \tau^t(U_i) - d_i^s \tau^s(U_i)}{d_i} - \frac{1}{d_i} \sum_{e \in E_i^t} \tau^s(U_{s(e)}) + \frac{1}{d_i} \sum_{e \in E_i^s} \tau^t(U_{t(e)}) \right)_{i \in V},$$

as pointed out in the main text. Thus, it suffices to show that

$$\frac{1}{d_i} \sum_{e \in E_i^t} \tau^s(U_{s(e)}) = o_p(1), \quad \frac{1}{d_i} \sum_{e \in E_i^s} \tau^t(U_{t(e)}) = o_p(1).$$

These averages are mean-zero and each variance is given by:

$$\frac{d_i^t}{d_i^2} \mathbb{E}[\tau^s(U_j)^2] = o(1), \quad \frac{d_i^s}{d_i^2} \mathbb{E}[\tau^t(U_j)^2] = o(1),$$

as $\tau^s(U_j)/\tau^t(U_j)$ is i.i.d across $j \in V_i$. Thus, the averages converge to zero in probability by Chebyshev's inequality.

By the continuous mapping theorem, we have

$$\hat{\alpha}_i - \alpha_i = \frac{d_i^t \tau^t(U_i) - d_i^s \tau^s(U_i)}{d_i} + O_p \left(\sqrt{\frac{\lambda_{n,F}}{d_i \lambda_{2,L} h_i}} \right),$$

as $d_i \rightarrow \infty$, which completes the proof. \square

E.2. Proof of Theorem 2.

Proof. Write $\hat{\tau}_i = \tau(U_i) + r_i$ where $r_i = O_p(\sqrt{1/d_i})$ by Proposition 2 for $i \in \mathcal{C}_n$. Let F_n be the empirical distribution of $\tau(U_i)$ for $i \in \mathcal{C}_n$. Then, observe that

$$\sup_{x \in \mathbb{R}} |\hat{F}_{n,\tau}(x) - F_\tau(x)| \leq \sup_{x \in \mathbb{R}} |\hat{F}_{n,\tau}(x) - F_{n,\tau}(x)| + \sup_{x \in \mathbb{R}} |F_{n,\tau}(x) - F_\tau(x)| \quad (\text{E.1})$$

The second term in (E.1) converges to zero in probability by the Glivenko-Cantelli theorem since $\tau(U_i)$ is i.i.d across i and $|\mathcal{C}_n| \rightarrow \infty$.

For the first term in (E.1), note that for any $\delta > 0$,

$$\mathbb{I}\{\hat{\tau}_i \leq x\} = \mathbb{I}\{\tau(U_i) + r_i \leq x\} \leq \mathbb{I}\{\tau(U_i) \leq x + \delta\} + \mathbb{I}\{|r_i| > \delta\},$$

for each $i \in \mathcal{C}_n$ and $x \in \mathbb{R}$. Since F_τ is non-decreasing, we have

$$\sup_{x \in \mathbb{R}} |\hat{F}_{n,\tau}(x) - F_{n,\tau}(x)| \leq \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{I}\{|r_i| > \delta\} + \sup_{x \in \mathbb{R}} |F_{n,\tau}(x + \delta) - F_{n,\tau}(x - \delta)| \quad (\text{E.2})$$

The second term in (E.2) is bounded by

$$\sup_{x \in \mathbb{R}} |F_{n,\tau}(x + \delta) - F_{n,\tau}(x - \delta)| \leq \sup_{x \in \mathbb{R}} |F_\tau(x + \delta) - F_\tau(x - \delta)| + 2 \sup_{x \in \mathbb{R}} |F_{n,\tau}(x) - F_\tau(x)|.$$

Since F_τ is continuous, the first term is arbitrarily small for sufficiently small δ . For the first term in (E.2), note that

$$\mathbb{E}[\mathbb{I}\{|r_i| > \delta\}] = \mathbb{P}(|r_i| > \delta) \leq \frac{\mathbb{E}[|r_i|]}{\delta} = O\left(\frac{1}{\sqrt{d_i}\delta}\right).$$

Therefore, by Markov's inequality, for any $\epsilon > 0$, we have

$$\mathbb{P}\left(\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{I}\{|r_i| > \delta\} > \epsilon\right) \leq \frac{1}{\epsilon |\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{E}[\mathbb{I}\{|r_i| > \delta\}] = O\left(\frac{\sqrt{\eta_n}}{\epsilon \delta}\right) = o(1),$$

as ϵ and δ are arbitrary and $\eta_n \rightarrow 0$ by the assumption.

Combining the above, we have,

$$\begin{aligned} & \sup_{x \in \mathbb{R}} |\hat{F}_{n,\tau}(x) - F_\tau(x)| \\ & \leq 3 \sup_{x \in \mathbb{R}} |F_{n,\tau}(x) - F_\tau(x)| + \sup_{x \in \mathbb{R}} |F_\tau(x + \delta) - F_\tau(x - \delta)| + \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \mathbb{I}\{|r_i| > \delta\} \\ & = o_p(1), \end{aligned}$$

as $n \rightarrow \infty$. This completes the proof. \square

E.3. Proof of Theorem 3.

Proof. First, we show the consistency of $\hat{\sigma}_\tau^2$:

Lemma 3. *Under the stated assumptions, we have*

$$\hat{\sigma}_\tau^2 \rightarrow_p \sigma_\tau^2.$$

Proof. Observe that

$$\hat{\tau}_i = \tau(U_i) + c_i^{-1} \underbrace{\left(\frac{1}{d_i} \sum_{e \in E_i} \hat{\epsilon}_e - c_i \tau(U_i) \right)}_{\equiv \tilde{u}_i}.$$

Then,

$$\hat{\sigma}_\tau^2 = \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \tau^2(U_i) \tag{E.3}$$

$$+ \frac{2}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \tau(U_i) \tilde{u}_i \tag{E.4}$$

$$+ \frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \tilde{u}_i^2. \tag{E.5}$$

The first term (E.3) converges to $\mathbb{E}[\tau^2(U_i)]$ by the law of large numbers as $\tau(U_i)$ is i.i.d across $i \in \mathcal{C}_n$ with a finite second moment.

The third term (E.5) converges to zero in probability. To see this, from the proof of Proposition 2, observe that

$$\mathbb{E}[\tilde{u}_i^2] = O\left(\frac{1}{d_i}\right),$$

since $c_i > c$. Thus, we have

$$\mathbb{E}[(E.5)] = O(\eta_n) = o(1),$$

under the hypothesis that $\eta_n \rightarrow 0$ as $n \rightarrow \infty$. Then, by Markov's inequality, (E.5) converges to zero in probability.

The second term (E.4) converges to zero in probability by the Cauchy-Schwarz inequality:

$$(E.4) \leq \sqrt{\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \tau^2(U_i)} \sqrt{\frac{1}{|\mathcal{C}_n|} \sum_{i \in \mathcal{C}_n} \tilde{u}_i^2} = O_p(\sqrt{\eta_n}) = o_p(1),$$

as shown above. This completes the proof of Lemma 3. \square

Observe that

$$\hat{V}_\alpha^{bc} - V_\alpha = 2 \frac{\boldsymbol{\epsilon}' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \boldsymbol{\alpha}}{n} \quad (\text{E.6})$$

$$+ 2 \frac{(\boldsymbol{\epsilon} - \mathbf{F} \boldsymbol{\tau})' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \boldsymbol{\tau}}{n} \quad (\text{E.7})$$

$$+ \frac{(\boldsymbol{\epsilon} - \mathbf{F} \boldsymbol{\tau})' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' (\boldsymbol{\epsilon} - \mathbf{F} \boldsymbol{\tau})}{n} \quad (\text{E.8})$$

$$+ \frac{(\mathbb{E}[\tau^2(U_i)] - \hat{\sigma}_\tau^2) \times \text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n} \quad (\text{E.9})$$

$$+ \frac{\boldsymbol{\tau}' \mathbf{F}' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \boldsymbol{\tau}}{n} - \frac{\mathbb{E}[\tau^2(U_i)] \times \text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \mathbf{F}')}{n} \quad (\text{E.10})$$

The first term (E.6) is mean-zero and its variance is proportional to

$$\begin{aligned} \text{Var} \left(\frac{\boldsymbol{\epsilon}' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \boldsymbol{\alpha}}{n} \right) &= \frac{1}{n^2} \mathbb{E}[\boldsymbol{\epsilon}' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \boldsymbol{\alpha} \boldsymbol{\alpha}' \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \boldsymbol{\epsilon}] \\ &= \frac{1}{n^2} \text{tr}(\mathbf{B} \mathbf{L}^* \mathbf{M}_n \boldsymbol{\alpha} \boldsymbol{\alpha}' \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbb{E}[\boldsymbol{\epsilon} \boldsymbol{\epsilon}']) \\ &\leq \frac{C \lambda_{n,F}}{n^2} \times \boldsymbol{\alpha}' \mathbf{L}^* \boldsymbol{\alpha} \\ &\leq \frac{C \lambda_{n,F}}{n} \times \frac{1}{n \lambda_{2,L}} \sum_{i=1}^n \frac{\alpha_i^2}{d_i} \\ &\leq \frac{C \times C_\alpha \times \lambda_{n,F}}{n} \times \frac{1}{n \lambda_{2,L}} \sum_{i=1}^n \frac{1}{d_i} \\ &= o(1), \end{aligned}$$

where the first inequality follows from Lemma 2 for some constant $C > 0$, the second inequality follows from $\mathbf{L}^* \leq \lambda_2^{-1} \mathbf{D}^{-1}$ and $\boldsymbol{\alpha}' \mathbf{L}^* \boldsymbol{\alpha} \leq \sum_{i=1}^n \frac{\alpha_i^2}{\lambda_{2,L} d_i}$, and the third inequality follows from the assumption that $\max_{i \in V} |\alpha_i| \leq C_\alpha$ for some absolute constant $C_\alpha > 0$. Thus,

$$(\text{E.6}) = o_p(1),$$

by Chebyshev's inequality.

The second term (E.7) is mean-zero and let $\mathbf{Q} \equiv \mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}'$. Then, we can write

$$(\text{E.7}) = 2 \frac{(\boldsymbol{\epsilon} - \mathbf{F} \boldsymbol{\tau})' \mathbf{Q} \mathbf{F} \boldsymbol{\tau}}{n}$$

and its variance is given by

$$\begin{aligned}
\text{Var}((\text{E.7})) &= \frac{4}{n^2} \mathbb{E}[(\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})' \mathbf{Q} \mathbf{F} \boldsymbol{\tau} (\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})' \mathbf{Q} \mathbf{F} \boldsymbol{\tau}] \\
&= \frac{4\mathbb{E}[\tau^2(U_i)]}{n^2} \text{tr}(\mathbf{Q} \mathbf{F} \mathbf{F}' \mathbf{Q}' \boldsymbol{\Omega}_2) \\
&\quad + \frac{4}{n^2} \sum_{e_1, e_2, e_3, e_4 \in E} \mathbf{Q}_{e_1, e_2} \mathbf{Q}_{e_3, e_4} \text{cum}((\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})_{e_1}, (\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})_{e_2}, (\mathbf{F}\boldsymbol{\tau})_{e_3}, (\mathbf{F}\boldsymbol{\tau})_{e_4}),
\end{aligned}$$

where

$$\text{cum}(x_1, x_2, x_3, x_4) = \mathbb{E}[x_1 x_2 x_3 x_4] - (\mathbb{E}[x_1 x_2] + \mathbb{E}[x_1 x_3] + \mathbb{E}[x_1 x_4]).$$

The first term on the right-hand side is bounded by

$$\begin{aligned}
\frac{\mathbb{E}[\tau^2(U_i)]}{n^2} \text{tr}(\mathbf{Q} \mathbf{F} \mathbf{F}' \mathbf{Q}' \boldsymbol{\Omega}_2) &\leq \frac{\mathbb{E}[\tau^2(U_i)]}{n^2} \times \tilde{\sigma}_n^2 \times \lambda_{n,F} \times \text{tr}(\mathbf{L}^*)^2 \\
&\leq \mathbb{E}[\tau^2(U_i)] \times \tilde{\sigma}_n^2 \times \lambda_{n,F} \times \left(\frac{1}{n} \sum_{i \in V} d_i^{-1} + \frac{1}{n\lambda_{2,L}} \sum_{i \in V} d_i^{-1} h_i^{-1} \right)^2 \\
&= o(1),
\end{aligned}$$

where the first inequality follows from $\boldsymbol{\Omega}_2 \leq \tilde{\sigma}_n \mathbf{I}_m$, $\mathbf{F} \mathbf{F}' \leq \lambda_n^2 \mathbf{I}_m$, and $\text{tr}(\mathbf{Q} \mathbf{Q}') = \text{tr}(\mathbf{L}^* \mathbf{L}^*) \leq \text{tr}(\mathbf{L}^*)^2$, the second inequality follows from $e_i' \mathbf{L}^* e_i \leq 1/d_i + 1/(d_i \lambda_{2,L} h_i)$ for each $i \in V$, and the last equality follows from the hypothesis. The second term on the right-hand side is bounded by

$$\begin{aligned}
&\frac{1}{n^2} \sum_{e_1, e_2, e_3, e_4 \in E} \mathbf{Q}_{e_1, e_2} \mathbf{Q}_{e_3, e_4} \times \text{cum}((\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})_{e_1}, (\boldsymbol{\epsilon} - \mathbf{F}\boldsymbol{\tau})_{e_2}, (\mathbf{F}\boldsymbol{\tau})_{e_3}, (\mathbf{F}\boldsymbol{\tau})_{e_4}) \\
&\leq \frac{C \left(\sum_{e_1, e_2} \mathbf{Q}_{e_1, e_2} \right)^2}{n^2} \\
&= \frac{C \times \text{tr}(\mathbf{Q}^2)}{n^2} \\
&\leq \frac{C \times \text{tr}(\mathbf{L}^*)^2}{n^2} \\
&\leq C \times \left(\frac{1}{n} \sum_{i \in V} d_i^{-1} + \frac{1}{n\lambda_{2,L}} \sum_{i \in V} d_i^{-1} h_i^{-1} \right)^2 \\
&= o(1),
\end{aligned}$$

where the first inequality follows from Assumption 1 with uniformly bounded ϵ_e , the equality follows from the fact that the square of the Frobenius norm is equal to the trace

of the square of the matrix, and the rest is similar to the first term. Thus, we have

$$\text{Var}((\text{E.7})) = o(1),$$

and by Chebyshev's inequality, we have $(\text{E.7}) \rightarrow_p 0$.

The third term (E.8)'s mean is given by

$$\begin{aligned} \mathbb{E}[(\text{E.8})] &= \frac{\text{tr}(\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{\Omega}_2)}{n} \\ &\leq \frac{\tilde{\sigma}_n^2 \times \text{tr}(\mathbf{L}^*)}{n} \\ &\leq \frac{\tilde{\sigma}_n^2}{n} \times \left(\sum_{i \in V} \frac{1}{d_i} + \sum_{i \in V} \frac{1}{\lambda_{2,L} h_i d_i} \right) \\ &= o(1), \end{aligned}$$

where the second inequality follows from $e_i' \mathbf{L}^* e_i \leq 1/d_i + 1/(d_i \lambda_{2,L} h_i)$ for each $i \in V$. Since (E.8) is non-negative ($\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}$ is positive semidefinite), this implies that $(\text{E.8}) \rightarrow_p 0$.

The fourth term (E.9) converges to zero in probability by Lemma (3) and the fact that

$$\begin{aligned} \frac{\text{tr}(\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}')}{n} &\leq \frac{C\lambda_{n,F}}{n} \times \text{tr}(\mathbf{L}^*) \\ &\leq \frac{C\lambda_{n,F}}{n} \times \left(\sum_{i \in V} \frac{1}{d_i} + \sum_{i \in V} \frac{1}{\lambda_{2,L} h_i d_i} \right) \\ &= O(1), \end{aligned}$$

where the second inequality follows from $e_i' \mathbf{L}^* e_i \leq 1/d_i + 1/(d_i \lambda_{2,L} h_i)$ for each $i \in V$.

The fifth term (E.10) is mean-zero, and letting $\mathbf{Q}_\mathbf{F} \equiv \mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}$, its variance is given by

$$\begin{aligned} &\text{Var} \left(\frac{\boldsymbol{\tau}' \mathbf{F}' \mathbf{B} \mathbf{L}^* \mathbf{M}_n \mathbf{L}^* \mathbf{B}' \mathbf{F} \boldsymbol{\tau}}{n} \right) \\ &= \frac{\mathbb{E}[\tau^2(U_i)]}{n^2} \times \text{tr}(\mathbf{Q}_\mathbf{F}^2) + \frac{(\mathbb{E}[\tau^4(U_i)] - 3\mathbb{E}[\tau^2(U_i)]^2)}{n^2} \times \sum_{i \in V} \mathbf{Q}_{\mathbf{F},i,i}^2 \\ &\leq \frac{\mathbb{E}[\tau^4(U_i)]}{n^2} \times \text{tr}(\mathbf{Q}_\mathbf{F}^2), \end{aligned}$$

where the inequality follows as $\text{tr}(\mathbf{Q}_{\mathbf{F}}^2) \geq \sum_{i \in V} \mathbf{Q}_{\mathbf{F},i,i}^2$. We bound $\text{tr}(\mathbf{Q}_{\mathbf{F}}^2)$ as follows:

$$\begin{aligned}
\text{tr}(\mathbf{Q}_{\mathbf{F}}^2) &= \text{tr}(\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}) \\
&\leq \text{tr}(\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{L}^*) \\
&\leq 2\text{tr}(\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}) \\
&\quad + 2\text{tr}(\mathbf{D}^{-1}\mathbf{A}\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{A}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{A}\mathbf{L}^*\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{A}\mathbf{D}^{-1}) \\
&\leq 2\text{tr}(\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}) \\
&\quad + \frac{2\lambda_{n,F}^2}{\lambda_{2,L}^2} \text{tr}(\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}) \\
&\leq 2\text{tr}(\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}) + 2 \left(\frac{\lambda_{n,F}}{\lambda_{2,L}} \text{tr}(\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}\mathbf{A}\mathbf{D}^{-1}) \right)^2 \\
&= 2\text{tr}(\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}) + 2 \left(\frac{\lambda_{n,F}}{\lambda_{2,L}} \sum_{i \in V} \frac{1}{d_i h_i} \right)^2
\end{aligned}$$

where the first inequality follows from the fact that $\|\mathbf{M}_n\|$ is bounded by 1, the second inequality follows from $\mathbf{B}\mathbf{L}^* = \mathbf{B}\mathbf{D}^{-1} + \mathbf{B}\mathbf{D}^{-1}\mathbf{A}\mathbf{L}^*$ and the Cauchy-Schwarz inequality, the third inequality follows from $\mathbf{L}^* \leq \lambda_{2,L}^{-1}\mathbf{D}^{-1}$ and $\mathbf{F}\mathbf{F}' \leq \lambda_{n,F}\mathbf{I}_m$, and the last inequality follows from $\text{tr}(\cdot^2) \leq \text{tr}(\cdot)^2$. For the first term on the far right-hand side, we have

$$\begin{aligned}
&\text{tr}(\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}\mathbf{D}^{-1}\mathbf{B}'\mathbf{F}\mathbf{F}'\mathbf{B}\mathbf{D}^{-1}) \\
&= \sum_{i \in V} \left(\frac{(d_i^t - d_i^s)^2 + \sum_{j \neq i} r_{i,j}^2}{d_i^2} \right)^2 + \sum_{i \neq j} \left(\frac{\sum_{k \in V} r_{i,k} r_{j,k}}{d_i d_j} \right)^2 \\
&\leq Cn(1 + \sum_{i \in V} d_i^{-2})
\end{aligned}$$

for some constant $C > 0$, where $r_{i,j} = \sum_{e \in E(i,j)} \mathbf{B}_{e,i}$ and the inequality follows from $\sum_{j \neq i} r_{i,j}^2 \leq C_A d_i$, $\sum_{k \in V} r_{i,k} r_{j,k} \leq C_A^2 d_i$. Thus,

$$\begin{aligned}
&\text{Var} \left(\frac{\boldsymbol{\tau}'\mathbf{F}'\mathbf{B}\mathbf{L}^*\mathbf{M}_n\mathbf{L}^*\mathbf{B}'\mathbf{F}\boldsymbol{\tau}}{n} \right) \\
&\leq C \left(\frac{1}{n} + \frac{1}{n} \sum_{i \in V} d_i^{-2} \right) + C \left(\frac{\lambda_{n,F}}{\lambda_{2,L}} \frac{1}{n} \sum_{i \in V} \frac{1}{d_i h_i} \right)^2 \\
&= o(1),
\end{aligned}$$

under the stated assumptions. Thus, by Chebyshev's inequality, we have (E.10) $\rightarrow_p 0$.

Finally, we combine all the results above to conclude that

$$\hat{V}_{\alpha}^{bc} - V_{\alpha} \rightarrow_p 0.$$

This completes the proof. □

E.4. Proof of Proposition 1.

Proof. First, observe that

$$\mathbf{\Omega}_1 \leq \sigma_{\tau}^2 \mathbf{F} \mathbf{F}' \leq \sigma_{\tau}^2 \lambda_{n,F} \mathbf{I}_m,$$

where the first inequality holds by $\mathbf{F} = \mathbf{F}^t + \mathbf{F}^s$. Then, the first term in the variance formula (2.7) is bounded by:

$$\mathbf{L}^* \mathbf{B}' \mathbf{\Omega}_1 \mathbf{B} \mathbf{L}^* \leq \sigma_{\tau}^2 \lambda_{n,F} \mathbf{L}^*,$$

where we used the fact that \mathbf{L}^* is a pseudo-inverse of $\mathbf{B}' \mathbf{B}$.

Next, since $\mathbf{\Omega}_2 \leq \tilde{\sigma}_n^2 \mathbf{I}_m$, the second term in the variance formula (2.7) is bounded by:

$$\mathbf{L}^* \mathbf{B}' \mathbf{\Omega}_2 \mathbf{B} \mathbf{L}^* \leq \tilde{\sigma}_n^2 \mathbf{L}^*.$$

Therefore, the variance of the fixed effect is bounded by:

$$\text{Var}(\hat{\alpha}) \leq (\tilde{\sigma}_n^2 + \sigma_{\tau}^2 \lambda_{n,F}) \times \mathbf{L}^*.$$

Thus, for a unit vector e_i with 1 at the i -th position and 0 elsewhere, we have:

$$\text{Var}(\hat{\alpha}_i) = e_i' \text{Var}(\hat{\alpha}) e_i \leq (\tilde{\sigma}_n^2 + \sigma_{\tau}^2 \lambda_{n,F}) \times e_i' \mathbf{L}^* e_i.$$

By the proof of Theorem 2 in Jochmans and Weidner (2019), we have

$$e_i' \mathbf{L}^* e_i \leq \frac{1}{d_i} \left(1 + \frac{1}{\lambda_{2,L} h_i} \right) - \frac{2}{m}.$$

Thus,

$$\text{Var}(\hat{\alpha}_i) \leq (\tilde{\sigma}_n^2 + \sigma_{\tau}^2 \lambda_{n,F}) \times \left\{ \frac{1}{d_i} \left(1 + \frac{1}{\lambda_{2,L} h_i} \right) - \frac{2}{m} \right\},$$

which completes the proof. □

E.5. Proof of Proposition 2.

Proof. Observe that from the transformed model (B.1), we can write the residuals as

$$\begin{aligned}\hat{\epsilon} &= \mathbf{y} - \mathbf{B}\hat{\alpha} = (\mathbf{I}_m - \mathbf{B}(\mathbf{B}'\mathbf{B})^*\mathbf{B}')\mathbf{y} = \mathbf{M}_\mathbf{B}\mathbf{y} \\ &= \mathbf{M}_\mathbf{B}\tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) + \mathbf{M}_\mathbf{B}\tilde{\epsilon}.\end{aligned}$$

Recall that the statistic of interest minus $c_i(\tau^t(U_i) + \tau^s(U_i))/2$ can be written as

$$\begin{aligned}&\frac{1}{d_i} \sum_{e \in E_i} \hat{\epsilon}_e - c_i \frac{\tau^t(U_i) + \tau^s(U_i)}{2} \\ &= \frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\epsilon}\end{aligned}\tag{E.11}$$

$$+ \frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) - c_i \frac{\tau^t(U_i) + \tau^s(U_i)}{2}\tag{E.12}$$

where $\mathbf{1}_{i,m}$ is the m -dimensional vector with $1_{i,m,e} = 1$ if $e \in E_i$ and 0 otherwise.

First, we show that (E.12) converges to zero in probability. By the definition of $\tilde{\epsilon}$, this term has mean zero. Its variance is given by

$$\begin{aligned}\text{Var} \left(\frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\epsilon} \right) &= \frac{1}{d_i^2} \mathbb{E} [\mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\epsilon} \tilde{\epsilon}' \mathbf{M}_\mathbf{B} \mathbf{1}_{i,m}] \\ &\leq \frac{\tilde{\sigma}_n^2}{d_i},\end{aligned}$$

where the inequality follows from $\|\mathbf{M}_\mathbf{B}\| \leq 1$ (since $\mathbf{M}_\mathbf{B}$ is idempotent), $\mathbb{E}[\tilde{\epsilon}\tilde{\epsilon}'] \leq \tilde{\sigma}_n^2 \mathbf{I}_m$, and $\mathbf{1}'_{i,m} \mathbf{1}_{i,m} = d_i$. Since $d_i \rightarrow \infty$, the variance converges to zero. Thus, by Chebyshev's inequality,

$$\frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\epsilon} = O_p \left(\frac{1}{\sqrt{d_i}} \right) = o_p(1),$$

as $d_i \rightarrow \infty$.

Second, we show that (E.11) converges in probability to zero. For $i \notin V^{st}$, this term is exactly zero since either $\mathbf{1}_{i,m} = \mathbf{b}_i$ or $\mathbf{1}_{i,m} = -\mathbf{b}_i$, where \mathbf{b}_i is the i -th column of \mathbf{B} ; as there are no in- or outflows for i , $\mathbf{1}_{i,m}$ is orthogonal to $\mathbf{M}_\mathbf{B}$ and thus the first term vanishes. Also, note that $c_0 = 0$ in this case.

For $i \in V^{st}$, since $\mathbb{E}[\tau^o(U_j)|U_i] = 0$ for $j \neq i$ and $o \in \{s, t\}$, we have

$$\mathbb{E} \left[\frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) - c_i \frac{\tau^t(U_i) + \tau^s(U_i)}{2} \middle| U_i \right] = 0$$

where $\tilde{\mathbf{f}}_i$ is the i -th column of $\tilde{\mathbf{F}}$, and $c_i = \frac{1}{d_i} \mathbf{1}'_{i,m} \mathbf{M}_\mathbf{B} \tilde{\mathbf{f}}_i$. Note $c_i \in [0, 1]$ since $\mathbf{M}_\mathbf{B}$ is a projection matrix and $\mathbf{1}'_{i,m} \tilde{\mathbf{f}}_i = d_i$.

Since

$$\begin{aligned} & \frac{1}{d_i} 1'_{i,m} \mathbf{M}_B \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) \\ &= \frac{1}{2d_i} \sum_{j \in V^{st}} (\tau^t(U_j) + \tau^s(U_j)) \cdot 1'_{i,m} \mathbf{M}_B \tilde{\mathbf{f}}_j, \end{aligned}$$

the conditional variance given U_i is

$$\begin{aligned} \text{Var} \left[\frac{1}{d_i} 1'_{i,m} \mathbf{M}_B \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) | U_i \right] &= \frac{\mathbb{E}[(\tau^t(U_i) + \tau^s(U_i))^2]}{4d_i^2} \sum_{j \neq i \in V^{st}} \left(1'_{i,m} \mathbf{M}_B \tilde{\mathbf{f}}_j \right)^2 \\ &\leq \frac{\mathbb{E}[(\tau^t(U_i) + \tau^s(U_i))^2]}{4d_i^2} \sum_{j \in V^{st}} |E_{(i,j)}|^2 \\ &\leq \frac{\max_{j \in V^{st}} |E_{(i,j)}|^2 \cdot \mathbb{E}[(\tau^t(U_i) + \tau^s(U_i))^2]}{4d_i} \\ &= O\left(\frac{1}{d_i}\right) = o(1), \end{aligned}$$

as $d_i \rightarrow \infty$, where the first inequality uses $\|\mathbf{M}_B\| \leq 1$ and $1'_{i,m} \tilde{\mathbf{f}}_j = |E_{(i,j)}|$, and the last equality follows by assumption. By the law of total variance and

$$\text{Var} \left[\frac{1}{d_i} 1'_{i,m} \mathbf{M}_B \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) - c_i \frac{\tau^t(U_i) + \tau^s(U_i)}{2} \right] = O\left(\frac{1}{d_i}\right) = o(1),$$

Thus, by Chebyshev's inequality,

$$\frac{1}{d_i} 1'_{i,m} \mathbf{M}_B \tilde{\mathbf{F}}(\tilde{\tau}^t/2 + \tilde{\tau}^s/2) - c_i \frac{\tau^t(U_i) + \tau^s(U_i)}{2} = O_p\left(\frac{1}{\sqrt{d_i}}\right) = o_p(1),$$

as $d_i \rightarrow \infty$.

Finally, combining the results from the two steps, we have

$$\frac{1}{d_i} \sum_{e \in E_i} \hat{\epsilon}_e - c_i(\tau^t(U_i) + \tau^s(U_i))/2 = O_p\left(\frac{1}{\sqrt{d_i}}\right) = o_p(1),$$

as $d_i \rightarrow \infty$ by the continuous mapping theorem. This completes the proof. \square

E.6. Proof of Proposition 3.

Proof. Validity of the confidence interval: First, we show that $\hat{c}_\alpha \rightarrow_p c_\alpha$, where c_α is the α -quantile of F_τ for any $\alpha \in (0, 1)$ such that $c_\alpha < \infty$. This is immediate from the uniform convergence of $\hat{F}_{n,\tau}$ to F_τ in Theorem 2: Since $\hat{F}_{n,\tau}$ weakly converges to F_τ in probability (Problem 23.1 in van der Vaart, 1998), $\hat{F}_{n,\tau}^{-1}(\cdot) = \inf\{x \in \mathbb{R} : \hat{F}_{n,\tau}(x) \geq \cdot\}$ weakly converges $F_\tau^{-1}(\cdot) = \inf\{x \in \mathbb{R} : F_\tau(x) \geq \cdot\}$ in probability (Lemma 21.2 in van der

Vaart, 1998). Since F_τ is continuous and strictly increasing around c_α , $F_\tau^{-1}(\alpha) = c_\alpha$ and $\hat{c}_\alpha = \hat{F}_{n,\tau}^{-1}(\alpha) \rightarrow_p c_\alpha$.

Next, observe that by Theorem 1 and $((d_i^t - d_i^s)/d_i)^{-1} = O(1)$,

$$\tilde{T} \equiv \left(\frac{d_i^t - d_i^s}{d_i} \right)^{-1} (\hat{\alpha}_i - \alpha_i) \rightarrow_p \tau(U_i),$$

as $d_i \rightarrow \infty$. Thus, $(\tilde{T}, \hat{c}_{\alpha/2}, \hat{c}_{1-\alpha/2})$ weakly converges to $(\tau(U_i), c_{\alpha/2}, c_{1-\alpha/2})$. By Slutsky's theorem, we have

$$\mathbb{P}(\hat{c}_{\alpha/2} \leq \tilde{T} \leq \hat{c}_{1-\alpha/2}) \rightarrow \mathbb{P}(c_{\alpha/2} \leq \tau(U_i) \leq c_{1-\alpha/2}) = F_\tau(c_{1-\alpha/2}) - F_\tau(c_{\alpha/2}) = 1 - \alpha.$$

This shows the asymptotic validity of the confidence interval.

Validity of the test: Observe that under the null hypothesis, by Theorem 1 and the continuous mapping theorem, we have

$$\begin{aligned} T &= \underbrace{((\tau(U_i))_{i \in V_0})' \mathbf{C}_n ((\tau(U_i))_{i \in V_0})}_{\equiv g_n((\tau(U_i))_{i \in V_0})} + o_p(1) \\ &\rightarrow_d \underbrace{((\tau(U_i))_{i \in V_0})' \mathbf{C} ((\tau(U_i))_{i \in V_0})}_{\equiv g((\tau(U_i))_{i \in V_0})} = T_\infty \end{aligned}$$

where $\mathbf{C} = \mathbf{D}_{V_0}^{st} \mathbf{M}_{n_0} \mathbf{D}_{V_0}^{st}$ and $\mathbf{D}_{V_0}^{st} = \lim_{n \rightarrow \infty} \mathbf{D}_{n,V_0}^{st}$. Note that the cdf of T_∞ is continuous. Let $\tau_i^{(m)}$ be m -th random draw from $\hat{F}_{n,\tau}$ (with replacement) for $m = 1, \dots, M$. Then, the simulated test statistic is given by

$$\tilde{T}^{(m)} = \underbrace{(\tau_i^{(m)})_{i \in V_0}' \mathbf{C}_n ((\tau_i^{(m)})_{i \in V_0})}_{\equiv g_n((\tau_i^{(m)})_{i \in V_0})},$$

$\mathbf{C}_n = \mathbf{D}_{n,V_0}^{st} \mathbf{M}_{n_0} \mathbf{D}_{n,V_0}^{st}$ and $\mathbf{D}_{n,V_0}^{st} = \text{diag}((d_i^t - d_i^s)/d_i, i \in V_0)$. Note that $\tilde{T}^{(m)}, m = 1, \dots, M$ are independent and identically distributed conditional on $\hat{F}_{n,\tau}$.

Let $\hat{F}_{n,T}^{\otimes n_0}$ be the distribution of $(\tau_i^{(1)})_{i \in V_0}$ given $\hat{F}_{n,\tau}$ and $F_\tau^{\otimes n_0}$ be the distribution of $(\tau(U_i))_{i \in V_0}$. Since n_0 is fixed, by the weak convergence of $\hat{F}_{n,\tau}$ to F_τ in probability, $\hat{F}_{n,T}^{\otimes n_0}$ weakly converges to $F_\tau^{\otimes n_0}$ in probability (Theorem 2.8 in Billingsley, 1999). As g_n and g are continuous, by the continuous mapping theorem (Theorem 3.27 in Kallenberg, 2002), $\hat{G}_{n,T}$ weakly converges to G_T in probability, where $\hat{G}_{n,T}$ is the distribution of $\tilde{T}^{(1)}$ given $\hat{F}_{n,\tau}$ and G_T is the distribution of T_∞ . Since G_T is continuous, by the same argument as in the proof of the validity of the confidence interval, we have $\hat{c}_\alpha \rightarrow_p c_\alpha$, where \hat{c}_α is the α -quantile of $\hat{G}_{n,T}$ and c_α is the α -quantile of G_T . Thus, (T, \hat{c}_α) weakly converges to

(T_∞, c_α) . By Slutsky's theorem, we have

$$\mathbb{P}(T > \hat{c}_{1-\alpha}) \rightarrow \mathbb{P}(T_\infty > c_{1-\alpha}) = 1 - G_T(c_{1-\alpha}) = \alpha,$$

which shows the asymptotic validity of the test with the known $\hat{G}_{n,T}$.

Finally, we show that the test remains valid when $\hat{G}_{n,T}$ is estimated by the empirical distribution of $\tilde{T}^{(m)}$, $m = 1, \dots, M$. Let $\hat{G}_{n,T,M}$ be the empirical distribution of $\tilde{T}^{(m)}$, $m = 1, \dots, M$ given $\hat{F}_{n,\tau}$. Let $\hat{c}_{\alpha,M}$ be the α -quantile of $\hat{G}_{n,T,M}$. Observe that, for any $\epsilon > 0$,

$$\begin{aligned} & |\mathbb{P}(T > \hat{c}_{1-\alpha,M}) - \mathbb{P}(T > c_{1-\alpha})| \\ &= \left| \mathbb{P}(\hat{G}_{n,T,M}(T) > 1 - \alpha) - \mathbb{P}(G_T(T) > 1 - \alpha) \right| \\ &\leq \mathbb{P}(|G_T(T) - 1 + \alpha| \leq \epsilon) + \mathbb{P}\left(|G_T(T) - \hat{G}_{n,T,M}(T)| > \epsilon\right) \\ &\leq \mathbb{P}(|G_T(T) - 1 + \alpha| \leq \epsilon) + \mathbb{P}\left(\sup_{x \in \mathbb{R}} |G_T(x) - \hat{G}_{n,T,M}(x)| > \epsilon\right), \end{aligned}$$

where the first equality follows from $\hat{G}_{n,T,M}$ and G_T being right-continuous, and the first inequality follows from the fact that $|\mathbb{I}(x \leq a) - \mathbb{I}(x \leq b)| \leq \mathbb{I}(|x - a| \leq \epsilon) + \mathbb{I}(|a - b| > \epsilon)$ for any $x, a, b \in \mathbb{R}$ and $\epsilon > 0$. By the continuous mapping theorem and the Portmanteau theorem, the first term on the right-hand side converges to $\mathbb{P}(|G_T(T) - 1 + \alpha| \leq \epsilon)$, which can be made arbitrarily small by choosing ϵ small enough. For the second term on the right-hand side, observe that

$$\sup_{x \in \mathbb{R}} |G_T(x) - \hat{G}_{n,T,M}(x)| \leq \sup_{x \in \mathbb{R}} |G_T(x) - \hat{G}_{n,T}(x)| + \sup_{x \in \mathbb{R}} |\hat{G}_{n,T}(x) - \hat{G}_{n,T,M}(x)|.$$

Since G_T is continuous, the first term on the right-hand side converges to zero in probability (Problem 23.1 in van der Vaart, 1998). The second term on the right-hand side converges to zero almost surely by the Glivenko-Cantelli theorem as $M \rightarrow \infty$ conditionally on $\hat{F}_{n,\tau}$. Thus,

$$\mathbb{P}\left(\sup_{x \in \mathbb{R}} |G_T(x) - \hat{G}_{n,T,M}(x)| > \epsilon\right) \rightarrow 0,$$

as $n, M \rightarrow \infty$. Since $\epsilon > 0$ is arbitrary, we have

$$\mathbb{P}(T > \hat{c}_{1-\alpha,M}) \rightarrow \mathbb{P}(T > c_{1-\alpha}) = \alpha,$$

as $n, M \rightarrow \infty$. This completes the proof. \square