

FIXED-EFFECT REGRESSIONS ON NETWORK DATA

Koen Jochmans*
University of Cambridge

Martin Weidner[‡]
University College London

April 1, 2019

Abstract

This paper considers inference on fixed effects in a linear regression model estimated from network data. An important special case of our setup is the two-way regression model. This is a workhorse technique in the analysis of matched data sets, such as employer-employee or student-teacher panel data. We formalize how the structure of the network affects the accuracy with which the fixed effects can be estimated. This allows us to derive sufficient conditions on the network for consistent estimation and asymptotically-valid inference to be possible. Estimation of moments is also considered. We allow for general networks and our setup covers both the dense and sparse case. We provide numerical results for the estimation of teacher value-added models and regressions with occupational dummies.

Keywords: connectivity, fixed effects, graph, Laplacian, limited mobility, teacher value-added, two-way regression model.

JEL classification: C23, C55

*Address: University of Cambridge, Faculty of Economics, Austin Robinson Building, Sidgwick Avenue, Cambridge CB3 9DD, U.K. E-mail: kj345@cam.ac.uk.

[‡]Address: University College London, Department of Economics, Gower Street, London WC1E 6BT, U.K., and CeMMAP. E-mail: m.weidner@ucl.ac.uk.

We are grateful to Ulrich Müller and three referees for constructive comments. We would also like to thank Nadine Geiger, Bryan Graham, Áureo de Paula, Fabien Postel-Vinay, Valentin Verdier, and, in particular, Jean-Marc Robin for discussion and for comments on earlier drafts of this paper. Valentin Verdier also provided indispensable support that enabled us to include our illustration on the estimation of teacher value-added.

Jochmans gratefully acknowledges financial support from the European Research Council through Starting Grant n° 715787. Weidner gratefully acknowledges support from the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice grant RES-589-28-0001 and from the European Research Council grant ERC-2014-CoG-646917-ROMIA.

The first version of this paper dates from August 4, 2016. All versions are available at <https://arxiv.org/abs/1608.01532>.

1 Introduction

Data on the interaction between agents are in increasing supply. A workhorse technique to analyze such data is a linear regression model with agent-specific parameters. It has been used to investigate a variety of questions. For example, application of a two-way regression model to matched employer-employee data decomposes (log) wages into worker heterogeneity, firm heterogeneity, and residual variation. Following [Abowd, Kramarz and Margolis \(1999\)](#) the correlation between the estimated worker and firm effects is regarded as a measure of assortative matching. A positive correlation indicates that high quality workers are employed in more productive firms. Using the same decomposition, [Card, Heining and Kline \(2013\)](#) study to what extent the evolution of wage inequality is due to changes in the variance of worker and firm heterogeneity. [Nimczik \(2018\)](#) reports the whole distribution of the estimated worker and firm effects. In a similar fashion, the literature on student achievement backs out student and teacher effects from test score data. The estimated teacher heterogeneity is interpreted as teacher value-added and their variance as a measure of their importance (see [Jackson, Rockoff and Staiger 2014](#) for an overview of this literature). These estimates are used to assess teachers and are important inputs to personnel evaluations and merit pay programs ([Rothstein, 2010](#)).¹

In spite of their widespread use, there is little to no work on the theoretical properties of such fixed-effect approaches. In fact, the few results that are available point to issues of downward bias in the estimation of the correlation between worker and firm effects, finding a spurious negative correlation in many data sets ([Andrews, Gill, Schank and Upward, 2008, 2012](#)), and upward bias in the estimator of the variance of teacher effects ([Rockoff, 2004](#)). The presence of bias here is not surprising. Indeed, the individual effects are estimated with

¹Fixed-effect regressions of this kind are now part of the standard toolkit of many empiricists in a variety of different areas. [Finkelstein, Gentzkow and Williams \(2016\)](#) and [Amiti and Weinstein \(2018\)](#) use them to separate supply and demand factors in healthcare utilization from data on patient migration, and in firm investment behavior from financial data on banks loans, respectively. [Chetty and Hendren \(2018\)](#) evaluate the importance of growing up in a specific neighborhood on labor market outcomes later on in life.

noise. Their sampling error then introduces bias in the estimator of nonlinear functionals. A more complicated issue is the assessment of the statistical precision with which the fixed effects are estimated and, more generally, the development of distribution theory. This is important as it allows to establish conditions for consistency and rates of convergence, and yields insight into whether standard test statistics can be expected to be approximately size correct and have non-trivial power. None of these issues has been addressed so far. Providing such theory so is not only relevant for inference on the individual effects and their moments, but may also serve as a stepping stone to address related problems. For example, without theory for the fixed-effect estimator the behavior of the falsification test for value-added models of [Rothstein \(2010\)](#) remains unknown, and correct standard errors for regressions of outcomes on estimated fixed effects ([Kettemann, Mueller and Zweimüller, 2017](#)) cannot be derived.

The data structure arising from interactions between agents is different from that of standard cross-section or panel data. It is typically difficult to see how the data carry information about certain parameters. In this paper we present sufficient conditions for consistency and asymptotic normality of least-squares estimators of fixed effects in linear regression models. We see the data as a network and represent it by a graph where agents are vertices and edges between vertices are present if these agents interact. It is intuitive that the structure of this graph should be a key determinant of the accuracy of statistical inference. We formalize this here. Our setup places no a priori restrictions on the graph structure and our results apply to both dense and sparse settings. A data structure of particular importance is that of a bipartite graph. Here, the data concerns two types of individuals and interactions only occur between the types but not within each type. This is the case in our motivating examples above and we treat this bipartite case in detail. In fact, while we deal with general graphs, our regression setup is designed to capture the main features of the prototypical two-way regression model. We focus on inference on the individual effects but our results also serve as a stepping stone for the analysis of estimators of other parameters, such as the variance and other moments of (the distribution of) the individual effects, and we provide some results on these as well. We do not discuss inference

on common slope coefficients. In contemporary work, [Verdier \(2018\)](#) provides such results for two-way regression models. The fixed-effect model for test scores, for example, can be used to assess the effect of class-size reductions on student achievement while controlling for student and teacher heterogeneity.

The ability to accurately estimate the individual effect of a given vertex depends on how well this vertex is connected to the rest of the network. Our theory involves both global and local measures of network connectivity. The main global connectivity measure we use is the smallest non-zero eigenvalue of the (normalized) Laplacian matrix of the graph.² It reflects how easy it is to disconnect a network by removing edges from it. The other measures of connectivity that we use are the degrees of the vertices as well as various harmonic means thereof. All of these measures arise naturally when studying the variance of the fixed-effect estimator. We highlight the interplay between them in deriving conditions for consistent estimation and for inference based on standard first-order asymptotics to be possible. As the network grows the smallest eigenvalue may approach zero, and so the graph may become more sparse, provided the relevant harmonic mean grows sufficiently fast. These findings mimic conditions on the bandwidth in nonparametric estimation problems, although they will typically show up in second-order terms here. This explains why estimation at the parametric rate may be feasible even in sparse networks. Our results also show that inference on averages over the individual effects is more demanding on the network structure and, even after bias reduction, may only be feasible in quite dense networks.

Our analysis shows that it is useful to inspect measures of global and local connectivity when interpreting estimation results from network data. We do so here for two data sets. The first is a large network of teachers in elementary schools in North Carolina, where

²The Laplacian matrix is similar to the adjacency matrix as a device to represent a graph and can be obtained from it. Both matrices are formally defined below. Eigenvalues and eigenvectors of these and related matrices have also been found of use in determining equilibrium conditions in games played on networks ([Bramoullé, Kranton and D’Amours, 2014](#)) and in (statistical) community detection ([Schiebinger, Wainwright and Yu, 2015](#)).

the object of interest would be teacher value-added. This is arguably one of the most important applications of the two-way regression model. This graph is only very weakly connected and our theory does not support the use of large-sample arguments. When a simple model with homoskedastic errors is applied to these data standard errors based on conventional first-order approximations for teacher value-added are, on average, about 40% smaller than the actual standard deviations. Further, the sample variance of the estimated teacher effects has a substantial upward bias. This bias translates into an overly optimistic view on the ability of teacher value-added to explain variation in test scores. To provide an example of a data set that yields a much stronger connected graph we also construct an occupational network from the British Household Panel Survey (BHPS). This graph would arise in the context of wage regressions with occupational dummies, for example. Here, our connectivity measures are much more supportive of standard inferential approaches and, indeed, again in a simple model, we find that conventional first-order approximations are quite accurate.

The remainder of this paper is organized as follows. Section 2 details the structure of the data under study and introduces the regression model of interest. Special attention is given to the bipartite graph and the two-way regression model. Section 3 provides distribution theory for the least-squares estimator of the individual effects and also discusses estimation of their moments. Section 4 contains details on our two numerical illustrations. Section 5 concludes. The supplementary material to this paper contains some additional results and illustrations, as well as the proofs of all theorems.

2 Regression analysis of network data

2.1 Data structure

Consider an undirected graph $\mathcal{G} := \mathcal{G}(V, E)$ where $m := |E|$ edges are placed between $n := |V|$ vertices. We allow for multiple edges between vertices (i.e., \mathcal{G} can be a multigraph) and the edges may be assigned a weight. We do not consider loops (i.e., no edge connects

a vertex with itself). Without loss of generality we label the vertices by natural numbers, so that V is $\{1, \dots, n\}$. The multiset E contains the $m > 0$ unordered pairs (i, j) from the product set $V \times V$ that are linked by an edge, possibly with repetition. The same pair (i, j) will appear multiple times in E if they share more than one edge; we let $E_{(i,j)} \subset E$ denote the set of edges between them. We have $E_{(i,j)} = E_{(j,i)}$ and may have $E_{(i,j)} = \emptyset$. We will label the edges by natural numbers; so, each edge e edge has assigned to it an integer $\varepsilon_e \in \{1, \dots, m\}$. For later use we note that vertices i and j are said to be connected if \mathcal{G} contains a path from i to j , and that the graph \mathcal{G} is said to be connected if every pair of vertices in the graph is connected.

For an edge $e \in E$ let $w_e > 0$ be its weight. An unweighted graph has $w_e = 1$ for all $e \in E$. The graph \mathcal{G} may be represented by its $m \times n$ (oriented) incidence matrix \mathbf{B} , with entries

$$(\mathbf{B})_{\varepsilon_e i} := \begin{cases} \sqrt{w_e} & \text{if } e \in E_{(i,j)} \text{ for some } j \in V \text{ and } i < j, \\ -\sqrt{w_e} & \text{if } e \in E_{(i,j)} \text{ for some } j \in V \text{ and } i > j, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Here, the choice of sign gives each edge e an orientation. As will become apparent, our analysis below is invariant to this choice of orientation. The graph may also be represented through its $n \times n$ adjacency matrix \mathbf{A} , which has elements

$$(\mathbf{A})_{ij} := \sum_{e \in E_{(i,j)}} w_e.$$

The incidence matrix and adjacency matrix are related through the $n \times n$ Laplacian matrix \mathbf{L} , as

$$\mathbf{L} := \mathbf{B}'\mathbf{B} = \mathbf{D} - \mathbf{A}, \quad (2)$$

for $\mathbf{D} := \text{diag}(d_1, \dots, d_n)$ the diagonal $n \times n$ (weighted) degree matrix, where the degree of vertex i is

$$d_i := \sum_{j=1}^n (\mathbf{A})_{ij}.$$

When \mathcal{G} is an unweighted graph, for example, d_i equals the number of edges that involve vertex i . For a vertex i we will let $[i] := \{j \in V : E_{(i,j)} \neq \emptyset\}$ denote the set of its direct

neighbors. Observe that d_i may be large even if i has few neighbors—i.e, when $|[i]|$ is small—as the edge weights $(\mathbf{A})_{ij}$ for $j \in [i]$ may be large. An example is a multigraph where many edges exist between i and some $j \in [i]$.

2.2 Regression model and least-squares estimator

Now, given a graph \mathcal{G} , for each edge $e \in E$ we observe an outcome y_{ε_e} and a p -vector of covariates $\mathbf{x}_{\varepsilon_e}$. Allowing \mathcal{G} to be a multigraph covers the (unbalanced) panel data case, where multiple outcomes are available for some vertex pairs. Collect all outcomes in the m -vector \mathbf{y} and all covariates in the $m \times p$ matrix \mathbf{X} . Let $\boldsymbol{\alpha} := (\alpha_1, \dots, \alpha_n)'$ be an n -vector of vertex-specific parameters and let $\boldsymbol{\beta} := (\beta_1, \dots, \beta_p)'$ be a p -vector of regression slopes. Our interest lies in estimating the model

$$\mathbf{y} = \mathbf{B}\boldsymbol{\alpha} + \mathbf{X}\boldsymbol{\beta} + \mathbf{u}, \tag{3}$$

where \mathbf{u} is an m -vector of regression errors.³ We will treat \mathbf{B} and \mathbf{X} as fixed throughout. This implies that we consider the network as non-random and exogenous.⁴ Our focus is on the vector $\boldsymbol{\alpha}$. In the two-way regression model of our motivating examples these are the worker and firm effects or the student and teacher effects, respectively. By the definition of \mathbf{B} one of these effects will enter (3) with a minus sign. While this may appear to be an unusual convention from an applied perspective, it is convenient for our theoretical analysis. As will be explained below, this sign convention is without loss of generality in the two-way model, where the underlying graph is bipartite.

In (3) the outcomes for a given vertex pair (i, j) depend on the individual effects through their difference $\alpha_i - \alpha_j$. This implies that our model is overparameterized. Indeed, we have $\mathbf{B}\boldsymbol{\iota}_n = \mathbf{0}$, where $\boldsymbol{\iota}_n := (1, \dots, 1)'$ is the n -vector of ones, as each row of \mathbf{B} sums up to zero.

³A change of edge orientation corresponds to a sign flip in the corresponding outcome and regressor matrices. This does not affect the least squares estimator of (3).

⁴Exogeneity of the network is the standard assumption in the literature building on [Abowd, Kramarz and Margolis \(1999\)](#). Accounting for endogenous network formation requires more complicated models and has started to receive some attention; see [Bonhomme, Lamadon and Manresa \(2018\)](#) and [Lentz, Piyapromdee and Robin \(2018\)](#).

It follows that the mean of the vertex-specific parameters cannot be learned from the data and a normalization is required. We impose that

$$\sum_{i=1}^n \sum_{j=1}^n (\mathbf{A})_{ij} (\alpha_i + \alpha_j) = 0, \quad (4)$$

which will prove a convenient choice for our purposes. Denoting the degree vector by $\mathbf{d} := (d_1, \dots, d_n)'$ we may write the constraint (4) compactly as $\mathbf{d}'\boldsymbol{\alpha} = 0$. A normalization can be dispensed with if interest lies in parameter differences, i.e. $\alpha_i - \alpha_j$, as in [Finkelstein, Gentzkow and Williams \(2016\)](#), for example. Results for such differences that parallel those developed below are given in the supplementary material.

The standard estimator of $\boldsymbol{\alpha}$ is the constrained least-squares estimator

$$\check{\boldsymbol{\alpha}} := (\check{\alpha}_1, \dots, \check{\alpha}_n)' = \underset{\mathbf{a} \in \{\mathbf{a} \in \mathbb{R}^n : \mathbf{d}'\mathbf{a} = 0\}}{\operatorname{arg\,min}} \|\mathbf{M}_X \mathbf{y} - \mathbf{M}_X \mathbf{B} \mathbf{a}\|^2, \quad (5)$$

where $\|\cdot\|$ denotes the Euclidean norm, $\mathbf{M}_X := \mathbf{I}_m - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$, and \mathbf{I}_m is the identity matrix of dimension $m \times m$. The following theorem gives conditions under which this estimator exists and is unique. For any square matrix \mathbf{C} we denote its Moore-Penrose pseudoinverse by \mathbf{C}^\dagger . When \mathbf{C} is $n \times n$ we let $\mathbf{C}^* := \mathbf{D}^{-1/2} (\mathbf{D}^{-1/2} \mathbf{C} \mathbf{D}^{-1/2})^\dagger \mathbf{D}^{-1/2}$. It is easily shown that $\mathbf{C} \mathbf{C}^* \mathbf{C} = \mathbf{C}$ and $\mathbf{C}^* \mathbf{C} \mathbf{C}^* = \mathbf{C}^*$. Therefore, \mathbf{C}^* is a pseudoinverse of \mathbf{C} .

Theorem 1 (Existence). *Let \mathcal{G} be connected, $\operatorname{rank}(\mathbf{X}) = p$, and $\operatorname{rank}((\mathbf{X}, \mathbf{B})) = p + n - 1$. Then*

$$\check{\boldsymbol{\alpha}} = (\mathbf{B}' \mathbf{M}_X \mathbf{B})^* \mathbf{B}' \mathbf{M}_X \mathbf{y}$$

and is unique.

The need for a pseudoinverse arises because $\mathbf{B}' \mathbf{M}_X \mathbf{B}$ is singular, which follows from the fact that $\mathbf{B} \mathbf{t}_n = 0$. The use of the particular pseudoinverse $(\mathbf{B}' \mathbf{M}_X \mathbf{B})^*$ is a consequence of our normalization $\mathbf{d}'\boldsymbol{\alpha} = 0$. A change of normalization would imply a different pseudoinverse in the statement of Theorem 1. The result of the theorem is intuitive and generalizes results in the literature on matched employer-employee data ([Abowd, Creedy and Kramarz 2002](#)). When the graph \mathcal{G} is disconnected a separate normalization of the form in (4) is needed

for each connected component of \mathcal{G} . Our results then apply to each of these subgraphs. In practice, the analysis is typically confined to the largest connected component of \mathcal{G} (see, for example, [Card, Heining and Kline 2013](#), p. 988).

While $\check{\alpha}$ is routinely used its statistical properties are not well understood. Our aim here is to shed light on how the structure of the network \mathcal{G} affects its sampling behavior and, with it, the reliability of standard inferential procedures based on $\check{\alpha}$. For our analysis edge-specific covariates mostly complicate notation and presentation. It will on occasion be convenient to first analyze (3) when β is treated as known and the outcome vector is redefined as $\mathbf{y} - \mathbf{X}\beta$. Then

$$\hat{\alpha} := (\mathbf{B}'\mathbf{B})^* \mathbf{B}'(\mathbf{y} - \mathbf{X}\beta)$$

is the least-squares estimator of α subject to (4). To appreciate how the structure of \mathcal{G} relates to our problem of estimating the parameter α suppose first that $\mathbf{u} \sim (\mathbf{0}, \sigma^2 \mathbf{I}_m)$. Then

$$\text{var}(\hat{\alpha}) = \sigma^2 \mathbf{L}^*. \tag{6}$$

So, up to a scale factor, the variance of $\hat{\alpha}$ is completely determined by the Laplacian of \mathcal{G} . If, in addition, we were to assume that $\mathbf{u} \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_m)$ we would be in the classical regression setting and, given unbiasedness of $\hat{\alpha}$, size-correct inference could be performed for any sample size. It is not clear, however, how one should proceed with non-classical regression errors.

The validity of standard large-sample arguments is not immediate here. From (6) we have

$$\text{var}(\hat{\alpha}_i) = \sigma^2 \frac{(\mathbf{S}^\dagger)_{ii}}{d_i}, \tag{7}$$

where

$$\mathbf{S} := \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2} = \mathbf{I}_n - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$$

is the normalized Laplacian. Equation (7) follows from the fact that $\mathbf{L}^* = \mathbf{D}^{-1/2} \mathbf{S}^\dagger \mathbf{D}^{-1/2}$, such that $(\mathbf{L}^*)_{ii} = (\mathbf{S}^\dagger)_{ii}/d_i$.⁵ While (7) shows the importance of the sample size in the

⁵Our choice of normalization (4) guarantees the appearance of the Moore-Penrose pseudoinverse of \mathbf{S}

variance of $\hat{\alpha}_i$ through the presence of the degree d_i , it does not imply that $\text{var}(\hat{\alpha}_i)$ shrinks as $d_i \rightarrow \infty$, nor would it give a convergence rate if it did. This is because the normalized Laplacian \mathbf{S} also changes when \mathcal{G} grows.

Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of \mathbf{S} . The spectrum of Laplacian matrices is well studied (see, e.g., [Chung 1997](#)). We have $0 \leq \lambda_i \leq 2$ for all i . We always have that $\lambda_1 = 0$, with $\mathbf{1}_n$ as eigenvector. The number of zero eigenvalues λ_i equals the number of connected components in \mathcal{G} . Hence, if \mathcal{G} is connected, then $\lambda_2 > 0$ is the smallest non-zero eigenvalue of the normalized Laplacian.

Our theory involves conditions on λ_2 and on the degree structure of the network through various harmonic means thereof. λ_2 can be seen as a measure of global connectivity of \mathcal{G} . To see this we note that it can be linked to the Cheeger constant,

$$C := \min_{U \in \{U \subset V: 0 < \sum_{i \in U} d_i \leq \sum_{i \notin U} d_i\}} \frac{\sum_{i \in U} \sum_{j \notin U} (\mathbf{A})_{ij}}{\sum_{i \in U} d_i}.$$

The constant $C \in [0, 1]$ measures how difficult it is to separate \mathcal{G} into two disconnected components by removing edges from it. The numerator in the definition of C is the total weight of the removed edges, the denominator is the total degree in the smallest of the two components. A larger value of C implies a more strongly-connected graph, and it is linked to λ_2 through the inequalities

$$2C \geq \lambda_2 \geq 1 - \sqrt{1 - C^2} \geq \frac{1}{2} C^2, \tag{8}$$

which are due to [Friedland and Nabben \(2002\)](#). Thus, like the Cheeger constant, λ_2 is a measure of global connectivity of the graph \mathcal{G} . Our results below allow for $\lambda_2 \rightarrow 0$ as \mathcal{G} grows, and so cover situations where the graph becomes increasingly more sparse. We will give explicit rates on λ_2 for consistent estimation to be possible.

Example 1 (Erdős-Rényi graph). Consider the [Erdős and Rényi \(1959\)](#) random-graph model, where edges between n vertices are formed independently with probability p_n . The

in \mathbf{L}^* , which is the main reason for that choice. We are grateful to Nadine Geiger for pointing out an inconsistency in our normalization in an earlier version of this paper.

threshold on p_n for \mathcal{G} to be connected is $\ln(n)/n$ (Hoffman, Kahle and Paquette 2013). That is, if $p_n = c \ln(n)/n$ for a constant c , then, as $n \rightarrow \infty$, with probability approaching one, \mathcal{G} is disconnected if $c < 1$ and connected if $c > 1$. In the former case, $\lambda_2 \rightarrow 0$ while, in the latter case, $\lambda_2 \rightarrow 1$, almost surely. \square

2.3 Two-way regression model on bipartite graph

To relate our model to our main motivating examples consider the case of a bipartite graph \mathcal{G} , i.e., $V = V_1 \cup V_2$ and $V_1 \cap V_2 = \emptyset$, and edges are formed only between the subsets V_1 and V_2 but not within. So, for an edge (i, j) we necessarily have that $i \in V_1$ and $j \in V_2$. A bipartite graph describes the interaction between two types of units, such as workers and firms or students and teachers. The outcome of interest here would typically be (log) wages or earnings and test scores, respectively. If we have panel data, so \mathcal{G} is a multigraph, we may observe workers match with different firms over time and observe students in different classrooms or across multiple subjects. In fact, in these applications, such longitudinal data are necessary for \mathcal{G} to be connected. A two-way regression model for such data takes the form

$$\mathbf{y} = \mathbf{B}_1 \boldsymbol{\mu} + \mathbf{B}_2 \boldsymbol{\eta} + \mathbf{X} \boldsymbol{\beta} + \mathbf{u}, \quad (9)$$

where $\boldsymbol{\mu} := (\mu_1, \dots, \mu_{n_1})'$ and $\boldsymbol{\eta} := (\eta_1, \dots, \eta_{n_2})'$ are the $n_1 := |V_1|$ and $n_2 := |V_2|$ parameter vectors for the two types of units, and the $m \times n_1$ and $m \times n_2$ matrices \mathbf{B}_1 and \mathbf{B}_2 have entries

$$(\mathbf{B}_1)_{\varepsilon_e i} := \begin{cases} 1 & \text{if } e \in E_{(i,j)} \text{ for some } j \in V_2 \\ 0 & \text{otherwise,} \end{cases}$$

$$(\mathbf{B}_2)_{\varepsilon_e j} := \begin{cases} 1 & \text{if } e \in E_{(i,j)} \text{ for some } i \in V_1 \\ 0 & \text{otherwise.} \end{cases}$$

This is a workhorse specification to capture heterogeneity across units in linked data sets. It can be cast into (3) by setting

$$\alpha_i = \begin{cases} \mu_i & \text{if } i \in V_1, \\ -\eta_i & \text{if } i \in V_2, \end{cases}$$

sorting the units in V by type so that we can write $\boldsymbol{\alpha} = (\boldsymbol{\mu}', -\boldsymbol{\eta}')$, and constructing the $m \times n$ matrix $\mathbf{B} = (\mathbf{B}_1, -\mathbf{B}_2)$ by concatenation. Choosing the sign in front of η_i is without loss of generality because links are only formed between, but never within, the subsets V_1 and V_2 . The need for a normalization built-in in our general specification arises here from the fact that (9) is invariant to reparametrizations of the form $(\mu_i, \eta_j) \mapsto (\mu_i + c, \eta_j - c)$ for any c .

The two-way regression model provides an interesting example where a weighted graph arises naturally. In many applications the researcher is primarily interested in learning the parameters of one type, say those $i \in V_2$. This is so in teacher value-added models, for example. There, interest lies in estimating the n_2 teacher effects while controlling for unobserved student-specific heterogeneity through the inclusion of n_1 student effects (see, e.g., [Jackson, Rockoff and Staiger 2014](#)). Partialling-out the vector $\boldsymbol{\mu}$ from the two-way model in (9) gives

$$\mathbf{M}_{\mathbf{B}_1} \mathbf{y} = (\mathbf{M}_{\mathbf{B}_1} \mathbf{B}_2) \boldsymbol{\eta} + (\mathbf{M}_{\mathbf{B}_1} \mathbf{X}) \boldsymbol{\beta} + \mathbf{M}_{\mathbf{B}_1} \mathbf{u}, \quad \mathbf{M}_{\mathbf{B}_1} := \mathbf{I}_m - \mathbf{B}_1 (\mathbf{B}'_1 \mathbf{B}_1)^* \mathbf{B}'_1. \quad (10)$$

From standard partitioned-regression theory, the least-squares estimator of $\boldsymbol{\eta}$ from this equation is numerically identical to the one obtained from joint estimation of $\boldsymbol{\mu}$ and $\boldsymbol{\eta}$ in (9). However, the formulation in (10) is helpful in understanding the behavior of the estimator of $\boldsymbol{\eta}$. The properties of the matrix $\mathbf{B}'_2 \mathbf{M}_{\mathbf{B}_1} \mathbf{B}_2$ drive the sampling behavior of $\check{\boldsymbol{\eta}}$. This matrix is the Laplacian of a weighted one-mode projection ([Newman, 2010](#), p. 124) of the bipartite graph \mathcal{G} on the n_2 vertices in V_2 .

It is instructive to discuss this one-mode projection in more detail and to formalize how it fits the general setup in (3). Projecting the bipartite graph $\mathcal{G} = \mathcal{G}(V_1 \cup V_2, E)$ on V_2 is done by suppressing the vertices in V_1 . This gives a new (unipartite) graph, say $\mathcal{G}' = \mathcal{G}(V_2, E')$. Each edge pair (e_1, e_2) with $e_1 \in E_{(i,j)}$ and $e_2 \in E_{(i,j')}$ in \mathcal{G} for some $i \in V_1$ and $j, j' \in V_2$ gives rise to a single edge $e = (e_1, e_2) \in E'_{(j,j')}$ in \mathcal{G}' . In the student-teacher example, two teachers j and j' are connected by an edge in \mathcal{G}' if and only if there exists at least one student i that they have both taught. Alternatively, the edge $e = (e_1, e_2) \in E'$ exists because $e_1 \in E$ and $e_2 \in E$ both connect to the same vertex $i \in V_1$. Given the edges

e_1, e_2 this connecting vertex i is unique; for later use we denote it by $c((e_1, e_2))$. In \mathcal{G}' we have

$$m' := |E'| = \sum_{i \in V_1} \frac{d_i(d_i - 1)}{2}$$

edges; m' need not equal m and, indeed, may be much larger. We again label $e' \in E'$ by natural numbers $\varepsilon'_{e'}$. The process of concatenating edges in E to form the new edge set E' can be described by the $m' \times m$ matrix \mathbf{Q} with entries

$$(\mathbf{Q})_{\varepsilon'_{e'} \varepsilon_{e_1}} := \begin{cases} 1 & \text{if } e' = (e_1, e_2) \text{ for some } e_2 \in E, \\ -1 & \text{if } e' = (e_2, e_1) \text{ for some } e_2 \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Choosing the orientation of the rows of \mathbf{Q} is without loss of generality. The matrix \mathbf{Q} has a first-differencing interpretation. Indeed, when applied to the two-way regression model (9) we get $\mathbf{Q}\mathbf{y} = \mathbf{Q}\mathbf{B}_2\boldsymbol{\eta} + \mathbf{Q}\mathbf{X}\boldsymbol{\beta} + \mathbf{Q}\mathbf{u}$, because $\mathbf{Q}\mathbf{B}_1 = 0$. Thus, \mathbf{Q} sweeps out the n_1 nuisance parameters $\boldsymbol{\mu}$ and transforms original outcomes into first differences. The matrix $\mathbf{Q}\mathbf{B}_2$ is the (oriented) incidence matrix of an unweighted graph, and this first-differenced regression equation fits (3).⁶ Applying least-squares directly to the first differences is inefficient and is not equivalent to estimation of the two-way regression model. Ordinary least-squares estimation of (10) is numerically equivalent to weighted least-squares estimation of the first-differenced equation. The relevant $m' \times m'$ diagonal weight matrix \mathbf{W} has entries $(\mathbf{W})_{\varepsilon'_e \varepsilon'_e} := 1/\sqrt{d_{c(e)}}$. Ordinary least-squares applied to (10) and to

$$\mathbf{W}\mathbf{Q}\mathbf{y} = \mathbf{W}\mathbf{Q}\mathbf{B}_2\boldsymbol{\eta} + \mathbf{W}\mathbf{Q}\mathbf{X}\boldsymbol{\beta} + \mathbf{W}\mathbf{Q}\mathbf{u} \tag{11}$$

yields the same result. Here, $\mathbf{W}\mathbf{Q}\mathbf{B}_2$ is the incidence matrix of a weighted one-mode projection of \mathcal{G} . This \mathcal{G}' determines the properties of the least-squares estimator. Its

⁶ $\mathbf{Q}\mathbf{B}_2$ will contain rows with only zero entries if there are differenced outcomes that do not depend on $\boldsymbol{\eta}$. This is at odds with the definition of an incidence matrix. Dropping these differences from $\mathbf{Q}\mathbf{y}$, however, restores the incidence matrix interpretation of $\mathbf{Q}\mathbf{B}_2$. This operation does not affect estimation of $\boldsymbol{\eta}$ and so is irrelevant for our purposes. However, the differenced outcomes may provide information on $\boldsymbol{\beta}$, which is why we prefer to work with $\mathbf{Q}\mathbf{B}_2$ as defined here.

Laplacian is

$$L' := B_2'(Q'W^2Q)B_2 = B_2'M_{B_1}B_2,$$

where we use the fact that $Q'W^2Q = M_{B_1}$.

The adjacency matrix of \mathcal{G}' is the $n_2 \times n_2$ matrix \mathbf{A}' with entries

$$(\mathbf{A}')_{jj'} := \begin{cases} \sum_{i \in [j] \cap [j']} \frac{|E_{(i,j)}| |E_{(i,j')}|}{\sum_{k \in V_2} |E_{(i,k)}|} & \text{for } j \neq j', \\ 0 & \text{for } j = j'. \end{cases}$$

Here, $[j] \cap [j']$ is the set of all vertices in V_1 that are connected to both $j \in V_2$ and $j' \in V_2$ in the original bipartite graph \mathcal{G} . In the student-teacher example two teachers are connected by an edge if there is at least one student who was taught by both teachers. The weight $(\mathbf{A}')_{jk}$ of the edge is larger the more students there are connecting teachers j and k , and the more courses they have taken from these teachers. \mathcal{G}' determines the accuracy with which teacher value-added can be estimated.

The matrix \mathbf{A}' is also the adjacency matrix of the simple graph obtained from \mathcal{G}' by replacing all edges $e \in E'_{(j,j')}$ by one weighted edge, with weight $(\mathbf{A}')_{jj'}$. Figure 1 provides an illustration of a simple bipartite graph for students (circular vertices) and teachers (square vertices), given in the left plot, and its induced weighted graph featuring only teachers, given in the right plot. The thickness of the edge between (j, j') in the latter plot reflects the magnitude of the weight $(\mathbf{A}')_{jj'}$.

The device of a one-mode projection highlights the importance of having movers in panel data. In matched worker-firm data sets workers do not frequently switch employer over the course of the sampling period. This lack of mobility is one cause of the substantial bias that is observed in the correlation coefficient between (estimated) worker and firm effects (Abowd, Kramarz, Lengermann and Perez-Duarte 2004, Andrews, Gill, Schank and Upward 2008, 2012). While this is now well recognized, limited mobility has consequences more broadly. Indeed, it implies that few workers connect firms in the one-mode firm projection. Therefore, the induced graph may be only weakly connected (and λ_2 will be close to zero) and the variance of the estimator of the firm effects may be large. This




Figure 1: A simple unweighted bipartite graph (left) with links between V_1 (circular vertices) and V_2 (square vertices), and the induced weighted graph (right) on V_2 alone resulting from profiling out the parameters associated with V_1 .

is not only detrimental for identifying sorting between workers and firms but, indeed, complicates estimation and inference of the firm effects as well as all their moments, such as their variance. Restricting attention to large firms need not resolve this problem. An analogous argument holds for teacher effects and their estimated variance, and so for our ability to infer the contribution of teacher value-added to observed variation in test scores. We illustrate this in our data below.

3 Variance bound and asymptotic analysis

3.1 Finite-sample bound

To work towards general distribution theory it is instructive to start with a finite-sample bound on the variance of the fixed-effect estimator when the errors are homoskedastic and uncorrelated. Let

$$h_i := \left(\frac{1}{d_i} \sum_{j \in [i]} \frac{(\mathbf{A})_{ij}^2}{d_j} \right)^{-1}. \quad (12)$$

This is a (weighted) harmonic mean of the (weighted) degrees $d_j/(\mathbf{A})_{ij}$ of all $j \in [i]$. Note that, for a given vertex i , h_i is increasing in the degree of its direct neighbors.

Theorem 2 (Variance bound). *Let \mathcal{G} be connected. Suppose that $\mathbf{u} \sim (\mathbf{0}, \sigma^2 \mathbf{I}_m)$. Then*

$$\frac{\sigma^2}{d_i} - \frac{2\sigma^2}{m} \leq \text{var}(\hat{\alpha}_i) \leq \frac{\sigma^2}{d_i} \left(1 + \frac{1}{\lambda_2 h_i}\right) - \frac{2\sigma^2}{m}.$$

Theorem 2 states that, for a given degree d_i and global connectivity measure λ_2 , the upper bound on the variance of $\hat{\alpha}_i$ is smaller if the direct neighbors of vertex i are themselves more strongly connected to other vertices in the network. The theorem provides insight into how the local connectivity structure of the network, around vertex i , affects statistical precision.

Example 1 (cont'd). Consider the Erdős and Rényi (1959) random-graph model with $p_n = c \ln(n)/n$ for $c > 1$. Let i be a randomly chosen vertex. Then, as $n \rightarrow \infty$, we have, almost surely,

$$\lambda_2 \rightarrow 1, \quad \frac{d_i}{\ln n} \rightarrow c, \quad \frac{h_i}{\ln n} \rightarrow c.$$

Consequently,

$$\text{var}(\hat{\alpha}_i) = \frac{\sigma^2}{d_i} + O(d_i^{-2})$$

follows from Theorem 2. □

Additional calculations for analytically-tractable cases where $\lambda_2 \rightarrow 0$ as the network grows are provided in the supplementary material.

Theorem 2 highlights the importance of $\lambda_2 h_i \rightarrow \infty$ as a sufficient condition for the parametric rate $d_i^{-1/2}$ to be attainable for estimation of $\hat{\alpha}_i$. This result carries over to the model with covariates. Let

$$\rho := \left\| (\mathbf{X}'\mathbf{X})^{-1/2} (\mathbf{X}'\mathbf{M}_B\mathbf{X}) (\mathbf{X}'\mathbf{X})^{-1/2} \right\|_2,$$

where $\|\cdot\|_2$ denotes the spectral norm and $\mathbf{M}_B := \mathbf{I}_m - \mathbf{B}(\mathbf{B}'\mathbf{B})^* \mathbf{B}'$. Note that $\rho \in [0, 1]$ is a measure of non-collinearity between the columns of \mathbf{X} and \mathbf{B} , with ρ close to zero indicating near-collinearity. Indeed, while $\mathbf{X}'\mathbf{X}$ measures the total variation in \mathbf{X} , $\mathbf{X}'\mathbf{M}_B\mathbf{X}$ captures the residual variation in \mathbf{X} after its linear dependence on \mathbf{B} has been partialled out. For $i \in V$ let \mathbf{b}_i be i^{th} column of \mathbf{B} , and let $\bar{\mathbf{x}}_i := \mathbf{X}'\mathbf{b}_i/d_i$ and $\boldsymbol{\Omega} := \mathbf{X}'\mathbf{X}/m$ in the following theorem.

Theorem 3 (Variance bound (cont'd.)). *Let \mathcal{G} be connected. Suppose that $\mathbf{u} \sim (\mathbf{0}, \sigma^2 \mathbf{I}_m)$, $\text{rank}(\mathbf{X}) = p$, and $\text{rank}((\mathbf{X}, \mathbf{B})) = p + n - 1$. Then*

$$|\text{var}(\tilde{\alpha}_i) - \text{var}(\hat{\alpha}_i)| \leq \frac{2\sigma^2}{\rho} \left(\frac{1 - \rho}{d_i(\lambda_2 h_i)} + \frac{\bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i}{m} \right),$$

for all $i \in V$.

This result shows that, if ρ is bounded away from zero, introducing covariates only has a higher-order effect on the statistical precision of the fixed-effect estimator. In particular we have

$$\text{var}(\tilde{\alpha}_i) = \frac{\sigma^2}{d_i} + o(d_i^{-1}), \quad (13)$$

provided that $\lambda_2 h_i \rightarrow \infty$ as \mathcal{G} grows. Furthermore, the parametric rate is achievable even if λ_2 is not treated as fixed, and \mathcal{G} becomes less dense as more vertices are added to the network.

3.2 Large-sample analysis

We now discuss asymptotic results under more general conditions on the regression errors. The following theorem provides a first-order representation of $\tilde{\alpha}_i$. Let $\boldsymbol{\Sigma} := \mathbb{E}(\mathbf{u}\mathbf{u}')$.

Theorem 4 (First-order representation). *Let \mathcal{G} be connected. Assume that $\text{rank}(\mathbf{X}) = p$ and $\text{rank}((\mathbf{X}, \mathbf{B})) = p + n - 1$. Suppose that $\mathbb{E}(\mathbf{u}) = \mathbf{0}$ and that $\|\boldsymbol{\Sigma}\|_2 \leq \bar{\sigma}^2 = O(1)$. Then*

$$\tilde{\alpha}_i - \alpha_i = \frac{\mathbf{b}_i' \mathbf{u}}{d_i} + \epsilon_i + \tilde{\epsilon}_i,$$

where ϵ_i and $\tilde{\epsilon}_i$ are zero-mean random variables that satisfy $\mathbb{E}(\epsilon_i^2) \leq \bar{\sigma}^2(1 + \rho)/(\rho d_i \lambda_2 h_i)$, and $\mathbb{E}(\tilde{\epsilon}_i^2) \leq \bar{\sigma}^2 \bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i / (\rho m)$.

From the definition of the incidence matrix in (1), the m -vector \mathbf{b}_i has as many non-zero entries as there are edges involving vertex i . Further, $\mathbf{b}_i' \mathbf{b}_i = d_i$. Hence, the term $\mathbf{b}_i' \mathbf{u} / d_i$ is a (weighted) sample mean of the regression errors associated with the edges that involve vertex i .

We next consider sequences of growing networks such that

$$\rho^{-1} = O(1), \quad d_i/m \rightarrow 0, \quad \bar{\mathbf{x}}_i' \boldsymbol{\Omega}^{-1} \bar{\mathbf{x}}_i = O(1). \quad (14)$$

These are relatively weak conditions that ensure that the fact that $\boldsymbol{\beta}$ is estimated can be ignored in large samples. Moreover, they imply that

$$\epsilon_i = O_p \left(\frac{1}{\sqrt{d_i(\lambda_2 h_i)}} \right)$$

and that $\tilde{\epsilon}_i = O_p(1/\sqrt{m}) = o_p(1/\sqrt{d_i})$. The main implication of the theorem is that, then, under the now familiar condition $\lambda_2 h_i \rightarrow \infty$, as $d_i \rightarrow \infty$,

$$(\check{\alpha}_i - \alpha_i) \xrightarrow{p} \frac{\mathbf{b}'_i \mathbf{u}}{d_i}.$$

This result allows the errors to be heteroskedastic and correlated.

With Theorem 4 in hand the limit distribution of $\check{\alpha}_i$ can be deduced under conventional conditions. As an example we do so next for independent but heterogeneously distributed (i.n.i.d.) regression errors.

Theorem 5 (Limit distribution for i.n.i.d. errors). *Let the assumptions of Theorem 4 and the conditions in (14) hold. Suppose that the regression errors are independent, have bounded fourth-order moments, and variances bounded away from zero, and that the edge weights are bounded away from zero and from infinity. Then*

$$\frac{\check{\alpha}_i - \alpha_i}{\sqrt{\mathbf{b}'_i \boldsymbol{\Sigma} \mathbf{b}_i / d_i}} \xrightarrow{d} N(0, 1),$$

as $d_i \rightarrow \infty$, provided that $\lambda_2 h_i \rightarrow \infty$.

When the errors \mathbf{u} are independent and homoskedastic we have $\mathbf{b}'_i \boldsymbol{\Sigma} \mathbf{b}_i = \sigma^2 \mathbf{b}'_i \mathbf{b}_i = \sigma^2 d_i$ and the variance in the theorem reduces to σ^2/d_i , which agrees with (13).

A plug-in estimator of $\mathbf{b}'_i \boldsymbol{\Sigma} \mathbf{b}_i$ is $\mathbf{b}'_i \check{\boldsymbol{\Sigma}} \mathbf{b}_i$, where $\check{\boldsymbol{\Sigma}} := \text{diag}(\check{\mathbf{u}} \check{\mathbf{u}}')/m$ and $\check{\mathbf{u}}$ are the residuals from the least-squares regression. This involves estimation of α_j for all $j \in [i]$. We have that

$$(\mathbf{b}'_i \check{\boldsymbol{\Sigma}} \mathbf{b}_i - \mathbf{b}'_i \boldsymbol{\Sigma} \mathbf{b}_i) / d_i \xrightarrow{p} 0$$

as $d_i \rightarrow \infty$, provided that, in addition to the conditions of Theorem 5 holding, we have that $\lambda_2 H_i$ is bounded away from zero, where

$$H_i := \left(\sum_{j \in [i]} \frac{(h_i/d_i)/d_j}{h_j} \right)^{-1}$$

is a weighted harmonic mean. At the heart of this result lies (a local version of) a global convergence rate on $\|\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha}\|$, which is interesting in its own right. More precisely, letting

$$h := \left(\frac{1}{n} \sum_{i=1}^n \frac{1}{d_i} \right)^{-1} \quad \text{and} \quad H := \left(\sum_{i=1}^n \frac{(h/n)/d_i}{h_i} \right)^{-1},$$

it is easy to see that

$$\|\check{\boldsymbol{\alpha}} - \boldsymbol{\alpha}\| = O_p(\sqrt{n/h}),$$

provided that $\lambda_2 H$ is bounded away from zero.

3.3 Estimation of moments

Suppose that the α_i are sampled from some distribution. One might be interested to learn the variance of this distribution—as in, say, [Rockoff \(2004\)](#) or [Card, Heining and Kline \(2013\)](#)—or some other moment. The typical estimator is the corresponding sample moment of the estimated effects. Sampling noise in the estimated individual effects will introduce bias in the moment estimator, however. To see this, consider estimation of the variance in a simple model without regressors. The sample variance of the estimated effects in this case is

$$\hat{\boldsymbol{\alpha}}' \mathbf{M}_{\boldsymbol{\iota}_n} \hat{\boldsymbol{\alpha}} / (n-1),$$

where $\mathbf{M}_{\boldsymbol{\iota}_n} := \mathbf{I}_n - \boldsymbol{\iota}_n \boldsymbol{\iota}_n' / n$ is the usual demeaning matrix. When $\mathbf{u} \sim (\mathbf{0}, \sigma^2 \mathbf{I}_m)$ its bias is

$$\sigma^2 \frac{\text{tr}(\mathbf{L}^*)}{(n-1)} = (n-1)^{-1} \sum_{i=1}^n \text{var}(\hat{\alpha}_i),$$

which clearly shows how imprecise estimation of α_i contributes to the bias in the variance estimator.

It is difficult to derive an exact expression for the bias for more general functionals. Theorem 5 is instrumental here. Suppose that $\tau := \mathbb{E}(\varphi(\alpha_i))$ is of interest. Its plug-in estimator is

$$\check{\tau} := n^{-1} \sum_{i=1}^n \varphi(\check{\alpha}_i).$$

Under the conditions of the theorem we can calculate the leading bias in this estimator as

$$b := n^{-1} \sum_{i=1}^n \mathbb{E} \left(\frac{\varphi''(\alpha_i)}{2} \frac{\mathbf{b}'_i \boldsymbol{\Sigma} \mathbf{b}_i / d_i}{d_i} \right) = O(h^{-1}),$$

where φ'' denotes the second derivative, provided $\lambda_2 H \rightarrow \infty$. Simple regularity conditions on φ for this bias result to hold are that it is differentiable with $\mathbb{E}(\varphi''(\alpha_i)^2) < \infty$ and bounded third derivative. So, quite generally, the bias will shrink like h^{-1} . Therefore, for the bias to vanish and $\check{\tau}$ to be consistent, we need that the degrees of the individual vertices grow with n for an increasing fraction of the vertices.

If the functional of interest is the variance, an exact bias correction can be performed (see Andrews, Gill, Schank and Upward 2008 and Kline, Saggio and Sølvesten 2018). For functionals like τ , a plug-in estimator of the leading-order bias b is easily formed and so an adjusted estimator is readily constructed. Its effectiveness as a bias-correction device will again depend on the connectivity structure of the graph. We postpone a detailed analysis to future work. In a recent contribution Kline, Saggio and Sølvesten (2018) present limit theory for quadratic forms in $\boldsymbol{\alpha}$.

4 Empirical illustrations

4.1 Teacher value-added

We construct a graph connecting teachers as the (weighted) one-mode projection from matched student-teacher data from the North Carolina Education Research Center. The projection of interest is the one discussed in Section 2.3. The full data set includes scores for a standardized test in reading in elementary schools in North Carolina and was used by Verdier (2018) to estimate the effect of class-size reduction on student performance. The

analysis conducted here is useful to assess the precision with which teacher value-added can be estimated. The data concern pupils in Grades 4 and 5 of elementary school over the period 2008–2012. The full teacher graph (with a single weighted edge between neighboring teachers, as in Figure 1) has 12,057 vertices and 53,741 edges and is disconnected. The largest connected component involves 41,612 edges between 11,945 teachers and we work with this subgraph. With $\lambda_2 = .0039$ the projected teacher graph is weakly connected. Its local connectivity is summarized in Table I. The table contains the mean, standard deviation, and deciles of the relevant degree distributions. Inspection reveals that the degrees are small for all teachers.

Table I: Summary statistics for the teacher graph

	mean	stdev	10 th %	20 th %	30 th %	40 th %	50 th %	60 th %	70 th %	80 th %	90 th %
d_i	13.87	10.76	3.00	5.50	7.50	9.00	11.00	14.00	17.50	21.50	27.50
h_i	7.15	7.13	2.43	3.30	4.01	4.72	5.48	6.36	7.44	9.12	12.56
H_i	36.48	58.59	3.03	5.72	10.48	14.76	19.81	26.20	35.67	50.65	83.48

The weak connectivity suggests that inference on teacher value-added will be difficult. To get a sense of the precision of a first-order asymptotic approach we can look at the ratio

$$\frac{(\mathbf{L}^*)_{ii}}{1/d_i} = (\mathbf{S}^\dagger)_{ii}.$$

This is the exact variance of $\hat{\alpha}_i$ to its large-sample approximation in a regression model with homoskedastic and uncorrelated errors. This ratio is free of σ^2 and can be computed directly from the graph. The left plot of Figure 2 shows the deciles of the distribution of $(\mathbf{S}^\dagger)_{ii}$. The asymptotic approximation is revealed to be widely inaccurate. On average, the actual variance is about 2.5 times larger than its approximation. Even the first decile equals 1.29. This implies that confidence intervals based on the large-sample arguments in Theorem 4 are overoptimistic. To illustrate this the right plot in Figure 2 gives the distribution of the width of 95% confidence intervals for the α_i using both the exact variance (solid line) and its large-sample approximation (dashed line) for the case $\sigma^2 = 1$. The former stochastically dominates the latter.




Figure 2: Deciles of the distribution of $(\mathbf{S}^\dagger)_{ii}$ (left plot) and empirical distributions of the width of 95% confidence bands (right plot). The width is calculated as $2 \times (1.96 \mathbf{L}_{ii})$ (solid curve) and $2 \times (1.96 d_i^{-1})$ (dashed curve).

The large variability in the estimators of teacher value-added implies a large bias in their estimated variance. We calculate

$$\text{tr}(\mathbf{L}^*)/(n - 1) = .3545,$$

so the bias in the plug-in estimator of the variance is about one-third of the error variance when $\mathbf{u} \sim (\mathbf{0}, \sigma^2 \mathbf{I}_m)$. The large-sample approximation to the bias here is proportional to h^{-1} . With $h = 5.4554$ this yields a bias of about 18%, roughly half the size of the exact bias.

One reason for the global connectivity of the teacher graph to be low is limited mobility of teachers between schools. If interest lies in comparing teacher effectiveness within a given school it suffices to restrict attention to that subgraph. Of course, the effective sample size from which teacher value-added is estimated will remain small unless additional years of data are collected. Making accurate comparisons between schools is more complicated as it requires many teachers to switch schools during the sampling period. Collecting additional years of data will not automatically lead to more precise estimates. [Mansfield \(2015\)](#) discusses the feasibility of ranking teachers within and between schools (see also [Mihaly, McCaffrey, Sass and Lockwood 2013](#) for related discussions). A previous draft of this paper contains versions of our main theoretical results specialized to within and

between decompositions of graphs.

4.2 Occupational network

Wage regressions on worker and occupational dummies (as in [Kambourov and Manovskii 2009](#), for example) provide an interesting example of a situation where more accurate results can be obtained. We use all 18 available waves from the BHPS (for a total of 132,097 observations) to construct the induced (weighted) occupational network. The Standard Occupational Classification (SOC90) in the BHPS distinguishes (at the three-digit level) between 374 occupations. We again focus on the largest connected component, which contains 365 occupations with 14,825 weighted edges between them. As a measure of global connectivity here we find $\lambda_2 = .3289$. Compared to traditional matched employer-employee data our occupational network does not suffer as much from limited mobility. One reason is that the number of occupations is relatively small compared to the number of workers. Another is that workers may switch occupation also if they remain employed by the same firm, for example due to internal promotions. Finally, as we are dealing with self-reported occupations there is also the possibility of spurious mobility due to misreporting. A look at the distributions summarized in [Table II](#) reveals that the degrees and harmonic means tend to be larger here than in the teacher graph.

Table II: Summary statistics for the occupation graph

	mean	stdev	10 th %	20 th %	30 th %	40 th %	50 th %	60 th %	70 th %	80 th %	90 th %
d_i	155.06	268.82	7.35	16.08	27.68	45.68	67.10	92.66	143.99	212.09	402.01
h_i	81.50	125.20	14.52	20.88	27.25	35.16	46.23	60.61	80.72	113.67	163.74
H_i	213.77	455.93	11.15	22.90	35.29	48.59	66.34	102.82	172.75	296.80	539.07

The distribution of $(\mathcal{S}^\dagger)_{ii}$ now places most of its mass in close vicinity of unity. Its mean and standard deviation are 1.034 and .0521. The median is 1.0202 while the first and ninth decile are 1.0046 and 1.0759, respectively. This suggests that, here, the large-sample approximation to the variance is a much more accurate reflection of actual estimation

uncertainty. Similarly, we may again calculate $\text{tr}(\mathbf{L}^*)/(n - 1) = .0577$, which is about 7 times smaller than in the previous example. Further, as $h^{-1} = .0566$, here, the bias approximation is quite accurate.

5 Conclusion

We have presented inference results on individual effects in a linear fixed-effect regression model when the underlying data structure constitutes a (weighted) graph. An important example is a two-way regression model on a bipartite graph. The main contribution of this paper is to quantify the dependence of statistical precision of the estimator on the connectivity structure of the graph. A key measure of global connectivity is the smallest non-zero eigenvalue of the (normalized) Laplacian matrix of the graph. It reflects the intuitive notion of mobility in the network. A small eigenvalue captures the presence of bottlenecks, which is detrimental to statistical precision. Several measures of local connectivity, such as the degree structure and various harmonic means thereof, also arise naturally in our analysis.

Our theoretical work highlights the importance of and the interplay between global and local measures of network connectivity for conventional inferential approaches to be reliable. The analysis points to a set of simple statistics that can be inspected to evaluate whether the network is sufficiently well connected in a given application. In an application to teacher value-added we find that this is not the case. We further find that conventional standard errors on teacher value-added estimates are much too small, resulting in a false sense of (statistical) precision on these parameter estimates. In an occupational network, on the other hand, we find much higher measures of connectivity and support for our large-sample approximations.

References

- Abowd, J., R. Creecy, and F. Kramarz (2002). Computing person and firm effects using linked longitudinal employer-employee data. U.S. Census Technical Paper TP-2002-06.
- Abowd, J., F. Kramarz, P. Lengermann, and S. Perez-Duarte (2004). Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States. Mimeo.
- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67, 251–333.
- Amiti, M. and D. E. Weinstein (2018). How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data. *Journal of Political Economy* 126, 525–587.
- 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* 171, 673–697.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters* 117, 824–827.
- Bonhomme, S., T. Lamadon, and E. Manresa (2018). A distributional framework for matched employer-employee data. Forthcoming in *Econometrica*.
- Bramoullé, Y., R. Kranton, and M. D’Amours (2014). Strategic interaction and networks. *American Economic Review* 104, 898–930.
- Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics* 128, 967–1015.
- Chetty, R. and N. Hendren (2018). The impacts of neighborhoods on intergenerational mobility II: County-level estimates. Forthcoming in *Quarterly Journal of Economics*.
- Chung, F. R. K. (1997). *Spectral Graph Theory*. Volume 92 of CBMS Regional Conference Series in Mathematics, American Mathematical Society.
- Erdős, P. and A. Rényi (1959). On random graphs. *Publicationes Mathematicae* 6, 290–297.
- Finkelstein, A., M. Gentzkow, and H. Williams (2016). Sources of geographic variation in health care: Evidence from patient migration. *Quarterly Journal of Economics* 131, 1681–1726.

- Friedland, S. and R. Nabben (2002). On Cheeger-type inequalities for weighted graphs. *Journal of Graph Theory* 41, 1–17.
- Hoffman, C., M. Kahle, and E. Paquette (2013). Spectral gaps of random graphs and applications to random topology. Mimeo.
- Jackson, C. K., J. E. Rockoff, and D. O. Staiger (2014). Teacher effects and teacher related policies. *Annual Review of Economics* 6, 801–825.
- Kambourov, G. and I. Manovskii (2009). Occupational specificity of human capital. *International Economic Review* 50, 63–115.
- Kettemann, A., A. I. Mueller, and J. Zweimüller (2017). Wages, workers and vacancy durations: Evidence from linked data. Mimeo.
- Kline, P., R. Saggio, and M. Sølvssten (2018). Leave-out estimation of variance components. Mimeo.
- Lentz, R., S. Piyapromdee, and J.-M. Robin (2018). On worker and firm heterogeneity in wages and employment mobility: Evidence from danish register data. Mimeo.
- Mansfield, R. K. (2015). Teacher quality and student inequality. *Journal of Labor Economics* 33, 751–788.
- Mihaly, K., D. McCaffrey, T. R. Sass, and J. R. Lockwood (2013). Where you come from or where you go? Distinguishing between school quality and the effectiveness of teacher preparation program graduates. *Education Finance and Policy* 8, 459–493.
- Newman, M. E. J. (2010). *Networks*. Oxford University Press.
- Nimczik, J. (2018). Job mobility networks and endogenous labor markets. Mimeo.
- Rockoff, J. E. (2004). The impact of individual teachers on student achievement: Evidence from panel data. *American Economic Review* 94, 247–252.
- Rothstein, J. (2010). Teacher quality in educational production: Tracking, decay, and student achievement. *Quarterly Journal of Economics* 125, 175–214.
- Schiebinger, G., M. J. Wainwright, and B. Yu (2015). The geometry of kernelized spectral clustering. *Annals of Statistics* 43, 819–846.
- Verdier, V. (2018). Estimation and inference for linear models with two-way unobserved heterogeneity and sparsely matched data. Forthcoming in *Review of Economics and Statistics*.