

NBER WORKING PAPER SERIES

LEAVE-OUT ESTIMATION OF VARIANCE COMPONENTS

Patrick Kline
Raffaele Saggio
Mikkel Sølvssten

Working Paper 26244
<http://www.nber.org/papers/w26244>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2019

We thank Isaiah Andrews, Bruce Hansen, Whitney Newey, Anna Mikusheva, Jack Porter, Andres Santos, Azeem Shaikh and seminar participants at UC Berkeley, CEMFI, Chicago, Harvard, UCLA, MIT, Northwestern, NYU, Princeton, Queens, UC San Diego, Wisconsin, the NBER Labor Studies meetings, and the CEME Interactions workshop for helpful comments. The data used in this study was generously provided by the Fondazione Rodolfo De Benedetti and originally developed by the Economics Department of the Università Ca Foscari Venezia under the supervision of Giuseppe Tattara. We thank the Berkeley Institute for Research on Labor and Employment for funding support and Schmidt Futures, which provided financial assistance for this project through the Labor Science Initiative at the Berkeley Opportunity Lab. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Patrick Kline, Raffaele Saggio, and Mikkel Sølvssten. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Leave-out Estimation of Variance Components
Patrick Kline, Raffaele Saggio, and Mikkel Solvsten
NBER Working Paper No. 26244
September 2019
JEL No. C1,J31

ABSTRACT

We propose leave-out estimators of quadratic forms designed for the study of linear models with unrestricted heteroscedasticity. Applications include analysis of variance and tests of linear restrictions in models with many regressors. An approximation algorithm is provided that enables accurate computation of the estimator in very large datasets. We study the large sample properties of our estimator allowing the number of regressors to grow in proportion to the number of observations. Consistency is established in a variety of settings where plug-in methods and estimators predicated on homoscedasticity exhibit first-order biases. For quadratic forms of increasing rank, the limiting distribution can be represented by a linear combination of normal and non-central χ^2 random variables, with normality ensuing under strong identification. Standard error estimators are proposed that enable tests of linear restrictions and the construction of uniformly valid confidence intervals for quadratic forms of interest. We find in Italian social security records that leave-out estimates of a variance decomposition in a two-way fixed effects model of wage determination yield substantially different conclusions regarding the relative contribution of workers, firms, and worker-firm sorting to wage inequality than conventional methods. Monte Carlo exercises corroborate the accuracy of our asymptotic approximations, with clear evidence of non-normality emerging when worker mobility between blocks of firms is limited.

Patrick Kline
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720
and NBER
pkline@econ.berkeley.edu

Raffaele Saggio
University of British Columbia
Vancouver, BC
Canada
rsaggio@mail.ubc.ca

Mikkel Solvsten
University of Wisconsin - Madison
Office 7430
Department of Economics
William H. Sewell Social Science Building
1180 Observatory Drive
Madison, WI 53706
soelvsten@wisc.edu

As economic datasets have grown large, so has the number of parameters employed in econometric models. Typically, researchers are interested in certain low dimensional summaries of these parameters that communicate the relative influence of the various economic phenomena under study. An important benchmark comes from Fisher (1925)’s foundational work on analysis of variance (ANOVA) which he proposed as a means of achieving a “separation of the variance ascribable to one group of causes, from the variance ascribable to other groups.”¹

A large experimental literature (Sacerdote, 2001; Graham, 2008; Chetty et al., 2011; Angrist, 2014) employs variants of Fisher’s ANOVA approach to infer the degree of variability attributable to peer or classroom effects. Related methods are often used to study heterogeneity across firms, workers, and schools in their responsiveness to exogenous regressors with continuous variation (Raudenbush and Bryk, 1986, 2002; Arellano and Bonhomme, 2011; Graham and Powell, 2012). In labor economics, log-additive models of worker and firm fixed effects are increasingly used to study worker-firm sorting and the dispersion of firm specific pay premia (Abowd et al., 1999; Card et al., 2013, 2018; Song et al., 2017; Sorkin, 2018) and analogous methods have been applied to settings in health economics (Finkelstein et al., 2016; Silver, 2016) and the economics of education (Arcidiacono et al., 2012).

This paper considers estimation of and inference on *variance components*, which we define broadly as quadratic forms in the parameters of a linear model. Notably, this definition yields an important connection to the recent literature on testing linear restrictions in models with many regressors (Anatolyev, 2012; Chao et al., 2014; Cattaneo et al., 2018). Traditional variance component estimators are predicated on the assumption that the errors in a linear model are identically distributed draws from a normal distribution. Standard references on this subject (e.g., Searle et al., 2009) suggest diagnostics for heteroscedasticity and non-normality, but offer little guidance regarding estimation and inference when these problems are encountered. A closely related literature on panel data econometrics proposes variance component estimators designed for fixed effects models that either restrict the dimensionality of the underlying group means (Bonhomme et al., 2019) or the nature of the heteroscedasticity governing the errors (Andrews et al., 2008; Jochmans and Weidner, 2016).

Our first contribution is to propose a new variance component estimator designed for unrestricted linear models with heteroscedasticity of unknown form. The estimator is finite sample unbiased and can be written as a naive “plug-in” variance component estimator plus a bias correction term that involves “cross-fit” (Newey and Robins, 2018) estimators of observation-specific error variances. We also develop a representation of the estimator in terms of a covariance between outcomes and a “leave-one-out” generalized prediction (e.g., as in Powell et al., 1989), which allows us to apply recent results on the behavior of second order U-statistics. Building on work

¹See Cochran (1980) for a discussion of the intellectual development of this early work.

by [Achlioptas \(2003\)](#), we propose a random projection method that enables computation of our estimator in very large datasets with little loss of accuracy.

We study the asymptotic behavior of the proposed leave-out estimator in an environment where the number of regressors may be proportional to the sample size: a framework that has alternately been termed “many covariates” ([Cattaneo et al., 2018](#)) or “moderate dimensional” ([Lei et al., 2018](#)) asymptotics. Verifiable design requirements are provided under which the estimator is consistent and we show in an Appendix that these conditions are weaker than those required by jackknife bias correction procedures ([Quenouille, 1949](#); [Hahn and Newey, 2004](#); [Dhaene and Jochmans, 2015](#)). A series of examples is discussed where the leave-out estimator is consistent, while estimators relying on jackknife or homoscedasticity-based bias corrections are not.

We present three sets of theoretical results that enable inference based upon our estimator in a variety of settings. The first result concerns inference on quadratic forms of fixed rank, a problem which typically arises when testing a few linear restrictions in a model with many covariates ([Cattaneo et al., 2018](#)). Familiar examples of such applications include testing that particular regressors are significant in a fixed effects model and conducting inference on the coefficients from a projection of fixed effects onto a low dimensional vector of covariates. Extending classic proposals by [Horn et al. \(1975\)](#) and [MacKinnon and White \(1985\)](#), we show that our leave-out approach can be used to construct an Eicker-White style variance estimator that is unbiased in the presence of unrestricted heteroscedasticity and that enables consistent inference on linear contrasts under weaker design restrictions than those considered by [Cattaneo et al. \(2018\)](#).

Next, we derive a result establishing asymptotic normality of quadratic forms of growing rank. Such quadratic forms typically arise when conducting analysis of variance but also feature in tests of model specification involving a large number of linear restrictions ([Anatolyev, 2012](#); [Chao et al., 2014](#)). The large sample distribution of the estimator is derived using a variant of the arguments in [Chatterjee \(2008\)](#) and [Sølvsten \(2019\)](#) and a standard error estimator is proposed that utilizes sample splitting formulations of the sort considered by [Newey and Robins \(2018\)](#). This standard error estimator is shown to enable consistent inference on quadratic forms of growing rank in the presence of unrestricted heteroscedasticity when the regressor design allows for sample splitting and to provide conservative inference otherwise.

Finally, we present conditions under which the large sample distribution of our estimator is non-pivotal and can be represented by a linear combination of normal and non-central χ^2 random variables, with the non-centralities of the χ^2 terms serving as weakly identified nuisance parameters. This distribution arises in a two-way fixed effects model when there are “bottlenecks” in the mobility network. Such bottlenecks are shown to emerge, for example, when worker mobility is governed by a stochastic block model with limited mobility between blocks. To construct asymptotically valid confidence intervals in the presence of nuisance parameters, we propose inverting a minimum distance test statistic. Critical values are obtained via an application of the procedure of [Andrews](#)

and Mikusheva (2016). The resulting confidence interval is shown to be valid uniformly in the values of the nuisance parameters and to have a closed form representation in many settings, which greatly simplifies its computation.

We illustrate our results with an application of the two-way worker-firm fixed effects model of Abowd et al. (1999) to Italian social security records. The proposed leave-out estimator finds a substantially smaller contribution of firms to wage inequality and much more assortativity in the matching of workers to firms than either the uncorrected plug-in estimator originally considered by Abowd et al. (1999) or the homoscedasticity-based correction procedure of Andrews et al. (2008). When studying panels of length greater than two, we allow for serial correlation in the errors by employing a generalization of our estimator that leaves out all the observations in a worker-firm match. Failing to account for this dependence is shown to yield over-estimates of the variance of firm effects.

Projecting firm effect estimates onto measures of worker age and firm size, we find that older workers tend to be employed at firms offering higher firm wage effects; however, this phenomenon is largely explained by the tendency of older workers to sort to bigger firms. Leave-out standard errors for the coefficients of these linear projections are found to be several times larger than a naive standard error predicated on the assumption that the estimated fixed effects are independent of each other. Stratifying our analysis by birth cohort, we formally reject the null hypothesis that older and younger workers face identical vectors of firm effects. However, the two sets of firm effects are estimated to have a correlation coefficient of nearly 0.9, while the plug-in estimate of correlation is only 0.54.

To assess the accuracy of our asymptotic approximations, we conduct a series of Monte Carlo exercises utilizing the realized mobility patterns of workers between firms. Clear evidence of non-normality arises in the sampling distribution of the estimated variance of firm effects in settings where the worker-firm mobility network is weakly connected. The proposed confidence regions are shown to provide reliable size control in both strongly and weakly identified settings.

1 Unbiased Estimation of Variance Components

Consider the linear model

$$y_i = x_i' \beta + \varepsilon_i \quad (i = 1, \dots, n)$$

where the regressors $x_i \in \mathbb{R}^k$ are non-random and the design matrix $S_{xx} = \sum_{i=1}^n x_i x_i'$ has full rank. The unobserved errors $\{\varepsilon_i\}_{i=1}^n$ are mutually independent and obey $\mathbb{E}[\varepsilon_i] = 0$, but may possess observation specific variances $\mathbb{E}[\varepsilon_i^2] = \sigma_i^2$.

Our object of interest is a quadratic form $\theta = \beta' A \beta$ for some known non-random symmetric

matrix $A \in \mathbb{R}^{k \times k}$ of rank r . Following [Searle et al. \(2009\)](#), when A is positive semi-definite θ is a *variance* component, while when A is non-definite θ may be referred to as a *covariance* component. Note that linear restrictions on the parameter vector β can be formulated in terms of variance components: for a non-random vector v , the null hypothesis $v'\beta = 0$ is equivalent to the restriction $\theta = 0$ when $A = vv'$. Examples from the economics literature where variance components are of direct interest are discussed in [Section 2](#).

1.1 Estimator

A naive plug-in estimator of θ is given by the quadratic form $\hat{\theta}_{\text{PI}} = \hat{\beta}'A\hat{\beta}$, where $\hat{\beta} = S_{xx}^{-1} \sum_{i=1}^n x_i y_i$ denotes the Ordinary Least Squares (OLS) estimator of β . Estimation error in $\hat{\beta}$ leads the plug-in estimator to exhibit a bias involving a linear combination of the unknown variances $\{\sigma_i^2\}_{i=1}^n$. Specifically, standard results on quadratic forms imply that $\mathbb{E}[\hat{\theta}] = \theta + \text{trace}(A\mathbb{V}[\hat{\beta}])$, where

$$\text{trace}\left(A\mathbb{V}[\hat{\beta}]\right) = \sum_{i=1}^n B_{ii}\sigma_i^2 \quad \text{and} \quad B_{ii} = x_i' S_{xx}^{-1} A S_{xx}^{-1} x_i.$$

As discussed in [Section 2](#), this bias can be particularly severe when the dimension of the regressors k is large relative to the sample size.

A bias correction can be motivated by observing that an unbiased estimator of the i -th error variance is

$$\hat{\sigma}_i^2 = y_i \left(y_i - x_i' \hat{\beta}_{-i} \right),$$

where $\hat{\beta}_{-i} = (S_{xx} - x_i x_i')^{-1} \sum_{\ell \neq i} x_\ell y_\ell$ denotes the leave- i -out OLS estimator of β . This insight suggests the following bias-corrected estimator of θ :

$$\hat{\theta} = \hat{\beta}'A\hat{\beta} - \sum_{i=1}^n B_{ii}\hat{\sigma}_i^2. \tag{1}$$

While [Newey and Robins \(2018\)](#) observe that “cross-fit” covariances relying on sample splitting can be used to remove bias of the sort considered here, we are not aware of existing estimators involving the leave-one-out estimators $\{\hat{\sigma}_i^2\}_{i=1}^n$.

One can also motivate $\hat{\theta}$ via a change of variables argument. Letting $\tilde{x}_i = A S_{xx}^{-1} x_i$ denote a vector of “generalized” regressors, we can write

$$\theta = \beta' A \beta = \beta' S_{xx} S_{xx}^{-1} A \beta = \sum_{i=1}^n \beta' x_i \tilde{x}_i' \beta = \sum_{i=1}^n \mathbb{E} [y_i \tilde{x}_i' \beta].$$

This observation suggests using the unbiased *leave-out* estimator

$$\hat{\theta} = \sum_{i=1}^n y_i \tilde{x}'_i \hat{\beta}_{-i}. \quad (2)$$

Note that direct computation of $\hat{\beta}_{-i}$ can be avoided by exploiting the representation

$$y_i - x'_i \hat{\beta}_{-i} = \frac{y_i - x'_i \hat{\beta}}{1 - P_{ii}}, \quad (3)$$

where $P_{ii} = x'_i S_{xx}^{-1} x_i$ gives the leverage of observation i . Applying the Sherman-Morrison-Woodbury formula (Woodbury, 1949; Sherman and Morrison, 1950), this representation also reveals that (1) and (2) are numerically equivalent:

$$y_i \tilde{x}'_i \hat{\beta}_{-i} = \underbrace{y_i \tilde{x}'_i S_{xx}^{-1} \sum_{\ell \neq i} x_\ell y_\ell}_{=y_i \tilde{x}'_i \hat{\beta} - B_{ii} y_i^2} + \underbrace{\frac{y_i \tilde{x}'_i S_{xx}^{-1} x_i x'_i S_{xx}^{-1}}{1 - P_{ii}} \sum_{\ell \neq i} x_\ell y_\ell}_{=B_{ii} y_i x'_i \hat{\beta}_{-i}} = y_i \tilde{x}'_i \hat{\beta} - B_{ii} \hat{\sigma}_i^2.$$

A similar combination of a change of variables argument and a leave-one-out estimator was used by Powell et al. (1989) in the context of weighted average derivatives. The JIVE estimators proposed by Phillips and Hale (1977) and Angrist et al. (1999) also use a leave-one-out estimator, though without the change of variables.²

Remark 1. The $\{\hat{\sigma}_i^2\}_{i=1}^n$ can also be used to construct an unbiased variance estimator

$$\hat{V}[\hat{\beta}] = S_{xx}^{-1} \left(\sum_{i=1}^n x_i x'_i \hat{\sigma}_i^2 \right) S_{xx}^{-1}.$$

Section 3 shows that $\hat{V}[\hat{\beta}]$ can be used to perform asymptotically valid inference on linear contrasts in settings where existing Eicker-White estimators fail. Specifically, $\hat{V}[\hat{\beta}]$ leads to valid inference under conditions where the MINQUE estimator of Rao (1970) and the MINQUE-type estimator of Cattaneo et al. (2018) do not exist (see, e.g., Horn et al., 1975; Verdier, 2017).

Remark 2. The quantity $\hat{V}[\hat{\beta}]$ is closely related to the HC2 variance estimator of MacKinnon and White (1985). While the HC2 estimator employs observation specific variance estimators $\hat{\sigma}_{i,HC2}^2 = \frac{(y_i - x'_i \hat{\beta})^2}{1 - P_{ii}}$, $\hat{V}[\hat{\beta}]$ relies instead on $\hat{\sigma}_i^2 = \frac{y_i (y_i - x'_i \hat{\beta})}{1 - P_{ii}}$

Remark 3. In some cases it may be important to allow dependence in the errors in addition to heteroscedasticity. A common case arises when the data are organized into mutually exclusive

² The object of interest in JIVE estimation is a *ratio* of quadratic forms $\beta_1' S_{xx} \beta_2 / \beta_2' S_{xx} \beta_2$ in the two-equation model $y_{ij} = x'_i \beta_j + \varepsilon_{ij}$ for $j = 1, 2$. When no covariates are present, using leave-out estimators of both the numerator and denominator of this ratio yields the JIVE1 estimator of Angrist et al. (1999).

and independent “clusters” within which the errors may be dependent (Moulton, 1986). The same change of variables argument implies that an estimator of the form $\sum_{i=1}^n y_i \tilde{x}_i' \hat{\beta}_{-c(i)}$ will be unbiased in such settings, where $\hat{\beta}_{-c(i)}$ is the OLS estimator obtained after leaving out all observations in the cluster to which observation i belongs.

1.2 Large Scale Computation

From (1) and (3), computation of $\hat{\theta}$ relies on the values $\{B_{ii}, P_{ii}\}_{i=1}^n$. Section 2 provides some canonical examples where these quantities can be computed in closed form. When closed forms are unavailable, a number of options exist for accelerating computation. For example, in the empirical application of Section 8, we make use of a preconditioned conjugate gradient algorithm suggested by Koutis et al. (2011) to compute exact leave-out variance decompositions in a two-way fixed effects model involving roughly one million observations and hundreds of thousands of parameters (see Appendix B.3 for details). However, in very large scale applications involving tens or hundreds of millions of parameters, exact computation of $\{B_{ii}, P_{ii}\}_{i=1}^n$ is likely to become infeasible. Fortunately, it is possible to quickly approximate $\hat{\theta}$ in such settings using a variant of the random projection method introduced by Achlioptas (2003). We refer to this method as the Johnson-Lindenstrauss approximation (JLA) for its connection to the work of Johnson and Lindenstrauss (1984).

JLA can be described by the following algorithm: fix a $p \in \mathbb{N}$ and generate the matrices $R_B, R_P \in \mathbb{R}^{p \times n}$, where (R_B, R_P) are composed of mutually independent Rademacher random variables that are independent of the data, i.e., their entries take the values 1 and -1 with probability $1/2$. Next decompose A into $A = \frac{1}{2}(A_1' A_2 + A_2' A_1)$ for $A_1, A_2 \in \mathbb{R}^{n \times k}$ where $A_1 = A_2$ if A is positive semi-definite.³ Let

$$\hat{P}_{ii} = \frac{1}{p} \left\| R_P X S_{xx}^{-1} x_i \right\|^2 \quad \text{and} \quad \hat{B}_{ii} = \frac{1}{p} \left(R_B A_1 S_{xx}^{-1} x_i \right)' \left(R_B A_2 S_{xx}^{-1} x_i \right)$$

where $X = (x_1, \dots, x_n)'$. The Johnson-Lindenstrauss approximation to $\hat{\theta}$ is

$$\hat{\theta}_{JLA} = \hat{\beta}' A \hat{\beta} - \sum_{i=1}^n \hat{B}_{ii} \hat{\sigma}_{i,JLA}^2,$$

where $\hat{\sigma}_{i,JLA}^2 = \frac{y_i (y_i - x_i' \hat{\beta})}{1 - \hat{P}_{ii}} \left(1 - \frac{1}{p} \frac{3\hat{P}_{ii}^3 + \hat{P}_{ii}^2}{1 - \hat{P}_{ii}} \right)$. The term $\frac{1}{p} \frac{3\hat{P}_{ii}^3 + \hat{P}_{ii}^2}{1 - \hat{P}_{ii}}$ removes a non-linearity bias introduced by approximating P_{ii} .

Section 1.5 establishes asymptotic equivalence between $\hat{\theta}_{JLA}$ and $\hat{\theta}$. Appendix B.3 discusses implementation details and numerically illustrates the trade-off between computation time and the

³Interpretable choices of A_1 and A_2 are typically suggested by the structure of the problem; see, for instance, the discussion in Example 4 of Section 2.

bias introduced by JLA for different choices of p under a range of sample sizes. Notably, we show that JLA allows us to accurately compute a variance decomposition in a two-way fixed effects model with roughly 15 million parameters – a scale comparable to the study of [Card et al. \(2013\)](#) – in under an hour. A MATLAB package ([Kline et al., 2019](#)) implementing both the exact and JLA versions of our estimator in the two-way fixed effects model is available online.

1.3 Relation to Existing Approaches

As discussed in Section 2, several literatures make use of bias corrections nominally predicated on homoscedasticity. A common “homoscedasticity-only” estimator takes the form

$$\hat{\theta}_{\text{HO}} = \hat{\beta}' A \hat{\beta} - \sum_{i=1}^n B_{ii} \hat{\sigma}_{\text{HO}}^2 \quad (4)$$

where $\hat{\sigma}_{\text{HO}}^2 = \frac{1}{n-k} \sum_{i=1}^n (y_i - x_i' \hat{\beta})^2$ is the degrees-of-freedom corrected variance estimator. A sufficient condition for unbiasedness of $\hat{\theta}_{\text{HO}}$ is that there be no empirical covariance between σ_i^2 and (B_{ii}, P_{ii}) . This restriction is in turn implied by the special cases of *homoscedasticity* where σ_i^2 does not vary with i or *balanced design* where (B_{ii}, P_{ii}) does not vary with i . In general, however, this estimator will tend to be biased (see, e.g., [Scheffe, 1959](#), chapter 10, or Appendix C.1.3).

A second estimator, closely related to $\hat{\theta}$, relies upon a jackknife bias-correction ([Quenouille, 1949](#)) of the plug-in estimator. This estimator can be written

$$\hat{\theta}_{\text{JK}} = n \hat{\theta}_{\text{PI}} - \frac{n-1}{n} \sum_{i=1}^n \hat{\theta}_{\text{PI},-i} \quad \text{where} \quad \hat{\theta}_{\text{PI},-i} = \hat{\beta}'_{-i} A \hat{\beta}_{-i}.$$

In Appendix C.1.3 we illustrate that jackknife bias-correction tends to over-correct and produce a first order bias in the opposite direction of the bias in the plug-in estimator. This is analogous to the upward bias in the jackknife estimator of $\mathbb{V}[\hat{\beta}]$ which was derived by [Efron and Stein \(1981\)](#) and shown by [El Karoui and Purdom \(2018\)](#) to be of first order importance for inference with many Gaussian regressors.

There are several proposed adaptations of the jackknife to long panels that can decrease bias under stationarity restrictions on the regressors. Letting $t(i) \in \{1, \dots, T\}$ denote the time period in which an observation is observed, we can write the panel jackknife of [Hahn and Newey \(2004\)](#) as

$$\hat{\theta}_{\text{PJK}} = T \hat{\theta}_{\text{PI}} - \frac{T-1}{T} \sum_{t=1}^T \hat{\theta}_{\text{PI},-t} \quad \text{where} \quad \hat{\theta}_{\text{PI},-t} = \hat{\beta}'_{-t} A \hat{\beta}_{-t}$$

and $\hat{\beta}_{-t} = (\sum_{i:t(i) \neq t} x_i x_i')^{-1} \sum_{i:t(i) \neq t} x_i y_i$ is the OLS estimator that excludes all observations from

period t . [Dhaene and Jochmans \(2015\)](#) propose a closely related split panel jackknife

$$\hat{\theta}_{\text{SPJK}} = 2\hat{\theta}_{\text{PI}} - \frac{\hat{\theta}_{\text{PI},1} + \hat{\theta}_{\text{PI},2}}{2} \quad \text{where} \quad \hat{\theta}_{\text{PI},j} = \hat{\beta}'_j A \hat{\beta}_j$$

and $\hat{\beta}_1$ (and $\hat{\beta}_2$) are OLS estimators based on the first half (and the last half) of an even number of time periods. In [Appendix C.1.3](#), we illustrate how short panels can lead these adaptations of the jackknife to produce first order biases in the opposite direction of the bias in the plug-in estimator.

1.4 Finite Sample Properties

We now study the finite sample properties of the leave-out estimator $\hat{\theta}$ and its infeasible analogue $\theta^* = \hat{\beta}' A \hat{\beta} - \sum_{i=1}^n B_{ii} \sigma_i^2$, which uses knowledge of the individual error variances. First, we note that $\hat{\theta}$ is unbiased whenever each of the leave-one-out estimators $\hat{\beta}_{-i}$ exists, which can equivalently be expressed as the requirement that $\max_i P_{ii} < 1$. This condition turns out to also be necessary for the existence of unbiased estimators, which highlights the need for additional restrictions on the model or sample whenever some leverages equal one.

Lemma 1. *1. If $\max_i P_{ii} < 1$, then $\mathbb{E}[\hat{\theta}] = \theta$.*

2. Unbiased estimators of $\theta = \beta' A \beta$ exist for all A if and only if $\max_i P_{ii} < 1$.

Next, we show that when the errors are normal, the infeasible estimator θ^* is a weighted sum of a series of non-central χ^2 random variables. This second result provides a useful point of departure for our asymptotic approximations and highlights the important role played by the matrix

$$\tilde{A} = S_{xx}^{-1/2} A S_{xx}^{-1/2},$$

which encodes features of both the target parameter (which is defined by A) and the design matrix S_{xx} .

Let $\lambda_1, \dots, \lambda_r$ denote the non-zero eigenvalues of \tilde{A} , where $\lambda_1^2 \geq \dots \geq \lambda_r^2$ and each eigenvalue appears as many times as its algebraic multiplicity. We use Q to refer to the corresponding matrix of orthonormal eigenvectors so that $\tilde{A} = Q D Q'$ where $D = \text{diag}(\lambda_1, \dots, \lambda_r)$. With these definitions we have

$$\hat{\beta}' A \hat{\beta} = \sum_{\ell=1}^r \lambda_{\ell} \hat{b}_{\ell}^2,$$

where $\hat{b} = (\hat{b}_1, \dots, \hat{b}_r)' = Q' S_{xx}^{1/2} \hat{\beta}$ contains r linear combinations of the elements in $\hat{\beta}$. The random vector \hat{b} and the eigenvalues $\lambda_1, \dots, \lambda_r$ are central to both the finite sample distribution provided below in [Lemma 2](#) and the asymptotic properties of $\hat{\theta}$ as studied in [Sections 3–5](#). Each eigenvalue

of \tilde{A} can be thought of as measuring how strongly θ depends on a particular linear combination of the elements in β relative to the difficulty of estimating that combination (as summarized by S_{xx}^{-1}). As discussed in Section 5, when a few of these eigenvalues are large relative to the others, a form of weak identification can arise.

Lemma 2. *If $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, then*

1. $\hat{b} \sim \mathcal{N}\left(b, \mathbb{V}[\hat{b}]\right)$ where $b = Q' S_{xx}^{1/2} \beta$,
2. $\theta^* = \sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell]\right)$

The distribution of θ^* is a sum of r potentially dependent non-central χ^2 random variables with non-centralities $b = (b_1, \dots, b_r)'$. In the special case of homoscedasticity ($\sigma_i^2 = \sigma^2$) and no signal ($b = 0$) we have that $\hat{b} \sim \mathcal{N}\left(0, \sigma^2 I_r\right)$, which implies that the distribution of θ^* is a weighted sum of r independent central χ^2 random variables. The weights are the eigenvalues of \tilde{A} , therefore consistency of θ^* follows whenever the sum of the squared eigenvalues converges to zero. The next subsection establishes that the leave-out estimator remains consistent when a signal is present ($b \neq 0$) and the errors exhibit unrestricted heteroscedasticity.

1.5 Consistency

We now drop the normality assumption and provide conditions under which $\hat{\theta}$ remains consistent. To accommodate high dimensionality of the regressors we allow all parts of the model to change with n :

$$y_{i,n} = x'_{i,n} \beta_n + \varepsilon_{i,n} \quad (i = 1, \dots, n)$$

where $x_{i,n} \in \mathbb{R}^{k_n}$, $S_{xx,n} = \sum_{i=1}^n x_{i,n} x'_{i,n}$, $\mathbb{E}[\varepsilon_{i,n}] = 0$, $\mathbb{E}[\varepsilon_{i,n}^2] = \sigma_{i,n}^2$ and $\theta_n = \beta_n' A_n \beta_n$ for some sequence of known non-random symmetric matrices $A_n \in \mathbb{R}^{k_n \times k_n}$ of rank r_n . By treating $x_{i,n}$ and A_n as sequences of constants, all uncertainty derives from the disturbances $\{\varepsilon_{i,n} : 1 \leq i \leq n, n \geq 1\}$. This *conditional* perspective is common in the statistics literatures on ANOVA (Scheffe, 1959; Searle et al., 2009) and allows us to be agnostic about the potential dependency among the $\{x_{i,n}\}_{i=1}^n$ and A_n .⁴ Following standard practice we drop the n subscript in what follows. All limits are taken as n goes to infinity unless otherwise noted.

Our analysis makes heavy use of the following assumptions.

⁴An *unconditional* analysis might additionally impose distributional assumptions on A_n and consider $\bar{\theta} = \beta' \mathbb{E}_{A_n}[A_n] \beta$ as the object of interest. The uncertainty in $\hat{\theta} - \bar{\theta}$ can always be decomposed into components attributable to $\hat{\theta} - \theta$ and $\theta - \bar{\theta}$. Because the behavior of $\theta - \bar{\theta}$ depends entirely on model choices, we leave such an analysis to future work.

Assumption 1. (i) $\max_i (\mathbb{E}[\varepsilon_i^4] + \sigma_i^{-2}) = O(1)$, (ii) there exist a $c < 1$ such that $\max_i P_{ii} \leq c$ for all n , and (iii) $\max_i (x_i' \beta)^2 = O(1)$.

Part (i) of this condition limits the thickness of the tails in the error distribution, as is typically required for OLS estimation (see, e.g., Cattaneo et al., 2018, page 10). The bounds on $(x_i' \beta)^2$ and P_{ii} imply that $\hat{\sigma}_i^2$ has bounded variance. Part (iii) is a technical condition that can be relaxed to allow $\max_i (x_i' \beta)^2$ to increase slowly with sample size as discussed further in Section 7. From (ii) it follows that $\frac{k}{n} \leq c < 1$ for all n .

The following Lemma establishes consistency of $\hat{\theta}$.

Lemma 3. *If Assumption 1 and one of the following conditions hold, then $\hat{\theta} - \theta \xrightarrow{p} 0$.*

(i) *A is positive semi-definite, $\theta = \beta' A \beta = O(1)$, and $\text{trace}(\tilde{A}^2) = \sum_{\ell=1}^r \lambda_{\ell}^2 = o(1)$.*

(ii) *$A = \frac{1}{2}(A_1' A_2 + A_2' A_1)$ where $\theta_1 = \beta' A_1' A_1 \beta$ and $\theta_2 = \beta' A_2' A_2 \beta$ satisfy (i).*

The first condition of Lemma 3 establishes consistency of variance components given boundedness of θ and a joint condition on the design matrix S_{xx} and the matrix A . The second condition shows that consistency of covariance components follows from consistency of variance components that dominate them via the Cauchy-Schwarz inequality, i.e., $\theta^2 = (\beta' A_1' A_2 \beta)^2 \leq \theta_1 \theta_2$. In several of the examples discussed in the next section, $\text{trace}(\tilde{A}^2)$ is of order r/n^2 , which is necessarily small in large samples. A more extensive discussion of primitive conditions that yield $\text{trace}(\tilde{A}^2) = o(1)$ is provided in Section 7.

We conclude this section by establishing asymptotic equivalence between the leave-out estimator $\hat{\theta}$ and its approximation $\hat{\theta}_{JLA}$ under the condition that p^4 is large relative to sample size.

Lemma 4. *If Assumption 1 is satisfied, $n/p^4 = o(1)$, $\mathbb{V}[\hat{\theta}]^{-1} = O(n)$, and one of the following conditions hold, then $\mathbb{V}[\hat{\theta}]^{-1/2}(\hat{\theta}_{JLA} - \hat{\theta} - B_p) = o_p(1)$ where $|B_p| \leq \frac{1}{p} \sum_{i=1}^n P_{ii}^2 |B_{ii}| \sigma_i^2$.*

(i) *A is positive semi-definite and $\mathbb{E}[\hat{\beta}' A \hat{\beta}] - \theta = \sum_{i=1}^n B_{ii} \sigma_i^2 = O(1)$.*

(ii) *$A = \frac{1}{2}(A_1' A_2 + A_2' A_1)$ where $\theta_1 = \beta' A_1' A_1 \beta$ and $\theta_2 = \beta' A_2' A_2 \beta$ satisfy (i) and $\frac{\mathbb{V}[\hat{\theta}_1] \mathbb{V}[\hat{\theta}_2]}{n \mathbb{V}[\hat{\theta}]^2} = O(1)$.*

Lemma 4 requires that $\hat{\theta}$ is not super-consistent and that the bias in the plug-in estimator is asymptotically bounded, assumptions which can be shown to be satisfied in the examples introduced in the next section. For variance components, the Lemma characterizes an approximation bias B_p in $\hat{\theta}_{JLA}$ of order $1/p$ and provides an interpretable bound on B_p : the approximation bias is at most $1/p$ times the bias in the plug in estimator $\hat{\beta}' A \hat{\beta}$. For covariance components, asymptotic equivalence follows when the variance components defined by $A_1' A_1$ and $A_2' A_2$ do not converge at substantially slower rates than $\hat{\theta}$. Under this condition, the approximation bias is at most $1/p$ times the average of the biases in the plug in estimators $\hat{\beta}' A_1' A_1 \hat{\beta}$ and $\hat{\beta}' A_2' A_2 \hat{\beta}$.

These bounds on the approximation bias suggests that a p of a few hundred should suffice for point estimation. However, unless $n/p^2 = o(1)$, the resulting approximation bias needs to be accounted for when conducting inference. Specifically, one can lengthen the tails of the confidence sets proposed in Sections 4 and 6 by $\frac{1}{p} \sum_{i=1}^n \hat{P}_{ii}^2 |\hat{B}_{ii}| \hat{\sigma}_{i,JLA}^2$ when relying on JLA.

2 Examples

We now consider four commonly encountered empirical examples where our proposed estimation strategy provides an advantage over existing methods.

Example 1 (Coefficient of determination).

Sewall Wright (1921) proposed measuring the explanatory power of a linear model using the coefficient of determination. When x_i includes an intercept, the object of interest and its corresponding plug-in estimator can be written

$$R^2 = \frac{\beta' A \beta}{\beta' A \beta + \frac{1}{n} \sum_{i=1}^n \sigma_i^2} = \frac{\sigma_{X\beta}^2}{\sigma_y^2} \quad \text{and} \quad \hat{R}_{\text{PI}}^2 = \frac{\hat{\beta}' A \hat{\beta}}{\frac{1}{n} \sum_{i=1}^n (y_i - \bar{y})^2} = \frac{\hat{\sigma}_{X\beta, \text{PI}}^2}{\hat{\sigma}_y^2}$$

where

$$A = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})', \quad \bar{x} = \frac{1}{n} \sum_{i=1}^n x_i, \quad \bar{y} = \frac{1}{n} \sum_{i=1}^n y_i.$$

Theil (1961) noted that the plug-in estimator of $\sigma_{X\beta}^2$ is biased and proposed an adjusted R^2 measure that utilizes the homoscedasticity-only estimator in (4). The above choice of A yields $B_{ii} = \frac{1}{n}(P_{ii} - \frac{1}{n})$, which implies $\sum_{i=1}^n B_{ii} = \frac{k-1}{n}$. Hence, Theil's proposal can be written

$$\hat{R}_{\text{adj}}^2 = \frac{\hat{\sigma}_{X\beta, \text{HO}}^2}{\hat{\sigma}_y^2} = \frac{\hat{\beta}' A \hat{\beta} - \frac{k-1}{n} \hat{\sigma}_{\text{HO}}^2}{\hat{\sigma}_y^2}.$$

A rearrangement gives the familiar representation $\frac{1 - \hat{R}_{\text{adj}}^2}{1 - \hat{R}_{\text{PI}}^2} = \frac{n-1}{n-k}$ which highlights that the adjusted estimator of R^2 relates to the unadjusted one through a degrees-of-freedom correction.

The leave-out estimator of $\sigma_{X\beta}^2$ allows for unrestricted heteroscedasticity and can be found by noting that $\tilde{x}_i = A S_{xx}^{-1} x_i = \frac{1}{n}(x_i - \bar{x})$, which yields

$$\hat{R}^2 = \frac{\hat{\sigma}_{X\beta}^2}{\hat{\sigma}_y^2} \quad \text{where} \quad \hat{\sigma}_{X\beta}^2 = \frac{1}{n} \sum_{i=1}^n y_i (x_i - \bar{x})' \hat{\beta}_{-i}.$$

In general, this estimator does not have an interpretation in terms of degrees-of-freedom corrections. Instead, the explanatory power of the linear model is assessed using the empirical covariance between leave-one-out predictions $(x_i - \bar{x})' \hat{\beta}_{-i}$ and the left out observation y_i .

Example 2 (Analysis of covariance).

Since the work of [Fisher \(1925\)](#), it has been common to summarize the effects of experimentally assigned treatments on outcomes with estimates of variance components. Consider a dataset comprised of observations on N groups with T_g observations in the g -th group. The “analysis of covariance” model posits that outcomes can be written

$$y_{gt} = \alpha_g + x'_{gt}\delta + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T_g \geq 2),$$

where α_g is a group effect and x_{gt} is a vector of strictly exogenous covariates.

A prominent example comes from [Chetty et al. \(2011\)](#) who study the adult earnings y_{gt} of $n = \sum_{g=1}^N T_g$ students assigned experimentally to one of N different classrooms. Each student also has a vector of predetermined background characteristics x_{gt} . The variability in student outcomes attributable to classrooms can be written:

$$\sigma_\alpha^2 = \frac{1}{n} \sum_{g=1}^N T_g (\alpha_g - \bar{\alpha})^2$$

where $\bar{\alpha} = \frac{1}{n} \sum_{g=1}^N T_g \alpha_g$ gives the (enrollment-weighted) mean classroom effect.

This model and object of interest can be written in the notation of the preceding section ($y_i = x'_i\beta + \varepsilon_i$ and $\sigma_\alpha^2 = \beta' A \beta$) by letting $i = i(g, t)$ where $i(\cdot, \cdot)$ is bijective with inverse denoted $(g(\cdot), t(\cdot))$, $y_i = y_{gt}$, $\varepsilon_i = \varepsilon_{gt}$,

$$x_i = (d'_i, x'_{gt})', \quad \beta = (\alpha', \delta')', \quad \alpha = (\alpha_1, \dots, \alpha_N)', \quad d_i = (\mathbf{1}_{\{g=1\}}, \dots, \mathbf{1}_{\{g=N\}})'$$

and

$$A = \begin{bmatrix} A_{dd} & 0 \\ 0 & 0 \end{bmatrix} \quad \text{where} \quad A_{dd} = \frac{1}{n} \sum_{i=1}^n (d_i - \bar{d})(d_i - \bar{d})', \quad \bar{d} = \frac{1}{n} \sum_{i=1}^n d_i.$$

[Chetty et al. \(2011\)](#) estimate σ_α^2 using a random effects ANOVA estimator (see e.g., [Searle et al., 2009](#)) which is of the homoscedasticity-only type given in (4). As discussed in Section 1 and Appendix C.1.3, this estimator is in general first order biased when the errors are heteroscedastic and group sizes are unbalanced.

Special Case: No Common Regressors When there are no common regressors ($x_{gt} = 0$ for all g, t), the leave-out estimator of σ_α^2 has a particularly simple representation:

$$\hat{\sigma}_\alpha^2 = \frac{1}{n} \sum_{g=1}^N \left(T_g (\hat{\alpha}_g - \hat{\hat{\alpha}})^2 - \left(1 - \frac{T_g}{n} \right) \hat{\sigma}_g^2 \right) \quad \text{for} \quad \hat{\sigma}_g^2 = \frac{1}{T_g - 1} \sum_{t=1}^{T_g} (y_{gt} - \hat{\alpha}_g)^2, \quad (5)$$

where $\hat{\alpha}_g = \frac{1}{T_g} \sum_{t=1}^{T_g} y_{gt}$, and $\hat{\hat{\alpha}} = \frac{1}{n} \sum_{g=1}^N T_g \hat{\alpha}_g$. This representation shows that if the model

consists only of group specific intercepts, then the leave-out estimator relies on group level degrees-of-freedom corrections. The statistic in (5) was analyzed by Akritas and Papadatos (2004) in the context of testing the null hypothesis that $\sigma_\alpha^2 = 0$ while allowing for heteroscedasticity at the group level.

Covariance Representation Another instructive representation of the leave-out estimator is in terms of the empirical covariance

$$\hat{\sigma}_\alpha^2 = \sum_{i=1}^n y_i \tilde{d}_i' \hat{\alpha}_{-i} \quad \text{where} \quad \hat{\beta}_{-i} = (\hat{\alpha}'_{-i}, \hat{\delta}'_{-i}).$$

The generalized regressor \tilde{d}_i can be described as follows: if there are no common regressors then $\tilde{d}_i = \frac{1}{n}(d_i - \bar{d})$, which is analogous to Example 1. If the model includes common regressors then $\tilde{d}_i = \frac{1}{n} \left((d_i - \bar{d}) - \hat{\Gamma}'(x_{g(i)t(i)} - \bar{x}_{g(i)}) \right)$ where $\bar{x}_g = \frac{1}{T_g} \sum_{t=1}^{T_g} x_{gt}$ and $\hat{\Gamma}$ is the coefficient vector from an instrumental variables (IV) regression of $d_i - \bar{d}$ on $x_{g(i)t(i)} - \bar{x}_{g(i)}$ using $x_{g(i)t(i)}$ as an instrument. The IV residual \tilde{d}_i is uncorrelated with $x_{g(i)t(i)}$ and the covariance between d_i and \tilde{d}_i is A_{dd} , which ensures that the empirical covariance between $y_i = d_i' \alpha + x_{g(i)t(i)}' \delta + \varepsilon_i$ and the generalized prediction $\tilde{d}_i' \hat{\alpha}_{-i}$ is an unbiased estimator of σ_α^2 .

Example 3 (Random coefficients).

Group memberships are often modeled as influencing slopes in addition to intercepts (Kuh, 1959; Hildreth and Houck, 1968; Raudenbush and Bryk, 2002; Arellano and Bonhomme, 2011; Graham and Powell, 2012; Graham et al., 2018). Consider the following “random coefficient” model:

$$y_{gt} = \alpha_g + z_{gt} \gamma_g + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T_g \geq 3).$$

An influential example comes from Raudenbush and Bryk (1986), who model student mathematics scores as a “hierarchical” linear function of socioeconomic status (SES) with school-specific intercepts ($\alpha_g \in \mathbb{R}$) and slopes ($\gamma_g \in \mathbb{R}$). Letting $\bar{\gamma} = \frac{1}{n} \sum_{g=1}^N T_g \gamma_g$ for $n = \sum_{g=1}^N T_g$, the student-weighted variance of slopes can be written:

$$\sigma_\gamma^2 = \frac{1}{n} \sum_{g=1}^N T_g (\gamma_g - \bar{\gamma})^2.$$

In the notation of the preceding section we can write $y_i = x_i' \beta + \varepsilon_i$ and $\sigma_\gamma^2 = \beta' A \beta$ where

$$x_i = (d_i', d_i' z_{gt}')', \quad \beta = (\alpha', \gamma')', \quad \gamma = (\gamma_1, \dots, \gamma_N)', \quad A = \begin{bmatrix} A_{dd} & 0 \\ 0 & 0 \end{bmatrix}$$

for y_i , ε_i , d_i , A_{dd} , and α as in the preceding example.

Raudenbush and Bryk (1986) use a maximum likelihood estimator of σ_γ^2 predicated upon nor-

mality and homoscedastic errors. [Swamy \(1970\)](#) considers an estimator of σ_γ^2 that relies on group-level degrees-of-freedom corrections and is unbiased when the error variance is allowed to vary at the group level, but not with the level of z_{gt} . By contrast, the leave-out estimator is unbiased under arbitrary patterns of heteroscedasticity.

Covariance Representation The leave-out estimator can be represented in terms of the empirical covariance

$$\hat{\sigma}_\gamma^2 = \sum_{i=1}^n y_i \tilde{z}_i \tilde{d}'_i \hat{\gamma}_{-i} \quad \text{where} \quad \tilde{d}_i = \frac{1}{n}(d_i - \bar{d}), \quad \tilde{z}_i = \frac{z_{g(i)t(i)} - \bar{z}_{g(i)}}{\sum_{t=1}^{T_{g(i)}} (z_{g(i)t} - \bar{z}_{g(i)})^2},$$

and $\bar{z}_g = \frac{1}{T_g} \sum_{t=1}^{T_g} z_{gt}$. Demeaning $z_{g(i)t(i)}$ at the group level makes $\tilde{d}_i \tilde{z}_i$ uncorrelated with d_i and scaling by the group variability in $z_{g(i)t}$ ensures that the covariance between $\tilde{d}_i \tilde{z}_i$ and $d_i z_{g(i)t(i)}$ is A_{dd} . This implies that the empirical covariance between $y_i = d'_i \alpha + z_{g(i)t(i)} d'_i \gamma + \varepsilon_i$ and the generalized prediction $\tilde{z}_i \tilde{d}'_i \hat{\gamma}_{-i}$ is an unbiased estimator of σ_γ^2 .

Example 4 (Two-way fixed effects).

Economists often study settings where units possess two or more group memberships, some of which can change over time. A prominent example comes from [Abowd et al. \(1999\)](#) (henceforth AKM) who propose a panel model of log wage determination that is additive in worker and firm fixed effects. This so-called “two-way” fixed effects model takes the form:

$$y_{gt} = \alpha_g + \psi_{j(g,t)} + x'_{gt} \delta + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T_g \geq 2) \quad (6)$$

where the function $j(\cdot, \cdot) : \{1, \dots, N\} \times \{1, \dots, \max_g T_g\} \rightarrow \{0, \dots, J\}$ allocates each of $n = \sum_{g=1}^N T_g$ person-year observations to one of $J+1$ firms. Here α_g is a “person effect”, $\psi_{j(g,t)}$ is a “firm effect”, x_{gt} is a time-varying covariate, and ε_{gt} is a time-varying error. In this context, the mean zero assumption on the errors ε_{gt} can be thought of as requiring both the common covariates x_{gt} and the firm assignments $j(\cdot, \cdot)$ to obey a strict exogeneity condition.

Interest in such models often centers on understanding how much of the variability in log wages is attributable to firms (see, e.g., [Card et al., 2013](#); [Song et al., 2017](#)). AKM summarize the firm contribution to wage inequality via the following two parameters:

$$\sigma_\psi^2 = \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} (\psi_{j(g,t)} - \bar{\psi})^2 \quad \text{and} \quad \sigma_{\alpha,\psi} = \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} (\psi_{j(g,t)} - \bar{\psi}) \alpha_g$$

where $\bar{\psi} = \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} \psi_{j(g,t)}$. The variance component σ_ψ^2 measures the contribution of firm wage variability to inequality, while the covariance component $\sigma_{\alpha,\psi}$ measures the additional contribution of systematic sorting of high wage workers to high wage firms.

To represent this model and the corresponding objects of interest in the notation of the preceding section ($y_i = x_i' \beta + \varepsilon_i$, $\sigma_\psi^2 = \beta' A_\psi \beta$, and $\sigma_{\alpha, \psi} = \beta' A_{\alpha, \psi} \beta$), let

$$x_i = (d_i', f_i', x_{gt}')', \quad \beta = (\alpha', \psi', \delta')', \quad \alpha = (\alpha_1, \dots, \alpha_N)' + \mathbf{1}'_N \psi_0, \quad \psi = (\psi_1 \dots, \psi_J)' - \mathbf{1}'_J \psi_0,$$

for y_i , ε_i , and d_i as in the preceding examples, $f_i = (\mathbf{1}_{\{j(g,t)=1\}}, \dots, \mathbf{1}_{\{j(g,t)=J\}})'$,

$$A_\psi = \begin{bmatrix} 0 & 0 & 0 \\ 0 & A_{ff} & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{where} \quad A_{ff} = \frac{1}{n} \sum_{i=1}^n (f_i - \bar{f})(f_i - \bar{f})', \quad \bar{f} = \frac{1}{n} \sum_{i=1}^n f_i,$$

and

$$A_{\alpha, \psi} = \frac{1}{2} \begin{bmatrix} 0 & A_{df} & 0 \\ A_{df}' & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{where} \quad A_{df} = \frac{1}{n} \sum_{i=1}^n (d_i - \bar{d})(f_i - \bar{f})'.$$

Computation of the Johnson-Lindenstrauss approximation can be facilitated using the representations $A_\psi = A_f' A_f$ and $A_{\alpha, \psi} = \frac{1}{2}(A_d' A_f + A_f' A_d)$ where

$$A_f' = \frac{1}{\sqrt{n}} \begin{bmatrix} 0 & 0 & 0 \\ f_1 - \bar{f} & \dots & f_n - \bar{f} \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad A_d' = \frac{1}{\sqrt{n}} \begin{bmatrix} d_1 - \bar{d} & \dots & d_n - \bar{d} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Addition and subtraction of ψ_0 in β amounts to the normalization, $\psi_0 = 0$, which has no effect on the variance components of interest. As [Abowd et al. \(1999, 2002\)](#) note, least squares estimation of (6) requires one normalization of the ψ vector within each set of firms connected by worker mobility. For simplicity, we assume all firms are connected so that only a single normalization is required.⁵

Covariance Representation [Abowd et al. \(1999\)](#) estimated σ_ψ^2 and $\sigma_{\alpha, \psi}$ using the naive plug-in estimators $\hat{\beta}' A_\psi \hat{\beta}$ and $\hat{\beta}' A_{\alpha, \psi} \hat{\beta}$ which are, in general, biased. [Andrews et al. \(2008\)](#) proposed the “homoscedasticity-only” estimators of (4). These estimators are unbiased when the errors ε_i are independent and have common variance. [Bonhomme et al. \(2019\)](#) propose a two-step estimation approach that is consistent in the presence of heteroscedasticity when the support of firm wage effects is restricted to a finite number of values and each firm grows large with the total sample size n . Our leave-out estimators, which avoid both the homoscedasticity requirement on the errors and any cardinality restrictions on the support of the firm wage effects, can be written compactly

⁵[Bonhomme et al. \(2019\)](#) study a closely related model where workers and firms each belong to one of a finite number of types and each pairing of worker and firm type is allowed a different mean wage. These mean wage parameters are shown to be identified when each worker type moves between each firm type with positive probability, enabling estimation even when many firms are not connected.

as covariances taking the form

$$\hat{\sigma}_\psi^2 = \sum_{i=1}^n y_i x_i' S_{xx}^{-1} A_\psi \hat{\beta}_{-i}, \quad \hat{\sigma}_{\alpha,\psi} = \sum_{i=1}^n y_i x_i' S_{xx}^{-1} A_{\alpha,\psi} \hat{\beta}_{-i}.$$

Notably, these estimators are unbiased whenever the leave out estimator $\hat{\beta}_{-i}$ can be computed, regardless of the distribution of firm sizes.

Special Case: Two time periods A simpler representation of $\hat{\sigma}_\psi^2$ is available in the case where only two time periods are available and no common regressors are present ($T_g = 2$ and $x_{gt} = 0$ for all g, t). Consider this model in first differences

$$\Delta y_g = \Delta f_g' \psi + \Delta \varepsilon_g \quad (g = 1, \dots, N)$$

where $\Delta y_g = y_{g2} - y_{g1}$, $\Delta \varepsilon_g = \varepsilon_{g2} - \varepsilon_{g1}$, and $\Delta f_g = f_{i(g,2)} - f_{i(g,1)}$. The leave-out estimator of σ_ψ^2 applied to this differenced representation of the model is:

$$\hat{\sigma}_\psi^2 = \sum_{g=1}^N \Delta y_g \Delta \tilde{f}_g' \hat{\psi}_{-g} \quad \text{where} \quad \Delta \tilde{f}_g = A_{ff} S_{\Delta f \Delta f}^{-1} \Delta f_g.$$

Note that the quantities $S_{\Delta f \Delta f}$ and $\hat{\psi}_{-g}$ correspond respectively to S_{xx} and $\hat{\beta}_{-i}$ in the first differenced model.

Remark 4. The leave-out representation above reveals that $\hat{\sigma}_\psi^2$ is not only unbiased under arbitrary heteroscedasticity and design unbalance, but also under arbitrary correlation between ε_{g1} and ε_{g2} . The same can be shown to hold for $\hat{\sigma}_{\alpha,\psi}$. Furthermore, this representation highlights that $\hat{\sigma}_\psi^2$ only depends upon observations with $\Delta f_g \neq 0$ (i.e., firm “movers”).

3 Inference on Quadratic Forms of Fixed Rank

While the previous section emphasized variance components where the rank r of A was increasing with sample size, we first study the case where r is fixed. Problems of this nature often arise when testing a few linear restrictions or when conducting inference on linear combinations of the regression coefficients, say $v'\beta$. In the case of two-way fixed effects models of wage determination, the quantity $v'\beta$ might correspond to the difference in mean values of firm effects between male and female workers (Card et al., 2015) or to the coefficient from a projection of firm effects onto firm size (Bloom et al., 2018). A third use case, discussed at length by Cattaneo et al. (2018), is where $v'\beta$ corresponds to a linear combination of a few common coefficients in a linear model with high dimensional fixed effects that are regarded as nuisance parameters.

To characterize the limit distribution of $\hat{\theta}$ when r is small, we rely on a representation of θ as a weighted sum of squared linear combinations of the data: $\hat{\theta} = \sum_{\ell=1}^r \lambda_{\ell} \left(\hat{b}_{\ell}^2 - \hat{\mathbb{V}}[\hat{b}_{\ell}] \right)$ where

$$\hat{b} = \sum_{i=1}^n w_i y_i \quad \text{and} \quad \hat{\mathbb{V}}[\hat{b}] = \sum_{i=1}^n w_i w_i' \hat{\sigma}_i^2$$

for $w_i = (w_{i1}, \dots, w_{ir})' = Q' S_{xx}^{-1/2} x_i$. The following theorem characterizes the asymptotic distribution of $\hat{\theta}$ while providing conditions under which \hat{b} is asymptotically normal and $\hat{\mathbb{V}}[\hat{b}]$ is consistent.

Theorem 1. *If Assumption 1 holds, r is fixed, and $\max_i w_i' w_i = o(1)$, then*

1. $\mathbb{V}[\hat{b}]^{-1/2}(\hat{b} - b) \xrightarrow{d} \mathcal{N}(0, I_r)$ where $b = Q' S_{xx}^{1/2} \beta$,
2. $\mathbb{V}[\hat{b}]^{-1} \hat{\mathbb{V}}[\hat{b}] \xrightarrow{p} I_r$,
3. $\hat{\theta} = \sum_{\ell=1}^r \lambda_{\ell} \left(\hat{b}_{\ell}^2 - \mathbb{V}[\hat{b}_{\ell}] \right) + o_p(\mathbb{V}[\hat{\theta}]^{1/2})$,

The high-level requirement of this theorem that $\max_i w_i' w_i = o(1)$ is a Lindeberg condition ensuring that no observation is too influential. One can think of $\max_i w_i' w_i$ as measuring the inverse effective sample size available for estimating b : when the weights are equal across i , the equality $\sum_{i=1}^n w_i w_i' = I_r$ implies that $w_{i\ell}^2 = \frac{1}{n}$. Since $\frac{1}{n} \sum_{i=1}^n w_i' w_i = \frac{r}{n}$, the requirement that $\max_i w_i' w_i = o(1)$ is implied by a variety of primitive conditions that limit how far a maximum is from the average (see, e.g., [Anatolyev, 2012](#), Appendix A.1). Note that Theorem 1 does not apply to settings where r is proportional to n because $\max_i w_i' w_i \geq \frac{r}{n}$.

In the special case where $A = vv'$ for some non-random vector v , Theorem 1 establishes that the variance estimator $\hat{\mathbb{V}}[\hat{\beta}] = S_{xx}^{-1} \left(\sum_{i=1}^n x_i x_i' \hat{\sigma}_i^2 \right) S_{xx}^{-1}$ enables consistent inference on the linear combination $v' \beta$ using the approximation

$$\frac{v'(\hat{\beta} - \beta)}{\sqrt{v' \hat{\mathbb{V}}[\hat{\beta}] v}} \xrightarrow{d} \mathcal{N}(0, 1). \quad (7)$$

To derive this result we assumed that $\max_i P_{ii} \leq c$ for some $c < 1$, whereas standard Eicker-White variance estimators generally require that $\max_i P_{ii} \rightarrow 0$ and [Cattaneo et al. \(2018\)](#) establish an asymptotically valid approach to inference in settings where $\max_i P_{ii} \leq 1/2$. Thus $\hat{\mathbb{V}}[\hat{\beta}]$ leads to valid inference under weaker conditions than existing versions of Eicker-White variance estimators.

Remark 5. Theorem 1 extends classical results on hypothesis testing of a few linear restrictions, say, $H_0 : R\beta = 0$, to allow for many regressors and heteroscedasticity. A convenient choice of A for testing purposes is $\frac{1}{r} R' (R S_{xx}^{-1} R')^{-1} R$ where r , the rank of $R \in \mathbb{R}^{r \times k}$, is fixed. Under H_0 , the asymptotic distribution of $\hat{\theta}$ is an equally weighted sum of r central χ^2 random variables. This distribution is known up to $\mathbb{V}[\hat{b}]$ and a critical value can be found through simulation. For a recent

contribution to this literature, see [Anatolyev \(2012\)](#) who allows for many regressors but considers the special case of homoscedastic errors.

4 Inference on Quadratic Forms of Growing Rank

We now turn to the more challenging problem of conducting inference on θ when r increases with n , as in the examples discussed in [Section 2](#). These results also enable tests of many linear restrictions. For example, in a model of gender-specific firm effects of the sort considered by [Card et al. \(2015\)](#), testing the hypothesis that men and women face identical sets of firm fixed effects entails as many equality restrictions as there are firms.

4.1 Limit Distribution

In order to describe the result we introduce $\tilde{x}_i = \sum_{\ell=1}^n M_{i\ell} \frac{B_{\ell\ell}}{1-P_{\ell\ell}} x_{\ell}$ where $M_{i\ell} = \mathbf{1}_{\{i=\ell\}} - x_i S_{xx}^{-1} x_{\ell}$. Note that \tilde{x}_i gives the residual from a regression of $\frac{B_{ii}}{1-P_{ii}} x_i$ on x_i . Therefore, $\tilde{x}_i = 0$ when the regressor design is balanced. The contribution of \tilde{x}_i to the behavior of $\hat{\theta}$ is through the estimation of $\sum_{i=1}^n B_{ii} \sigma_i^2$, which can be ignored in the case where the rank of A is bounded. When the rank of A is large, as implied by condition (ii) of [Theorem 2](#) below, this estimation error can resurface in the asymptotic distribution. One can think of the eigenvalue ratio in (ii) as the inverse effective rank of \tilde{A} : when all the eigenvalues are equal $\frac{\lambda_1^2}{\sum_{\ell=1}^r \lambda_{\ell}^2} = \frac{1}{r}$.

Theorem 2. *Recall that $\tilde{x}_i = AS_{xx}^{-1} x_i$ where $\hat{\theta} = \sum_{i=1}^n y_i \tilde{x}_i' \hat{\beta}_{-i}$. If [Assumption 1](#) holds and the following conditions are satisfied*

$$(i) \mathbb{V}[\hat{\theta}]^{-1} \max_i \left((\tilde{x}_i' \beta)^2 + (\tilde{x}_i' \beta)^2 \right) = o(1), \quad (ii) \frac{\lambda_1^2}{\sum_{\ell=1}^r \lambda_{\ell}^2} = o(1),$$

then $\mathbb{V}[\hat{\theta}]^{-1/2} (\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, 1)$.

The proof of [Theorem 2](#) relies on a variation of Stein's method developed in [Solvsten \(2019\)](#) and a representation of $\hat{\theta}$ as a second order U-statistic, i.e.,

$$\hat{\theta} = \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell} y_i y_{\ell} \tag{8}$$

where $C_{i\ell} = B_{i\ell} - 2^{-1} M_{i\ell} \left(M_{ii}^{-1} B_{ii} + M_{\ell\ell}^{-1} B_{\ell\ell} \right)$ and $B_{i\ell} = x_i' S_{xx}^{-1} A S_{xx}^{-1} x_{\ell}$. The proof shows that the ‘kernel’ $C_{i\ell}$ varies with n in such a way that $\hat{\theta}$ is asymptotically normal whether or not $\hat{\theta}$ is a degenerate U-statistic (i.e., whether or not β is zero).

One representation of the variance appearing in Theorem 2 is

$$\mathbb{V}[\hat{\theta}] = \sum_{i=1}^n (2\tilde{x}'_i\beta - \check{x}'_i\beta)^2 \sigma_i^2 + 2 \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell}^2 \sigma_i^2 \sigma_\ell^2.$$

Note that this variance is bounded from below by $\min_i \sigma_i^2 \sum_{i=1}^n (2\tilde{x}'_i\beta)^2 + (\check{x}'_i\beta)^2$ since $\sum_{i=1}^n \tilde{x}'_i\beta\check{x}'_i\beta = 0$. Therefore (i) will be satisfied whenever $\max_i \left((\tilde{x}'_i\beta)^2 + (\check{x}'_i\beta)^2 \right)$ is not too large compared to $\sum_{i=1}^n (\tilde{x}'_i\beta)^2 + (\check{x}'_i\beta)^2$. As in Theorem 1, (i) is implied by a variety of primitive conditions that limit how far a maximum is from the average, but since (i) involves a one dimensional function of x_i it can also be satisfied when r is large. A particularly simple case where (i) is satisfied is when $\beta = 0$; further cases are discussed in Section 7.

Remark 6. Theorem 2 can be used to test a large system of linear restrictions of the form $H_0 : R\beta = 0$ where $r \rightarrow \infty$ is the rank of $R \in \mathbb{R}^{r \times k}$. Under this null hypothesis, choosing $A = \frac{1}{r} R' (RS_{xx}^{-1} R')^{-1} R$ implies $\mathbb{V}[\hat{\theta}]^{-1/2} \hat{\theta} \xrightarrow{d} \mathcal{N}(0, 1)$ since all the non-zero eigenvalues of \tilde{A} are equal to $\frac{1}{r}$. The existing literature allows for either heteroscedastic errors and moderately few regressors (Donald et al., 2003, $k^3/n \rightarrow 0$) or homoscedastic errors and many regressors (Anatolyev, 2012, $k/n \leq c < 1$). When coupled with the estimator of $\mathbb{V}[\hat{\theta}]$ presented in the next subsection, this result enables tests with heteroscedastic errors and many regressors.

Remark 7. Theorem 2 extends some common results in the literature on many and many weak instruments (see, e.g., Chao et al., 2012) where the estimators are asymptotically equivalent to quadratic forms. The structure of that setting is such that $\tilde{A} = I_r/r$ and $r \rightarrow \infty$, in which case condition (ii) of Theorem 2 is automatically satisfied.

4.2 Variance Estimation

In order to conduct inference based on the normal approximation in Theorem 2 we now propose an estimator of $\mathbb{V}[\hat{\theta}]$. The U-statistic representation of $\hat{\theta}$ in (8) implies that the variance of $\hat{\theta}$ is

$$\mathbb{V}[\hat{\theta}] = 4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} x'_\ell \beta \right)^2 \sigma_i^2 + 2 \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell}^2 \sigma_i^2 \sigma_\ell^2.$$

Naively replacing $\{x'_i\beta, \sigma_i^2\}_{i=1}^n$ with $\{y_i, \hat{\sigma}_i^2\}_{i=1}^n$ in the above formula to form a plug-in estimator of $\mathbb{V}[\hat{\theta}]$ will, in general, lead to invalid inferences as $\hat{\sigma}_i^2 \hat{\sigma}_\ell^2$ is a biased estimator of $\sigma_i^2 \sigma_\ell^2$. For this reason, we consider estimators of the error variances that rely on leaving out more than one observation. Since this approach places additional restrictions on the design, Appendix C.5.1 describes a simple adjustment which leads to conservative inference in settings where these restrictions do not hold.

Sample Splitting Our specific proposal is an estimator that exploits two independent unbiased

estimators of $x'_i\beta$ that are also independent of y_i . We denote these estimators $\widehat{x'_i\beta}_{-i,s} = \sum_{\ell \neq i}^n P_{i\ell,s} y_\ell$ for $s = 1, 2$, where $P_{i\ell,s}$ does not (functionally) depend on the $\{y_i\}_{i=1}^n$. To ensure independence between $\widehat{x'_i\beta}_{-i,1}$ and $\widehat{x'_i\beta}_{-i,2}$, we require that $P_{i\ell,1}P_{i\ell,2} = 0$ for all ℓ . Employing these split sample estimators, we create a new set of unbiased estimators for σ_i^2 :

$$\tilde{\sigma}_i^2 = \left(y_i - \widehat{x'_i\beta}_{-i,1} \right) \left(y_i - \widehat{x'_i\beta}_{-i,2} \right) \quad \text{and} \quad \hat{\sigma}_{i,-\ell}^2 = \begin{cases} y_i(y_i - \widehat{x'_i\beta}_{-i,1}), & \text{if } P_{i\ell,1} = 0, \\ y_i(y_i - \widehat{x'_i\beta}_{-i,2}), & \text{if } P_{i\ell,1} \neq 0, \end{cases}$$

where $\hat{\sigma}_{i,-\ell}^2$ is independent of y_ℓ and $\tilde{\sigma}_i^2$ is a cross-fit estimator of the form considered in [Newey and Robins \(2018\)](#). These cross-fit estimators can be used to construct an estimator of $\sigma_i^2\sigma_\ell^2$ that, under certain design conditions, will be unbiased. Letting $P_{im,-\ell} = P_{im,1}1_{\{P_{i\ell,1}=0\}} + P_{im,2}1_{\{P_{i\ell,1}\neq 0\}}$ denote the weight observation m receives in $\hat{\sigma}_{i,-\ell}^2$ and $\tilde{C}_{i\ell} = C_{i\ell}^2 + 2\sum_{m=1}^n C_{mi}C_{m\ell}(P_{mi,1}P_{m\ell,2} + P_{mi,2}P_{m\ell,1})$, we define

$$\widehat{\sigma_i^2\sigma_\ell^2} = \begin{cases} \hat{\sigma}_{i,-\ell}^2 \cdot \hat{\sigma}_{\ell,-i}^2, & \text{if } P_{im,-\ell}P_{\ell m,-i} = 0 \text{ for all } m, \\ \tilde{\sigma}_i^2 \cdot \hat{\sigma}_{\ell,-i}^2, & \text{else if } P_{i\ell,1} + P_{i\ell,2} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot \tilde{\sigma}_\ell^2, & \text{else if } P_{\ell i,1} + P_{\ell i,2} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot (y_\ell - \bar{y})^2 \cdot 1_{\{\tilde{C}_{i\ell} < 0\}}, & \text{otherwise.} \end{cases}$$

The first three cases in the above definition correspond respectively to pairs where (i) $\hat{\sigma}_{i,-\ell}^2$ and $\hat{\sigma}_{\ell,-i}^2$ are independent, (ii) $\widehat{x'_i\beta}_{-i,1}$ and $\widehat{x'_i\beta}_{-i,2}$ are independent of y_ℓ , and (iii) $\widehat{x'_\ell\beta}_{-\ell,1}$ and $\widehat{x'_\ell\beta}_{-\ell,2}$ are independent of y_i . When any of these three cases apply, we obtain an unbiased estimator of $\sigma_i^2\sigma_\ell^2$. For the remaining set of pairs $\mathcal{B} = \{(i, \ell) : P_{im,-\ell}P_{\ell m,-i} \neq 0 \text{ for some } m, P_{i\ell,1} + P_{i\ell,2} \neq 0, P_{\ell i,1} + P_{\ell i,2} \neq 0\}$ that comprise the fourth case we rely on an unconditional variance estimator which leads to a biased estimator of $\sigma_i^2\sigma_\ell^2$ and conservative inference.

Design Requirements Constructing the above split sample estimators places additional requirements on the design matrix S_{xx} . We briefly discuss these requirements in the context of [Examples 2, 3, and 4](#). In the ANOVA setup of [Example 2](#), leave-one-out estimation requires a minimum group size of two, whereas existence of $\{\widehat{x'_i\beta}_{-i,s}\}_{s=1,2}$ requires groups sizes of at least three. Conservative inference can be avoided (i.e., the set \mathcal{B} will be empty) when the minimum group size is at least four. In the random coefficients model of [Example 3](#), minimum group sizes of three and five are sufficient to ensure feasibility of leave-one-out estimation and existence of $\{\widehat{x'_i\beta}_{-i,s}\}_{s=1,2}$, respectively. Conservativeness can be avoided with a minimum group size of seven.

In the first differenced two-way fixed effects model of [Example 4](#), the predictions $\{\widehat{x'_i\beta}_{-i,s}\}_{s=1,2}$ are associated with particular paths in the worker-firm mobility network and independence requires that these paths be edge-disjoint. Menger's theorem ([Menger, 1927](#)) implies that $\{\widehat{x'_i\beta}_{-i,s}\}_{s=1,2}$

exists if the design matrix has full rank when any two observations are dropped. Menger's theorem also implies that conservativeness can be avoided if the design matrix has full rank when any three observations are dropped. In our application, we use Dijkstra's algorithm to find the paths that generate $\{\widehat{x'_i\beta_{-i,s}}\}_{s=1,2}$ (see Appendix B.4 for further details).

Consistency The following lemma shows that $\widehat{\sigma_i^2\sigma_\ell^2}$ can be utilized to construct an estimator of $\mathbb{V}[\hat{\theta}]$ that delivers consistent inference when sufficiently few pairs fall into \mathcal{B} and provides conservative inference otherwise.

Lemma 5. For $s = 1, 2$, suppose that $\widehat{x'_i\beta_{-i,s}}$ satisfies (unbiasedness) $\sum_{\ell \neq i}^n P_{i\ell,s} x'_\ell \beta = x'_i \beta$, (sample splitting) $P_{i\ell,1} P_{i\ell,2} = 0$ for all ℓ , and (projection property) $\lambda_{\max}(P_s P'_s) = O(1)$ where $P_s = (P_{i\ell,s})_{i,\ell}$ is the hat-matrix corresponding to $\widehat{x'_i\beta_{-i,s}}$. Let

$$\widehat{\mathbb{V}}[\hat{\theta}] = 4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} y_\ell \right)^2 \tilde{\sigma}_i^2 - 2 \sum_{i=1}^n \sum_{\ell \neq i} \tilde{C}_{i\ell} \widehat{\sigma_i^2 \sigma_\ell^2}.$$

1. If the conditions of Theorem 2 hold and $|\mathcal{B}| = O(1)$, then $\frac{\hat{\theta} - \theta}{\widehat{\mathbb{V}}[\hat{\theta}]^{1/2}} \xrightarrow{d} \mathcal{N}(0, 1)$.
2. If the conditions of Theorem 2 hold, then $\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \left[\hat{\theta} \pm z_\alpha \widehat{\mathbb{V}}[\hat{\theta}]^{1/2} \right] \right) \geq 1 - \alpha$ where z_α^2 denotes the $(1 - \alpha)$ 'th quantile of a central χ_1^2 random variable.

In the formula for $\widehat{\mathbb{V}}[\hat{\theta}]$, the first term can be seen as a plug-in estimator and standard results for quartic forms imply that the expectation of this term is $\mathbb{V}[\hat{\theta}] + 2 \sum_{i=1}^n \sum_{\ell \neq i} \tilde{C}_{i\ell} \sigma_i^2 \sigma_\ell^2$. Hence, the second term is a bias correction which completely removes the bias when $\mathcal{B} = \emptyset$ and leaves a positive bias otherwise. In Appendix C.5.1 we establish validity of an adjustment to $\widehat{\mathbb{V}}[\hat{\theta}]$ that utilizes an upward biased unconditional variance estimator for observations where it is not possible to construct $\{\widehat{x'_i\beta_{-i,s}}\}_{s=1,2}$.

Remark 8. The purpose of the condition $|\mathcal{B}| = O(1)$ in the above lemma is to ensure that the bias of $\widehat{\mathbb{V}}[\hat{\theta}]$ grows small with the sample size. Because the bias of $\widehat{\mathbb{V}}[\hat{\theta}]$ is non-negative, inference based on $\widehat{\mathbb{V}}[\hat{\theta}]$ remains valid even when this condition fails, as stated in the second part of Lemma 5. In practice, it may be useful for researchers to calculate the fraction of pairs that belong to \mathcal{B} to gauge the extent to which inference might be conservative. Similarly, it may be useful to compute the share of observations where it is not possible to construct $\{\widehat{x'_i\beta_{-i,s}}\}_{s=1,2}$ to investigate whether upward bias in the standard error could lead to power concerns.

5 Weakly Identified Quadratic Forms of Growing Rank

In some settings where r grows with the sample size, condition (ii) of Theorem 2 may not apply. For example in two-way fixed effects models, it is possible that “bottlenecks” arise in the mobility network that lead the largest eigenvalues to dominate the others.

This section provides a theorem which covers the case where some of the squared eigenvalues $\lambda_1^2, \dots, \lambda_r^2$ are large relative to their sum $\sum_{\ell=1}^r \lambda_\ell^2$. To motivate this assumption, note that each eigenvalue of \tilde{A} measures how strongly θ depends on a particular linear combination of the elements of β relative to the difficulty of estimating that combination (as summarized by S_{xx}^{-1}). From Lemma 3, $\text{trace}(\tilde{A}^2) = \sum_{\ell=1}^r \lambda_\ell^2$ governs the total variability in $\hat{\theta}$. Therefore, Theorem 3 covers the case where θ depends strongly on a few linear combinations of β that are imprecisely estimated relative to the overall sampling uncertainty in $\hat{\theta}$. The following assumption formalizes this setting.

Assumption 2. *There exist a $c > 0$ and a known and fixed $q \in \{1, \dots, r - 1\}$ such that*

$$\frac{\lambda_{q+1}^2}{\sum_{\ell=1}^r \lambda_\ell^2} = o(1) \quad \text{and} \quad \frac{\lambda_q^2}{\sum_{\ell=1}^r \lambda_\ell^2} \geq c \quad \text{for all } n.$$

Assumption 2 defines q as the number of squared eigenvalues that are large relative to their sum. Equivalently, q indexes the number of nuisance parameters in b that are *weakly identified* relative to their influence on θ and the uncertainty in $\hat{\theta}$. The assumption that q is known is motivated by our discussion of Examples 1–4 in Section 7 and the theoretical literature on weak identification, which typically makes an ex-ante distinction between strongly and weakly identified parameters (e.g., Andrews and Cheng, 2012). In Section 6.2 we offer some guidance on choosing q in settings where it is unknown.

5.1 Limit Distribution

Given knowledge of q , we can split $\hat{\theta}$ into a known function of $\hat{\mathbf{b}}_q = (\hat{b}_1, \dots, \hat{b}_q)'$ and $\hat{\theta}_q$ where $\hat{b}_1, \dots, \hat{b}_q$ are OLS estimators of the weakly identified nuisance parameters:

$$\begin{aligned} \hat{\mathbf{b}}_q &= \sum_{i=1}^n \mathbf{w}_{iq} y_i, & \mathbf{w}_{iq} &= (w_{i1}, \dots, w_{iq})', \\ \hat{\theta}_q &= \hat{\theta} - \sum_{\ell=1}^q \lambda_\ell (\hat{b}_\ell^2 - \hat{\mathbb{V}}[\hat{b}_\ell]), & \hat{\mathbb{V}}[\hat{b}] &= \sum_{i=1}^n w_i w_i' \hat{\sigma}_i^2. \end{aligned}$$

The main difficulty in proving the following Theorem is to show that the joint distribution of $(\hat{\mathbf{b}}_q, \hat{\theta}_q)'$ is normal, which we do using the same variation of Stein’s method that was employed for Theorem 2. The high-level conditions involve \tilde{x}_{iq} and \check{x}_{iq} which are the parts of \tilde{x}_i and \check{x}_i that

pertain to $\hat{\theta}_q$ and are defined in the proof of Theorem 3.

Theorem 3. *If $\max_i w'_{iq} w_{iq} = o(1)$, $\mathbb{V}[\hat{\theta}_q]^{-1} \max_i \left((\hat{x}'_{iq} \beta)^2 + (\hat{x}'_{iq} \beta)^2 \right) = o(1)$, and Assumptions 1 and 2 hold, then*

1. $\mathbb{V}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)']^{-1/2} \left((\hat{\mathbf{b}}'_q, \hat{\theta}_q)' - \mathbb{E}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] \right) \xrightarrow{d} \mathcal{N}(0, I_{q+1})$
2. $\hat{\theta} = \sum_{\ell=1}^q \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right) + \hat{\theta}_q + o_p(\mathbb{V}[\hat{\theta}]^{1/2})$

Theorem 3 provides an approximation to $\hat{\theta}$ in terms of a quadratic function of q asymptotically normal random variables and a linear function of one asymptotically normal random variable. Here, the non-centralities $\mathbb{E}[\hat{\mathbf{b}}_q] = (b_1, \dots, b_q)'$ serve as nuisance parameters that influence both θ and the shape of the limiting distribution of $\hat{\theta} - \theta$. The next section proposes an approach to dealing with these nuisance parameters that provides asymptotically valid inference on θ for any value of q .

5.2 Variance Estimation

In Theorem 3 the relevant variance is $\Sigma_q := \mathbb{V}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)']$,

$$\Sigma_q = \sum_{i=1}^n \begin{bmatrix} w_{iq} w'_{iq} \sigma_i^2 & 2w_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \right) \sigma_i^2 \\ 2w'_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \right) \sigma_i^2 & 4 \left(\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \right)^2 \sigma_i^2 + 2 \sum_{\ell \neq i} C_{i\ell q}^2 \sigma_i^2 \sigma_\ell^2 \end{bmatrix},$$

where $C_{i\ell q}$ is defined in the proof of Theorem 3. Our estimator of this variance reuses the split sample estimators introduced for Theorem 2:

$$\hat{\Sigma}_q = \sum_{i=1}^n \begin{bmatrix} w_{iq} w'_{iq} \hat{\sigma}_i^2 & 2w_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right) \tilde{\sigma}_i^2 \\ 2w'_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right) \tilde{\sigma}_i^2 & 4 \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right)^2 \tilde{\sigma}_i^2 - 2 \sum_{\ell \neq i} \tilde{C}_{i\ell q}^2 \tilde{\sigma}_i^2 \tilde{\sigma}_\ell^2 \end{bmatrix}$$

where $\tilde{C}_{i\ell q}$ and $\tilde{\sigma}_i^2 \tilde{\sigma}_\ell^2$ are defined in the proof of the next lemma which shows consistency of this variance estimator.

Lemma 6. *For $s = 1, 2$, suppose that $\widehat{x}'_{i\beta_{-i,s}}$ satisfies $\sum_{\ell \neq i} P_{i\ell, s} x'_{\ell} \beta = x'_i \beta$, $P_{i\ell, 1} P_{i\ell, 2} = 0$ for all ℓ , and $\lambda_{\max}(P_s P'_s) = O(1)$. If the conditions of Theorem 3 hold and $|\mathcal{B}| = O(1)$, then $\Sigma_q^{-1} \hat{\Sigma}_q \xrightarrow{p} I_{q+1}$.*

Remark 9. As in the case of variance estimation for Theorem 2, it may be that the design does not allow for construction of the predictions $\widehat{x}'_{i\beta_{-i,1}}$ and $\widehat{x}'_{i\beta_{-i,2}}$ used in $\hat{\Sigma}_q$. For such cases, Appendix C.5.1 proposes an adjustment to $\hat{\Sigma}_q$ which has a positive definite bias and therefore leads to valid (but conservative) inference when coupled with the inference method discussed in the next section.

6 Inference with Nuisance Parameters

In this section, we develop a two-sided confidence interval for θ that delivers asymptotic size control conditional on a choice of q . Our proposal involves inverting a minimum distance statistic in $\hat{\mathbf{b}}_q$ and $\hat{\theta}_q$, which Theorem 3 implies are jointly normally distributed. To avoid the conservatism associated with standard projection methods (e.g., Dufour and Jasiak, 2001), we seek to adjust the critical value downwards to deliver size control on θ rather than $\mathbb{E}[(\hat{\mathbf{b}}_q', \hat{\theta}_q)']$. However, unlike in standard projection problems (e.g., the problem of subvector inference), θ is a nonlinear function of $\mathbb{E}[\hat{\mathbf{b}}_q]$. To accommodate this complication, we use a critical value proposed by Andrews and Mikusheva (2016) that depends on the curvature of the problem.

6.1 Inference With Known q

The confidence interval we consider is based on inversion of a minimum-distance statistic for $(\hat{\mathbf{b}}_q', \hat{\theta}_q)'$ using the critical value proposed in Andrews and Mikusheva (2016). For a specified level of confidence, $1 - \alpha$, we consider the interval

$$\hat{C}_{\alpha,q}^{\theta} = \left[\min_{(\hat{b}_1, \dots, \hat{b}_q, \hat{\theta}_q)' \in \hat{\mathbb{E}}_{\alpha,q}} \sum_{\ell=1}^q \lambda_{\ell} \hat{b}_{\ell}^2 + \hat{\theta}_q, \max_{(\hat{b}_1, \dots, \hat{b}_q, \hat{\theta}_q)' \in \hat{\mathbb{E}}_{\alpha,q}} \sum_{\ell=1}^q \lambda_{\ell} \hat{b}_{\ell}^2 + \hat{\theta}_q \right]$$

where

$$\hat{\mathbb{E}}_{\alpha,q} = \left\{ (\mathbf{b}'_q, \theta_q)' \in \mathbb{R}^{q+1} : \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}' \hat{\Sigma}_q^{-1} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix} \leq z_{\alpha, \hat{\kappa}_q}^2 \right\}.$$

The critical value function, $z_{\alpha, \kappa}$, depends on the maximal curvature, κ , of a certain manifold (exact definitions of $z_{\alpha, \kappa}$ and κ are given in Appendix C.6). Heuristically, κ can be thought of as summarizing the influence of the nuisance parameter $\mathbb{E}[\hat{\mathbf{b}}_q]$ on the shape of $\hat{\theta}$'s limiting distribution. Accordingly, $z_{\alpha}^2 := z_{\alpha, 0}^2$ is equal to the $(1 - \alpha)$ 'th quantile of a central χ_1^2 random variable. As $\kappa \rightarrow \infty$, $z_{\alpha, \kappa}^2$ approaches the $(1 - \alpha)$ 'th quantile of a central χ_{q+1}^2 random variable. This upper limit on $z_{\alpha, \kappa}$ is used in the projection method in its classical form as popularized in econometrics by Dufour and Jasiak (2001), while the lower limit z_{α} would yield size control if θ were linear in $\mathbb{E}[\hat{\mathbf{b}}_q]$.

When $q = 0$, the maximal curvature is zero and \hat{C}_0^{θ} simplifies to $[\hat{\theta} \pm z_{\alpha} \hat{\mathbb{V}}[\hat{\theta}]^{1/2}]$. When $q = 1$, the maximal curvature is $\hat{\kappa}_1 = \frac{2|\lambda_1| \hat{\mathbb{V}}[\hat{b}_1]}{\hat{\mathbb{V}}[\hat{\theta}_1]^{1/2} (1 - \hat{\rho}^2)^{1/2}}$ where $\hat{\rho}$ is the estimated correlation between \hat{b}_1 and $\hat{\theta}_1$. This curvature measure is intimately related to eigenvalue ratios previously introduced, as $\hat{\kappa}_1^2$ is approximately equal to $\frac{2\lambda_1^2}{\sum_{\ell=2}^r \lambda_{\ell}^2}$ when the error terms are homoscedastic and $\beta = 0$. A closed form expression for the $q = 1$ confidence interval is provided in Appendix C.6. When $q > 1$, inference relies on solving two quadratic optimization problems that involve $q + 1$ unknowns, which can be achieved reliably using standard quadratic programming routines.

The following lemma shows that a consistent variance estimator as proposed in Lemma 6 suffices for asymptotic validity under the conditions of Theorem 3 and Appendix C.5.1 establishes validity when only a conservative variance estimator is available.

Lemma 7. *If $\Sigma_q^{-1} \hat{\Sigma}_q \xrightarrow{p} I_{q+1}$ and the conditions of Theorem 3 hold, then*

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \hat{C}_{\alpha, q}^\theta \right) \geq 1 - \alpha.$$

The confidence interval studied in Lemma 7 constructs a $q + 1$ dimensional ellipsoid $\hat{E}_{\alpha, q}$ and maps it through the quadratic function $(\dot{b}_1, \dots, \dot{b}_q, \dot{\theta}_q) \mapsto \sum_{\ell=1}^q \lambda_\ell \dot{b}_\ell^2 + \dot{\theta}_q$. This approach ensures uniform coverage over any possible values of the nuisance parameters b_1, \dots, b_q which are imprecisely estimated relative to overall sampling uncertainty in $\hat{\theta}$.

Remark 10. An alternative to Lemma 7 is to conduct inference using a first-order Taylor expansion of $\sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \hat{\theta}_q$. This so-called ‘‘Delta method’’ approach is asymptotically equivalent to using the confidence interval $[\hat{\theta} \pm z_\alpha \hat{V}[\hat{\theta}]^{1/2}]$ studied in Section 4. However, the Delta method is not uniformly valid in the presence of nuisance parameters as approximate linearity can fail when $\min_{\ell \leq q} \lambda_\ell \hat{b}_\ell^2 = O(1)$. Section 7 introduces a stochastic block model with $q = 1$ and characterizes b_1^2 as the squared difference in average firm effects across two blocks multiplied by the number of between block movers. Thus the Delta method will potentially undercover unless there are strong systematic differences between the two blocks.

6.2 Choosing q

The preceding discussion of inference considered a setting where the number of weakly identified parameters was known in advance. In some applications, it may not be clear ex ante what value q takes. In such situations researchers may wish to report confidence intervals for two consecutive values of q (or their union). This heuristic serves to minimize the influence of the specific value of q picked, and both our simulations and empirical application suggest that $\hat{C}_{\alpha, q}^\theta$ barely varies with q when $\frac{\lambda_{q+1}^2}{\sum_{\ell=1}^r \lambda_\ell^2} < \frac{1}{10}$. Consequently, little power is sacrificed by taking the union.

This observation also suggests a heuristic threshold for choosing q ; namely, to let q be such that $\frac{\lambda_q^2}{\sum_{\ell=1}^r \lambda_\ell^2} \geq \frac{1}{10}$ and $\frac{\lambda_{q+1}^2}{\sum_{\ell=1}^r \lambda_\ell^2} < \frac{1}{10}$, with $q = 0$ when $\frac{\lambda_1^2}{\sum_{\ell=1}^r \lambda_\ell^2} < \frac{1}{10}$. A similar threshold rule can be motivated under a slight strengthening of Assumption 2 which allows one to learn q from the data.

Assumption 2’. *There exist a $c > 0$, an $\epsilon > 0$, and a fixed $q \in \{1, \dots, r - 1\}$ such that*

$$\frac{\lambda_{q+1}^2}{\sum_{\ell=1}^r \lambda_\ell^2} = O(r^{-\epsilon}) \quad \text{and} \quad \frac{\lambda_q^2}{\sum_{\ell=1}^r \lambda_\ell^2} \geq c \quad \text{for all } n.$$

A threshold based choice of q is the unique \hat{q} for which

$$\frac{\lambda_{\hat{q}+1}^2}{\sum_{\ell=1}^r \lambda_{\ell}^2} < c_r \quad \text{and} \quad \frac{\lambda_{\hat{q}}^2}{\sum_{\ell=1}^r \lambda_{\ell}^2} \geq c_r \quad \text{for some } c_r \rightarrow 0,$$

with $\hat{q} = 0$ when $\frac{\lambda_1^2}{\sum_{\ell=1}^r \lambda_{\ell}^2} < c_r$. Under Assumption 2', $\hat{q} = q$ in sufficiently large samples provided that c_r is chosen so that $c_r r^{\varepsilon} \rightarrow \infty$. This condition is satisfied when c_r shrinks slowly to zero, e.g., when $c_r \propto 1/\log(r)$.

7 Verifying Conditions

We now revisit the examples of Section 2 and verify the conditions required to apply our theoretical results. Appendix C.7 provides further details on these calculations.

Example 1. (Coefficient of determination, continued) Recall that $\theta = \sigma_{X\beta}^2 = \beta' A \beta$ where $A = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})'$ and $\tilde{A} = \frac{1}{n}(I_k - nS_{xx}^{-1/2} \bar{x} \bar{x}' S_{xx}^{-1/2})$. Suppose Assumption 1 holds.

Consistency Consistency follows from Lemma 3 since $\lambda_{\ell} = \frac{1}{n}$ for $\ell = 1, \dots, r$ where $r = \dim(x_i) - 1$. Thus $\text{trace}(\tilde{A}^2) = r/n^2 \leq 1/n = o(1)$.

Limit Distribution If $\dim(x_i)$ is fixed, then $w_i' w_i = P_{ii} - \frac{1}{n}$ and Theorem 1 applies under the standard “textbook” condition that $\max_i P_{ii} = o(1)$. If $\dim(x_i) \rightarrow \infty$, then Theorem 2 applies if $\mathbb{V}[\hat{\theta}]^{-1} \max_i (\tilde{x}_i' \beta)^2 = o(1)$ which follows if, e.g., $\max_i \frac{1}{\sqrt{r}} \sum_{\ell=1}^n |M_{i\ell}| = o(1)$ where $M_{i\ell} = \mathbf{1}_{\{i=\ell\}} - x_i' S_{xx}^{-1} x_{\ell}$ (this condition holds in the next two examples). Equality among all eigenvalues excludes the weak identification setting of Theorem 3.

Unbounded Mean Function Inspection of the proofs reveal that Assumption 1(iii), $\max_i (x_i' \beta)^2 = O(1)$, can be dropped if the above conditions are strengthened to $\max_{i,\ell} P_{ii} (x_{\ell}' \beta)^2 = o(1)$ when $\dim(x_i)$ is fixed or $\max_{i,j} \frac{|x_j' \beta| (1 + \sum_{\ell=1}^n |M_{i\ell}|)}{\sqrt{r}} = o(1)$ when $\dim(x_i) \rightarrow \infty$.

Example 2. (Analysis of covariance, continued) Recall that $\theta = \sigma_{\alpha}^2 = \frac{1}{n} \sum_{g=1}^N T_g (\alpha_g - \bar{\alpha})^2$ where $y_{gt} = \alpha_g + x_{gt}' \delta + \varepsilon_{gt}$, g index the N groups, and T_g is group size.

No Common Regressors This is a special case of the previous example with $r = N - 1$, $P_{ii} = T_{g(i)}^{-1}$ and $\tilde{x}_i = 0$. Assumption 1(ii),(iii) requires $T_g \geq 2$ and $\max_g \alpha_g^2 = O(1)$. Theorem 1 applies if the number of groups is fixed and $\min_g T_g \rightarrow \infty$, while Theorem 2 applies if the number of groups is large. Theorem 3 cannot apply as all eigenvalues are equal to $\frac{1}{n}$.

Common Regressors To accommodate common regressors of fixed dimension, assume $\|\delta\|^2 + \max_{g,t} \|x_{gt}\|^2 = O(1)$ and that $\frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} (x_{gt} - \bar{x}_g)(x_{gt} - \bar{x}_g)'$ converges to a positive definite limit. This is a standard assumption in basic panel data models (see, e.g., Wooldridge, 2010, Chapter 10). Allowing such common regressors does not alter the previous conclusions: Theorem 1 applies if N is fixed and $\min_g T_g \rightarrow \infty$ since $w_i' w_i \leq P_{ii} = T_{g(i)}^{-1} + O(n^{-1})$, Theorem 2 applies if

$N \rightarrow \infty$ since $\sum_{\ell=1}^n |M_{i\ell}| = O(1)$, and Theorem 3 cannot apply since $n\lambda_\ell \in [c_1, c_2]$ for $\ell = 1, \dots, r$ and some $c_2 \geq c_1 > 0$ not depending on n .

Unbounded Mean Function All conclusions continue to hold if $\max_{g,t} \alpha_g^2 + \|x_{gt}\|^2 = O(1)$ is replaced with $\frac{\max_{g,t} \alpha_g^2 + \|x_{gt}\|^2}{\max\{N, \min_g T_g\}} = o(1)$ and $\sigma_\alpha^2 + \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} \|x_{gt}\|^2 = O(1)$.

Example 3. (Random coefficients, continued) For simplicity, consider the *uncentered* second moment $\theta = \frac{1}{n} \sum_{g=1}^N T_g \gamma_g^2$ where $y_{gt} = \alpha_g + z'_{gt} \gamma_g + \varepsilon_{gt}$. Suppose Assumption 1 holds and assume that $\max_{g,t} \alpha_g + \gamma_g^2 + z_{gt}^2 = O(1)$ and $\min_g S_{zz,g} \geq c > 0$ where $S_{zz,g} = \sum_{t=1}^{T_g} (z_{gt} - \bar{z}_g)^2$. Note that $\min_g S_{zz,g} > 0$ is equivalent to full rank of S_{xx} and $S_{zz,g}$ indexes how precisely γ_g can be estimated.

Consistency The N eigenvalues of \tilde{A} are $\lambda_g = \frac{T_g}{n} S_{zz,g}^{-1}$ for $g = 1, \dots, N$ where the group indexes are ordered so that $\lambda_1 \geq \dots \geq \lambda_N$. Consistency follows from Lemma 3 if $\lambda_1^{-1} = n \frac{S_{zz,1}}{T_1} \rightarrow \infty$. This is automatically satisfied with many groups of bounded size.

Limit Distribution If N is fixed and $\min_g S_{zz,g} \rightarrow \infty$, then Theorem 1 applies. If $\frac{\sqrt{N}}{T_1} S_{zz,1} \rightarrow \infty$, then Theorem 2 applies. If $\frac{\sqrt{N}}{T_2} S_{zz,2} \rightarrow \infty$, $\frac{\sqrt{N}}{T_1} S_{zz,1} = O(1)$, and $S_{zz,1} \rightarrow \infty$, then Theorem 3 applies with $q = 1$. In this case, γ_1 is weakly identified relative to its influence on θ and the overall variability of $\hat{\theta}$. This is expressed through the condition $\frac{\sqrt{N}}{T_1} S_{zz,1} = O(1)$ where $S_{zz,1}$ is the identification strength of γ_1 , T_1 provides the influence of γ_1 on θ and $1/\sqrt{N}$ indexes the variability of $\hat{\theta}$.

Example 4. (Two-way fixed effects, continued) In this final example, we restrict attention to the first-differenced setting $\Delta y_g = \Delta f'_g \psi + \Delta \varepsilon_g$ with $T_g = 2$ and a large number of firms, $J \rightarrow \infty$. Our target parameter is the variance of firm effects $\theta = \sigma_\psi^2 = \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^{T_g} (\psi_{j(g,t)} - \bar{\psi})^2$ and we consider Assumption 1 satisfied; in particular, $\max_j |\psi_j| = O(1)$.

Leverages The leverage P_{gg} of observation g is less than one if the origin and destination firms of worker g are connected by a path not involving g . Letting n_g denote the number of edges in the shortest such path, one can show that $P_{gg} \leq \frac{n_g}{1+n_g}$. Therefore, if $\max_g n_g < 100$ then Assumption 1(ii) is satisfied with $\max P_{gg} \leq .99$. In our application we find $\max_g n_g = 12$, leading to a somewhat smaller bound on the maximal leverage. The same consideration implies a bound on the model in levels since $P_{i(g,t)i(g,t)} = \frac{1}{2}(1 + P_{gg})$.

Eigenvalues The eigenvalues of \tilde{A} satisfy the equality

$$\lambda_\ell = \frac{1}{n \dot{\lambda}_{J+1-\ell}} \quad \text{for } \ell = 1, \dots, J$$

where $\dot{\lambda}_1 \geq \dots \geq \dot{\lambda}_J$ are the non-zero eigenvalues of the matrix $E^{1/2} \mathcal{L} E^{1/2}$. \mathcal{L} is the normalized Laplacian of the employer mobility network and connectedness of the network is equivalent to full rank of S_{xx} (see Appendix C.7 for definitions). E is a diagonal matrix of employer specific ‘‘churn rates’’, i.e., the number of moves in and out of a firm divided by the total number of employees

in the firm. E and \mathcal{L} interact in determining the eigenvalues of \tilde{A} . In Example 3, the quantities $\{T_\ell^{-1}S_{zz,\ell}\}_{\ell=1}^N$ played a role directly analogous to the churn rates in E , so in this example we focus on the role of \mathcal{L} by assuming that the diagonal entries of E are all equal to one.

Strongly Connected Network The employer mobility network is *strongly connected* if $\sqrt{J}\mathcal{C} \rightarrow \infty$ where $\mathcal{C} \in (0, 1]$ is Cheeger’s constant for the mobility network (see, e.g., Mohar, 1989; Jochmans and Weidner, 2016). Intuitively, \mathcal{C} measures the most severe “bottleneck” in the network, where a bottleneck is a set of movers that upon removal from the data splits the mobility network into two disjoint blocks. The severity of the bottleneck is governed by the number of movers removed divided by the smallest number of movers in either of the two disjoint blocks. The inequalities $\lambda_J \geq 1 - \sqrt{1 - \mathcal{C}^2}$ (Chung, 1997, Theorem 2.3) and $\lambda_1^2 / \sum_{\ell=1}^J \lambda_\ell^2 \leq 4(\sqrt{J}\lambda_J)^{-2}$ imply that a strongly connected network yields $q = 0$, which rules out application of Theorem 3. Furthermore, a strongly connected network is sufficient (but not necessary) for consistency of $\hat{\theta}$ as $\sum_{\ell=1}^J \lambda_\ell^2 \leq \frac{J}{n}(\sqrt{n}\lambda_J)^{-2}$.

Weakly Connected Network When $\sqrt{J}\mathcal{C}$ is bounded, the network is *weakly connected* and can contain a sufficiently severe bottleneck that a linear combination of the elements of ψ is estimated imprecisely relative to its influence on θ and the total uncertainty in $\hat{\theta}$. The weakly identified linear combination in this case is a difference in average firm effects across the two blocks on either side of the bottleneck, which contributes a χ^2 term to the asymptotic distribution. Below we use a stochastic block model to further illustrate this phenomenon. Our empirical application demonstrates that weakly connected networks can appear in practically relevant settings.

Stochastic Block Model Consider a stochastic block model of network formation where firms belong to one of two blocks and a set of workers switch firms, possibly by moving between blocks. Workers’ mobility decisions are independent: with probability p_b a worker moves between blocks and with probability $1 - p_b$ she moves within block. For simplicity, we further assume that the two blocks contain equally many firms and consider a semi-sparse network where $\frac{J \log(J)}{n} + \frac{\log(J)}{np_b} \rightarrow 0$.⁶ In this model the asymptotic behavior of $\hat{\theta}$ is governed by p_b : the most severe bottleneck is between the two blocks and has a Cheeger’s constant proportional to p_b . In Appendix C.7, we use this model to verify the high-level conditions leading to Theorems 2 and 3 and show that Theorem 2 applies when $\sqrt{J}p_b \rightarrow \infty$, while Theorem 3 applies with $q = 1$ otherwise. The argument extends to any finite number of blocks, in which case q is the number of blocks minus one. Finally, we show that $\hat{\theta}$ is consistent even when the network is weakly connected. To establish consistency we only impose $\frac{\log(J)}{np_b} \rightarrow 0$, which requires that the number of movers across the two blocks is large.

⁶The semi-sparse stochastic block model is routinely employed in the statistical literature on spectral clustering, see, e.g, Sarkar and Bickel (2015).

8 Application

Consider again the problem of estimating variance components in two-way fixed effect models of wage determination. [Card et al. \(2018\)](#) note that plug-in wage decompositions of the sort introduced by AKM typically attribute 15%–25% of overall wage variance to variability in firm fixed effects. Given the bias and potential sampling variability associated with plug-in estimates, however, it has been difficult to infer whether firm effects play a differentially important role in certain markets or among particular demographic groups.

In this section, we use Italian social security records to compute leave-out estimates of the AKM wage decomposition and contrast them with estimates based upon the plug-in estimator of [Abowd et al. \(1999\)](#) and the homoscedasticity-corrected estimator of [Andrews et al. \(2008\)](#). We then investigate whether the variance components that comprise the AKM decomposition differ across age groups. While it is well known that wage inequality increases with age ([Mincer et al., 1974](#); [Lemieux, 2006](#)), less is known about the extent to which firm pay premia mediate this phenomenon. Standard wage posting models (e.g., [Burdett and Mortensen, 1998](#)) suggest older workers have had more time to climb (and fall off) the job ladder and to receive outside offers ([Bagger et al., 2014](#)), which may result in more dispersed firm wage premia. But older workers have also had more time to develop professional reputations revealing their relative productivity, which should generate a large increase in the variance of person effects ([Gibbons and Katz, 1992](#); [Gibbons et al., 2005](#)). The tools developed in this paper allow us to formally study these hypotheses.

8.1 Sample Construction

The data used in our analysis come from the Veneto Worker History (VWH) file, which provides the annual earnings and days worked associated with each covered employment spell taking place in the Veneto region of Northeast Italy over the years 1984–2001. The VWH data have been used in a number of recent studies ([Card et al., 2014](#); [Bartolucci et al., 2018](#); [Serafinelli, 2019](#); [Devicienti et al., 2019](#)) and are well suited to the analysis of age differences because they provide precise information on dates of birth. These data are also notable for being publicly available, making the costs of replicating our analysis unusually low.⁷

Our baseline sample consists of workers with employment spells taking place in the years 1999 and 2001, which provides us with a three year horizon over which to measure job mobility. In [Section 8.4](#) we analyze a longer unbalanced sample spanning the years 1996–2001 and find that it yields similar results. For each worker-year pair, we retain the unique employment spell yielding the highest earnings in that year. Wages in each year are defined as earnings in the selected spell

⁷See http://www.frdb.org/page/data/scheda/inps-data-veneto-workers-histories-vwh/doc_pk/11145 for information on obtaining the VWH.

divided by the spell length in days. Workers are divided into two groups of roughly equal size according to their year of birth: “younger” workers born in the years 1965-1983 (aged 18-34 in 1999) and “older” workers born in the years 1937-1964 (aged 35-64 in 1999). Further details on our processing of the VWH records is provided in Appendix A.1.

Table 1 reports the number of person-year observations available among workers employed by firms in the region’s largest connected set, along with the largest connected set for each age group. Workers are classified as “movers” if they switch firms between 1999 and 2001. Comparing the number of movers to half the number of person-year observations reveals that roughly 21% of all workers are movers. The movers share rises to 26% among younger workers while only 16% of older workers are movers, reflecting the tendency of mobility rates to decline with age. The average number of movers per connected firm ranges from nearly 3 in the pooled sample to roughly 2 in the thinner age-specific samples, suggesting that many firms are associated with only a single mover.

Our leave-out estimation strategy requires that each firm effect remain estimable after removing any single observation. The second panel of Table 1 enforces this requirement by restricting to firms that remain connected when any mover is dropped (see Appendix B.1 for computational details). Pruning the sample in this way drops roughly half of the firms but less than a third of the movers and eliminates roughly 30% of all workers regardless of their mobility status. These additional restrictions raise mean wages by roughly 5% and lower the variance of wages by 5–10% depending on the sample.

To assess the potential influence of these sample restrictions on our estimands of interest, we construct a third sample that further requires the firm effects to remain estimable after removing any two observations.⁸ This “leave-two-out connected set” is also of theoretical interest because it provides a setting where the requirements for consistency of the variance estimator of Lemma 5 appear to be satisfied. On average, the leave-two-out connected sets have roughly half as many firms and 20% fewer movers than the corresponding leave-one-out sets, and the average number of movers per firm ranges from approximately 5.6 in the sample of older workers to 4.3 in the sample of younger workers. Restricting the sample in this way further raises mean wages by 3–4% but yields negligible changes in variance, except among the sample of older workers, which experiences a nearly 7% *increase* in variance. We investigate below the extent to which these changes in unconditional variances reflect changes in the variance of underlying firm wage effects.

8.2 AKM Model and Design Diagnostics

Consider the following simplified version of the AKM model:

$$y_{gt} = \alpha_g + \psi_{j(g,t)} + \varepsilon_{gt}. \quad (g = 1, \dots, N, t = 1, 2)$$

⁸We thank an anonymous referee for this suggestion.

We fit models of this sort to the VHW data after having pre-adjusted log wages for year effects in a first step. This adjustment is obtained by estimating an augmented version of the above model by OLS that includes a dummy control for the year 2001. Hence, y_{gt} gives the log wage in year t minus a year 2001 dummy times its estimated coefficient. This two-step approach simplifies computation without compromising consistency because the year effect is estimated at a \sqrt{N} rate.

The bottom of Table 1 reports for each sample the maximum leverage ($\max_i P_{ii}$) of any person-year observation (Appendix B.3 discusses the computation of these leverages). While our pruning procedure ensures $\max_i P_{ii} < 1$, it is noteworthy that $\max_i P_{ii}$ is still quite close to one, indicating that certain person-year observations remain influential on the parameter estimates. This finding highlights the inadequacy of asymptotic approximations that require the dimensionality of regressors to grow slower than the sample size, which would lead the maximum leverage to tend to zero.

The asymptotic results of Section 5 emphasize the importance of not only the maximal leverage, but the number and severity of any bottlenecks in the mobility network. Figure 1 illustrates the leave-two-out connected set for older workers. Each firm is depicted as a dot, with the size of the dot proportional to the total number of workers employed at the firm over the years 1999 and 2001. Dots are connected when a worker moves between the corresponding pair of firms. The figure highlights the two most severe bottlenecks in this network, which divide the firms into three distinct blocks. Each block’s firms have been shaded a distinct color. The blue block consists of only five firms, four of which are quite small, which limits its influence on the asymptotic behavior of our estimator. However, the green block has 51 firms with a non-negligible employment share of 9.5%. Theorem 3 and the discussion in Section 7 therefore suggest that the bottleneck between the green and the larger red block will generate weak identification and asymptotic non-normality, predictions we explore in detail below.

8.3 Variance Decompositions

Table 2 reports the results of applying to our samples three estimators of the AKM variance decomposition: the naive plug-in (PI) estimator $\hat{\theta}_{PI}$ originally proposed by AKM, the homoscedasticity-only (HO) estimator $\hat{\theta}_{HO}$ of Andrews et al. (2008), and the leave-out (KSS) estimator $\hat{\theta}$. The PI estimator finds that the variance of firm effects in the pooled leave-one-out connected set accounts for roughly 20% of the total variance of wages, while among younger workers firm effect variability is found to account for 31% of overall wage variance. Among older workers, variability in firm effects is estimated to account for only 16% of the variance of wages in the leave-one-out connected set.

Are these age differences driven by biases attributable to estimation error? Applying the HO estimator of Andrews et al. (2008) reduces the estimated variances of firm effects by roughly 18% in

the age-pooled sample, 27% in the sample of younger workers, and 16% in the sample of older workers. However, the KSS estimator yields further, comparably sized, reductions in the estimated firm effect variance relative to the HO estimator, indicating the presence of substantial heteroscedasticity in these samples. For instance, in the pooled leave-one-out sample, the KSS estimator finds a variance of firm effects that accounts for only 13% of the overall variance of wages, while the HO estimator finds that firm effects account for 16% of wage variance.

Moreover, while the plug-in estimates suggested that the firm effect variance was greater among older than younger workers, the KSS estimator finds the opposite pattern. The KSS estimator also finds that the pooled variance of firm effects exceeds the corresponding variance in either age-specific sample, a sign that mean firm effects differ by age. We explore this between age group component of firm variability in greater depth below.

A potential concern with analyzing the leave-one-out connected set is that worker and firm behavior in this sample may be non-representative of the broader (just-)connected set. To assess this possibility, we also report estimates for the leave-two-out connected set. Remarkably, the KSS estimator finds negligible differences in the variance of firm effects between the leave-one-out and leave-two-out samples for both the pooled sample and the sample of younger workers. Among older workers the estimated firm effect variance falls by about 11% in the leave-two-out sample, though we show below that this difference may be attributable to sampling variation. The broad similarity between leave-one-out and leave-two-out KSS estimates is likely attributable to the fact that trimmed firms tend to be small and therefore contribute little to the person-year weighted variance of firm effects that has been the focus of the literature.

PI estimates of person effect variances are much larger than the corresponding estimates of firm effect variance, accounting for 66%–88% of the total variance of wages depending on the sample. The PI estimator also finds that person effects are much more dispersed among older than younger workers, which is in accord with standard models of human capital accumulation and employer learning. The estimated ratio of older to younger person effect variances in the leave-one-out sample is roughly 2.6. Applying the HO estimator reduces the magnitude of the person effect variance among all age groups, but boosts the ratio of older to younger person effect variances to 3.2. The KSS estimator yields further downward corrections to estimated person effect variances, leading the contribution of person effect variability to range from only 50% to 80% of total wage variance. Proportionally, however, the variability of older workers remains stable at 3.2 times that of younger workers.

PI estimates of the covariance between worker and firm effects are negative in both age-restricted samples, though not in the pooled sample. When converted to correlations, these figures suggest there is mild negative assortative matching of workers to firms. Applying the HO estimator leads the covariances to change sign in both age-specific samples, while generating a mild increase in the estimated covariance of the pooled sample. In all three samples, however, the HO estimates indicate

very small correlations between worker and firm effects. By contrast, the KSS estimator finds a rather strong positive correlation of 0.21 among younger workers, 0.27 among older workers, and 0.28 in the pooled leave-one-out sample, indicating the presence of non-trivial positive assortative matching between workers and firms. While the patterns in the leave-two-out sample are broadly similar, the KSS correlation estimate among older workers is substantially smaller in the leave-two-out than the leave-one-out sample (0.18 vs 0.27).

Finally, we examine the overall fit of the two-way fixed effects model using the coefficient of determination. The PI estimator of R^2 suggests the two-way fixed effects model explains more than 95% of wage variation in the pooled sample, 91% in the sample of younger workers, and 97% in the sample of older workers. The HO estimator of R^2 is equivalent to the adjusted R^2 measure of [Theil \(1961\)](#). The adjusted R^2 indicates that the two-way fixed effects model explains roughly 90% of the variance of wages in the pooled sample, which is quite close to the figures reported in [Card et al. \(2013\)](#) for the German labor market. Applying the KSS estimator yields very minor changes in estimated explanatory power relative to the HO estimates. Interestingly, a sample size weighted average of the age group specific KSS R^2 estimates lies slightly below the pooled KSS estimate of R^2 , which suggests allowing firm effects to differ by age group fails to appreciably improve the model’s fit. We examine this hypothesis more carefully in [Section 8.5](#).

8.4 Multiple Time Periods and Serial Correlation

Thus far, our analysis has relied upon panels with only two time periods. [Table 3](#) reports KSS estimates of the variance of firm effects in an unbalanced panel spanning the years 1996–2001. To analyze this longer panel, we expand our set of time varying covariates to include unrestricted year effects and a third order polynomial in age normalized to have slope zero at age 40 as discussed in [Card et al. \(2018\)](#).⁹ Allowing up to six wage observations per worker yields a substantially larger estimation sample with roughly three times more person-year observations in the age-pooled leave-one-out connected set than was found in [Table 1](#). For older workers, who have especially low mobility rates, allowing more time periods raises the number of person-year observations in the leave-one-out connected set by a factor of roughly 5.7 and more than triples the number of firms.

While these additional observations will tend to reduce the bias in the plug-in estimator, using longer panels may present two distinct sets of complications. First, the equivalence discussed in [Remark 4](#) no longer holds, which implies that leaving a single person-year observation out is unlikely to remove the bias in estimates of the variance of firm effects when the errors are serially correlated. Second, pooling many years of data may change the target parameter if firm or person

⁹Pre-adjusting for age has negligible effects on the variance decompositions reported in [Table 2](#) but is quantitatively more important in this longer panel. Age adjustments are particularly pronounced among younger workers who generally exhibit greater wage growth and tend to move rapidly to higher paying firms.

effects “drift” with time. The bottom rows of Table 3 probe for the importance of serial correlation by leaving out “clusters” of observations – as described in Remark 3 – defined successively as all observations within the same worker-firm “match” and all observations belonging to the same worker; see Appendix B.3.1 for computational details. Because worker g ’s person effect is not estimable when leaving that worker’s entire wage history out, we estimate a within-transformed specification that eliminates the person effects in a first step.

Leaving out the match yields an important reduction in the variance of firm effects relative to leaving out a single person-year observation, indicating the presence of substantial serial correlation within match. By contrast, leaving out the worker turns out to have negligible effects on the estimated variance of firm effects, suggesting that serial correlation across-matches is negligible. As expected, pooling several years of data reduces the bias of the PI estimator: the magnitude of the difference between the PI estimates of the variance of firm effects and the leave-worker-out estimates tends to be smaller than the corresponding difference between the PI and KSS estimates of the variance of firm effects reported in Table 2.

Remarkably, the firm effect variance estimates that result from leaving out either the match or worker are nearly identical to the KSS estimates reported in Table 2 for both the age-pooled samples and the samples of younger workers, suggesting the firm effects are relatively stable over this longer horizon. Among older workers, the leave-cluster-out estimates of the variance of firm effects are higher than those reported in Table 2, which is unsurprising given that the number of firms under consideration more than tripled in this longer panel. Reassuringly, however, Table 3 reveals that the KSS estimates of the variance of firm effects among older workers in the leave-one-out and leave-two-out connected sets are very close to one another. The general stability of the KSS estimates of firm effect variances to alternate panel lengths may be attributable to the relatively placid macroeconomic conditions present in Veneto over this period, see the discussion in [Devicienti et al. \(2019\)](#).

Our leave-cluster-out exercises suggest researchers seeking to analyze longer panels may be able to avoid biases stemming from serial correlation by simply collapsing the data to match means in a first step and then analyzing these means using the leave-one-observation-out estimator. This two-step approach should substantially reduce computational time while generating only mild efficiency losses due to equal weighting of matches. In what follows, we revert to our baseline sample with exactly two observations per worker.

8.5 Sorting and Wage Structure

The KSS estimates reported in Table 2 indicate that older workers exhibit somewhat less variable firm effects and a stronger correlation between person and firm effects than younger workers. These findings might reflect lifecycle differences in the sorting of workers to firms or differences in the

structure of firm wage effects across the two age groups.

Table 4 explores the sorting channel by projecting the pooled firm effects from the leave-one-out sample onto a constant, an indicator for being an older worker, the log of firm size, and the interaction of the indicator with log firm size. Because these projection coefficients are linear combinations of the estimated firm effects, we use the KSS standard errors proposed in equation (7) and analyzed in Theorem 1. For comparison, we also report a naive standard error that treats the firm effect estimates as independent observations and computes the usual Eicker-White “robust” standard errors. In all cases, the KSS standard error is at least twice the corresponding naive standard error and in one case roughly 24 times larger. In light of the consistency results of Theorem 1, this finding suggests the standard practice of regressing firm effect estimates on observables in a second step without adjusting the standard errors for correlation across firm effects can yield highly misleading inferences.

The first column of Table 4 shows that older workers tend to work at firms with higher average firm effects. Evidently older workers do occupy the upper rungs of the job ladder. The second column shows that this sorting relationship is largely mediated by firm size. An older worker at a firm with a single employee is estimated to have a mean firm wage effect 0.16 log points lower than a younger worker at a firm of the same size, an economically insignificant difference that is also revealed to be statistically insignificant when using the KSS standard error. As firm size grows, older workers begin to enjoy somewhat larger firm wage premia. Evaluated at the median firm size of 12 workers, the predicted gap between older and younger workers rises to 0.54 log points, a gap that we can distinguish from zero at the 5% level using the KSS standard error but is still quite modest. We conclude that the tendency of older workers to be employed at larger firms is a quantitatively important driver of the firm wage premia they enjoy.

Figure 2 investigates to what extent the firm wage effects differ between age groups. Using the age-restricted leave-one-out connected sets, we obtain a pair of age group specific firm effect estimates $\{\hat{\psi}_j^Y, \hat{\psi}_j^O\}_{j \in \mathcal{J}}$ for the set \mathcal{J} of 8,578 firms present in both samples (see Appendix B.5 for details). Figure 2 plots the person-year weighted averages of $\hat{\psi}_j^Y$ and $\hat{\psi}_j^O$ within each centile bin of $\hat{\psi}_j^O$. A person-year weighted projection of $\hat{\psi}_j^Y$ onto $\hat{\psi}_j^O$ yields a slope of only 0.501. To correct this plug-in slope estimate for attenuation bias, we multiply the unadjusted slope by the ratio of the PI estimate of the person-year weighted variance of ψ_j^O to the corresponding KSS estimate of this quantity. Remarkably, this exercise yields a projection slope of 0.987, suggesting that, were it not for the estimation error in $\hat{\psi}_j^O$, the conditional averages depicted in Figure 2 would be centered around the dashed 45 degree line. Converting this slope into a correlation using the KSS estimate of the person-year weighted variance of ψ_j^Y yields a person-year weighted correlation between the two sets of firm effects of 0.89, which indicates the underlying (ψ_j^Y, ψ_j^O) pairs are tightly clustered around this 45 degree line.

Theorem 2 allows us to formally test the joint null hypothesis that the two sets of firm effects

are actually identical, i.e., that both the slope and R^2 from a projection of ψ_j^Y onto ψ_j^O are one. We can state this hypothesis as $H_0 : \psi_j^O = \psi_j^Y$ for all $j \in \mathcal{J}$. Using the test suggested in Remark 6 we obtain a realized test statistic of 3.95 which, when compared to the right tail of a standard normal distribution, yields a p-value on H_0 of less than 0.1%. Hence, we can decisively reject the null hypothesis that older and younger workers face exactly the same vectors of firm effects. However, our earlier correlation results suggest that H_0 nonetheless provides a fairly accurate approximation to the structure of firm effects, at least among those firms that employ movers of both age groups.

8.6 Inference

We now study more carefully the problem of inference on the variance of firm effects. For convenience, the top row of Table 5 reprints our earlier KSS estimates of the variance of firm effects in each sample. Below each estimate of firm effect variance is a corresponding standard error estimate, computed according to the approach described in Lemma 5. As noted in Remark 8, these standard errors will be somewhat conservative when there is a large share of observations for which no split sample predictions can be created. In the leave-one-out samples this share varies between 15% and 22%, indicating that the standard errors are likely upward biased. In the leave-two-out samples, however, this source of bias is not present as the split sample predictions always exist. The standard errors will also tend to be conservative when there is a large share of observation pairs in the set \mathcal{B} , for which there is upward bias in the estimator of the error variance product. However, for both the leave-one-out and leave-two-out samples, this share varies between only 0.03% and 0.46%, suggesting only a small degree of upward bias stems from this source.

The next panel of Table 5 reports the 95% confidence intervals that arise from setting $q = 0$, $q = 1$, or $q = 2$. While the first interval employs a normal approximation, the latter two allow for weak identification by employing non-standard limiting distributions involving linear combinations of normal and χ^2 random variables. We also report estimates of the curvature parameters (κ_1, κ_2) used to construct the weak identification robust intervals. In the pooled samples both curvature parameters are estimated to be quite small, indicating that a normal approximation is likely to be accurate. Accordingly, setting $q > 0$ has little discernible effect on the resulting confidence intervals in these samples. However, among older workers, particularly in the leave-two-out sample, we find stronger curvature coefficients suggesting weak identification may be empirically relevant. Setting $q > 0$ in this sample widens the confidence interval somewhat and also changes its shape: mildly shortening the lower tail of the interval but lengthening the upper tail.

Treating the samples of younger and older workers as independent, the fact that the confidence intervals for the two age group samples overlap implies we cannot reject the null hypothesis that the firm effect variances are identical at the $(1 - 0.95^2) \times 100 = 9.75\%$ level. The significance of the 0.23 log point difference between the leave-one-out and leave-two-out estimates of firm effect

variance in the sample of older workers turns out to be more difficult to assess. By the Cauchy-Schwartz inequality, the covariance between the leave-one-out and leave-two-out estimators is at most $(0.0026)^2(0.0014)^2 = 3.64 \times 10^{-6}$. Hence the standard error on the difference between the two estimators is at least 0.0012, which implies a maximal t-statistic of 1.92. Therefore, even when using a normal approximation, we find rather weak evidence against the null that the leave-one-out and leave-two-out estimands are equal. However, because the leave-one-out standard error estimator is likely upward biased, this finding is somewhat less conclusive than would typically be the case.

Theorem 3 suggests two important diagnostics for the asymptotic behavior of our estimator are the Lindeberg statistics $\{\max_i w_{is}^2\}_{s=1,2}$ and the top eigenvalue shares $\{\lambda_s^2 / \sum_{\ell=1}^r \lambda_\ell^2\}_{s=1,2,3}$. The bottom panel of Table 5 reports these statistics for each sample. The top eigenvalue shares are fairly small in the pooled sample and among younger workers. A small top eigenvalue share indicates the estimator does not depend strongly on any particular linear combination of firm effects and hence that a normal distribution should provide a suitable approximation to the estimator's asymptotic behavior (i.e. that $q = 0$). Accordingly, we find that the confidence intervals are virtually identical for all values of q in both the pooled samples and the two samples of younger workers.

Among older workers the top eigenvalue share is 31% in the leave-one-out sample and 58% in the leave-two-out sample. The next largest eigenvalue share is, in both cases, less than 5%, which suggests this is a setting where $q = 1$. In line with this view, confidence intervals based upon the $q = 1$ and $q = 2$ approximations are nearly identical in both samples of older workers. The accuracy of these weak-identification robust confidence intervals hinges on the Lindeberg condition of Theorem 3 being satisfied. One can think of the Lindeberg statistic $\max_i w_{is}^2$ as giving an inverse measure of effective sample size available for estimating the linear combination of firm effects associated with the s 'th largest eigenvalue. The fact that these statistics are all less than or equal to 0.05 implies an effective sample size of at least 20. We study in the Monte Carlo exercises below whether this effective sample size is sufficient to provide accurate coverage. Reassuringly, the sum of squared eigenvalues is quite small in all six samples considered, indicating that the leave out estimator is consistent also in our weakly identified settings.

8.7 Monte Carlo Experiments

We turn now to studying the finite sample behavior of the leave-out estimator of firm effect variance and its associated confidence intervals under a particular data generating process (DGP). Data were generated from the following first differenced model based upon equation (4):

$$\Delta y_g = \Delta f_g' \hat{\psi}^{scale} + \Delta \varepsilon_g, \quad (g = 1, \dots, N).$$

Here $\hat{\psi}^{scale}$ gives the $J \times 1$ vector of OLS firm effect estimates found in the pooled leave-one-out sample, rescaled to match the KSS estimate of firm effect variance for that sample. The errors $\Delta\varepsilon_g$ were drawn independently from a normal distribution with variances given by the following model of heteroscedasticity:

$$\mathbb{V}[\Delta\varepsilon_g] = \exp(a_0 + a_1 B_{gg} + a_2 P_{gg} + a_3 \ln L_{g2} + a_4 \ln L_{g1}),$$

where L_{gt} gives the size of the firm employing worker g in period t . To choose the coefficients of this model, we estimated a nonlinear least squares fit to the $\hat{\sigma}_g^2$ in the pooled leave-one-out sample, which yielded the following estimates:

$$\hat{a}_0 = -3.3441, \quad \hat{a}_1 = 1.3951, \quad \hat{a}_2 = -0.0037, \quad \hat{a}_3 = -0.0012, \quad \hat{a}_4 = -0.0086.$$

For each sample, we drew from the above DGP 1,000 times while holding firm assignments fixed at their sample values.

Table 6 reports the results of this Monte Carlo experiment. In accord with theory, the KSS estimator of firm effect variances is unbiased while the PI and HO estimators are biased upwards. As expected, the KSS standard error estimator exhibits a modest upward bias in the leave-one-out samples ranging from 15% in the sample of older workers to 44% among younger workers. In the leave-two-out sample, however, the standard error estimator exhibits biases of only 6% or less. Unsurprisingly then, the $q = 0$ confidence interval over-covers in both the pooled leave-one-out sample and the leave-one-out sample of younger workers. In the corresponding leave-two-out samples, however, coverage is very near its nominal level, both for the normal based ($q = 0$) and the weak identification robust ($q = 1$) intervals.

In the samples of older workers, the normal distribution provides a poor approximation to the shape of the estimator's sampling distribution, which is to be expected given the large top eigenvalues found in these designs. This non-normality generates substantial under-coverage by the $q = 0$ confidence interval in the leave-two-out sample. Applying the weak identification robust interval in the leave-two-out sample of older workers yields coverage very close to nominal levels despite the fact that the effective sample size available for the top eigenvector is only about 20.

In sum, the Monte Carlo experiments demonstrate that confidence intervals predicated on the assumption that $q = 1$ can provide accurate size control in leave-two-out samples when the realized mobility network exhibits a severe bottleneck. We also achieved size control in leave-one-out samples, albeit at the cost of moderate over-coverage. Hence, in applications where statistical power is a first-order consideration, it may be attractive to restrict attention to leave-two-out samples, which tend to yield estimates of variance components very close to those found in leave-one-out samples but with substantially less biased standard errors.

9 Conclusion

We propose a new estimator of quadratic forms with applications to several areas of economics. The estimator is finite sample unbiased in the presence of unrestricted heteroscedasticity and can be accurately approximated in very large datasets via random projection methods. Consistency is established under verifiable design requirements in an environment where the number of regressors may grow in proportion to the sample size. The estimator enables tests of linear restrictions of varying dimension under weaker conditions than have been explored in previous work. A new distributional theory highlights the potential for the proposed estimator to exhibit deviations from normality when some linear combinations of coefficients are imprecisely estimated relative to others.

In an application to Italian worker-firm data, we showed that ignoring heteroscedasticity can substantially bias conclusions about the relative contribution of workers, firms, and worker-firm sorting to wage inequality. Accounting for serial correlation within a worker-firm match was found to be empirically important, while across match correlation appears to be negligible. Consequently, those studying longer panels may wish to collapse their data down to match level means and then apply the leave-observation-out estimator. Alternately, researchers can simply extract and analyze separately balanced panels of length two, which also facilitates analysis of the temporal stability of the firm and person effect variances.

Leave-out standard error estimates for the coefficients of a linear projection of firm effects onto worker and firm observables were found to be several times larger than standard errors that naively treat the estimated firm effects as independent. These results strongly suggest that researchers seeking to identify the observable correlates of high-dimensional fixed effects should consider employing the proposed standard errors, including when studying settings falling outside the traditional worker-firm setup (e.g., [Finkelstein et al., 2016](#); [Chetty and Hendren, 2018](#)). Stratifying our analysis by birth cohort, we formally rejected the null hypothesis that older and younger workers face identical vectors of firm effects but found that the two sets of firm effects were highly correlated. Corresponding techniques can be used to study multivariate models.

A Monte Carlo analysis demonstrated that bottlenecks in the worker-firm mobility network can generate quantitatively important deviations from normality. The proposed inference procedure captured these deviations accurately with a weak identification robust confidence interval. In cases where the mobility network was strongly connected, accurate inferences were obtained with a normal approximation. Our results suggest that in typical worker-firm applications, the normal approximation is likely to suffice. However, when studying small areas, or sub-populations with limited mobility, accounting for weak identification can be quantitatively important.

References

- Abowd, J. M., R. H. Creedy, F. Kramarz, et al. (2002). Computing person and firm effects using linked longitudinal employer-employee data. Technical report, Center for Economic Studies, US Census Bureau.
- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67(2), 251–333.
- Achlioptas, D. (2003). Database-friendly random projections: Johnson-lindenstrauss with binary coins. *Journal of computer and System Sciences* 66(4), 671–687.
- Akritas, M. G. and N. Papadatos (2004). Heteroscedastic one-way anova and lack-of-fit tests. *Journal of the American Statistical Association* 99(466), 368–382.
- Anatolyev, S. (2012). Inference in regression models with many regressors. *Journal of Econometrics* 170(2), 368–382.
- Andrews, D. W. K. and X. Cheng (2012). Estimation and inference with weak, semi-strong, and strong identification. *Econometrica* 80(5), 2153–2211.
- Andrews, I. and A. Mikusheva (2016). A geometric approach to nonlinear econometric models. *Econometrica* 84(3), 1249–1264.
- 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(3), 673–697.
- Angrist, J., G. Imbens, and A. Krueger (1999). Jackknife instrumental variables estimation. *Journal of Applied Econometrics* 14(1), 57–67.
- Angrist, J. D. (2014). The perils of peer effects. *Labour Economics* 30, 98–108.
- Arcidiacono, P., G. Foster, N. Goodpaster, and J. Kinsler (2012). Estimating spillovers using panel data, with an application to the classroom. *Quantitative Economics* 3(3), 421–470.
- Arellano, M. and S. Bonhomme (2011). Identifying distributional characteristics in random coefficients panel data models. *The Review of Economic Studies* 79(3), 987–1020.
- 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(6), 1551–96.

- Bartolucci, C., F. Devicienti, and I. Monzón (2018). Identifying sorting in practice. *American Economic Journal: Applied Economics* 10(4), 408–38.
- Bloom, N., F. Guvenen, B. S. Smith, J. Song, and T. von Wachter (2018). The disappearing large-firm wage premium. In *AEA Papers and Proceedings*, Volume 108, pp. 317–22.
- Bonhomme, S. (2017). Econometric analysis of bipartite networks. *Econometric Analysis of Network data edited by B. Graham and A. De Paula*.
- Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica* 87(3), 699–739.
- Burdett, K. and D. T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, 257–273.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics* 36(S1), S13–S70.
- Card, D., A. R. Cardoso, and P. Kline (2015). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly Journal of Economics* 131(2), 633–686.
- Card, D., F. Devicienti, and A. Maida (2014). Rent-sharing, holdup, and wages: Evidence from matched panel data. *The Review of Economic Studies* 81(1), 84–111.
- Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics* 128(3), 967–1015.
- Cattaneo, M. D., M. Jansson, and W. K. Newey (2018). Inference in linear regression models with many covariates and heteroscedasticity. *Journal of the American Statistical Association* 113(523), 1350–1361.
- Chao, J. C., J. A. Hausman, W. K. Newey, N. R. Swanson, and T. Woutersen (2014). Testing overidentifying restrictions with many instruments and heteroskedasticity. *Journal of Econometrics* 178, 15–21.
- Chao, J. C., N. R. Swanson, J. A. Hausman, W. K. Newey, and T. Woutersen (2012). Asymptotic distribution of jive in a heteroskedastic iv regression with many instruments. *Econometric Theory* 28(01), 42–86.
- Chatterjee, S. (2008). A new method of normal approximation. *The Annals of Probability* 36(4), 1584–1610.

- Chetty, R., J. N. Friedman, N. Hilger, E. Saez, D. W. Schanzenbach, and D. Yagan (2011). How does your kindergarten classroom affect your earnings? evidence from project star. *The Quarterly Journal of Economics* 126(4), 1593–1660.
- Chetty, R. and N. Hendren (2018). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics* 133(3), 1163–1228.
- Chung, F. R. (1997). *Spectral graph theory*. Number 92. American Mathematical Soc.
- Cochran, W. G. (1980). Fisher and the analysis of variance. In *RA Fisher: An Appreciation*, pp. 17–34. Springer.
- Devicienti, F., B. Fanfani, and A. Maida (2019). Collective bargaining and the evolution of wage inequality in italy. *British Journal of Industrial Relations* 57(2), 377–407.
- Dhaene, G. and K. Jochmans (2015). Split-panel jackknife estimation of fixed-effect models. *The Review of Economic Studies* 82(3), 991–1030.
- Donald, S. G., G. W. Imbens, and W. K. Newey (2003). Empirical likelihood estimation and consistent tests with conditional moment restrictions. *Journal of Econometrics* 117(1), 55–93.
- Dufour, J.-M. and J. Jasiak (2001). Finite sample limited information inference methods for structural equations and models with generated regressors. *International Economic Review* 42(3), 815–844.
- Efron, B. and C. Stein (1981, 05). The jackknife estimate of variance. *Ann. Statist.* 9(3), 586–596.
- El Karoui, N. and E. Purdom (2018). Can we trust the bootstrap in high-dimensions? the case of linear models. *The Journal of Machine Learning Research* 19(1), 170–235.
- 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(4), 1681–1726.
- Fisher, R. A. (1925). *Statistical methods for research workers*. Genesis Publishing Pvt Ltd.
- Gibbons, R. and L. Katz (1992). Does unmeasured ability explain inter-industry wage differentials? *The Review of Economic Studies* 59(3), 515–535.
- Gibbons, R., L. F. Katz, T. Lemieux, and D. Parent (2005). Comparative advantage, learning, and sectoral wage determination. *Journal of labor economics* 23(4), 681–724.
- Graham, B. S. (2008). Identifying social interactions through conditional variance restrictions. *Econometrica* 76(3), 643–660.

- Graham, B. S., J. Hahn, A. Poirier, and J. L. Powell (2018). A quantile correlated random coefficients panel data model. *Journal of Econometrics* 206(2), 305–335.
- Graham, B. S. and J. L. Powell (2012). Identification and estimation of average partial effects in irregular correlated random coefficient panel data models. *Econometrica* 80(5), 2105–2152.
- Hahn, J. and W. Newey (2004). Jackknife and analytical bias reduction for nonlinear panel models. *Econometrica* 72(4), 1295–1319.
- Hildreth, C. and J. P. Houck (1968). Some estimators for a linear model with random coefficients. *Journal of the American Statistical Association* 63(322), 584–595.
- Horn, S. D., R. A. Horn, and D. B. Duncan (1975). Estimating heteroscedastic variances in linear models. *Journal of the American Statistical Association* 70(350), 380–385.
- Jochmans, K. and M. Weidner (2016). Fixed-effect regressions on network data. *arXiv preprint arXiv:1608.01532*.
- Johnson, W. B. and J. Lindenstrauss (1984). Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics* 26(189-206), 1.
- Kline, P., R. Saggio, and M. Sølvssten (2019). **LeaveOutTwoWay**: A matlab package for leave out estimation of variance components in two way fixed effects models. <https://github.com/rsaggio87/LeaveOutTwoWay>.
- Koutis, I., G. L. Miller, and D. Tolliver (2011). Combinatorial preconditioners and multilevel solvers for problems in computer vision and image processing. *Computer Vision and Image Understanding* 115(12), 1638–1646.
- Kuh, E. (1959). The validity of cross-sectionally estimated behavior equations in time series applications. *Econometrica*, 197–214.
- Lei, L., P. J. Bickel, and N. El Karoui (2018). Asymptotics for high dimensional regression m-estimates: fixed design results. *Probability Theory and Related Fields* 172(3-4), 983–1079.
- Lemieux, T. (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *American Economic Review* 96(3), 461–498.
- MacKinnon, J. G. and H. White (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of econometrics* 29(3), 305–325.
- Menger, K. (1927). Zur allgemeinen kurventheorie. *Fundamenta Mathematicae* 10(1), 96–115.

- Mincer, J. A. et al. (1974). Schooling, experience, and earnings. *NBER Books*.
- Mohar, B. (1989). Isoperimetric numbers of graphs. *Journal of Combinatorial Theory, Series B* 47(3), 274–291.
- Moulton, B. R. (1986). Random group effects and the precision of regression estimates. *Journal of econometrics* 32(3), 385–397.
- Newey, W. K. and J. R. Robins (2018). Cross-fitting and fast remainder rates for semiparametric estimation. *arXiv preprint arXiv:1801.09138*.
- Oliveira, R. I. (2009). Concentration of the adjacency matrix and of the laplacian in random graphs with independent edges. *arXiv preprint arXiv:0911.0600*.
- Phillips, G. D. A. and C. Hale (1977). The bias of instrumental variable estimators of simultaneous equation systems. *International Economic Review*, 219–228.
- Powell, J. L., J. H. Stock, and T. M. Stoker (1989). Semiparametric estimation of index coefficients. *Econometrica: Journal of the Econometric Society*, 1403–1430.
- Quenouille, M. H. (1949). Approximate tests of correlation in time-series. *Journal of the Royal Statistical Society. Series B (Methodological)* 11(1), 68–84.
- Rao, C. R. (1970). Estimation of heteroscedastic variances in linear models. *Journal of the American Statistical Association* 65(329), 161–172.
- Raudenbush, S. and A. S. Bryk (1986). A hierarchical model for studying school effects. *Sociology of education*, 1–17.
- Raudenbush, S. W. and A. S. Bryk (2002). *Hierarchical linear models: Applications and data analysis methods*, Volume 1. Sage.
- Sacerdote, B. (2001). Peer effects with random assignment: Results for dartmouth roommates. *The Quarterly journal of economics* 116(2), 681–704.
- Sarkar, P. and P. J. Bickel (2015). Role of normalization in spectral clustering for stochastic blockmodels. *The Annals of Statistics* 43(3), 962–990.
- Scheffe, H. (1959). *The analysis of variance*. John Wiley & Sons.
- Searle, S. R., G. Casella, and C. E. McCulloch (2009). *Variance components*, Volume 391. John Wiley & Sons.

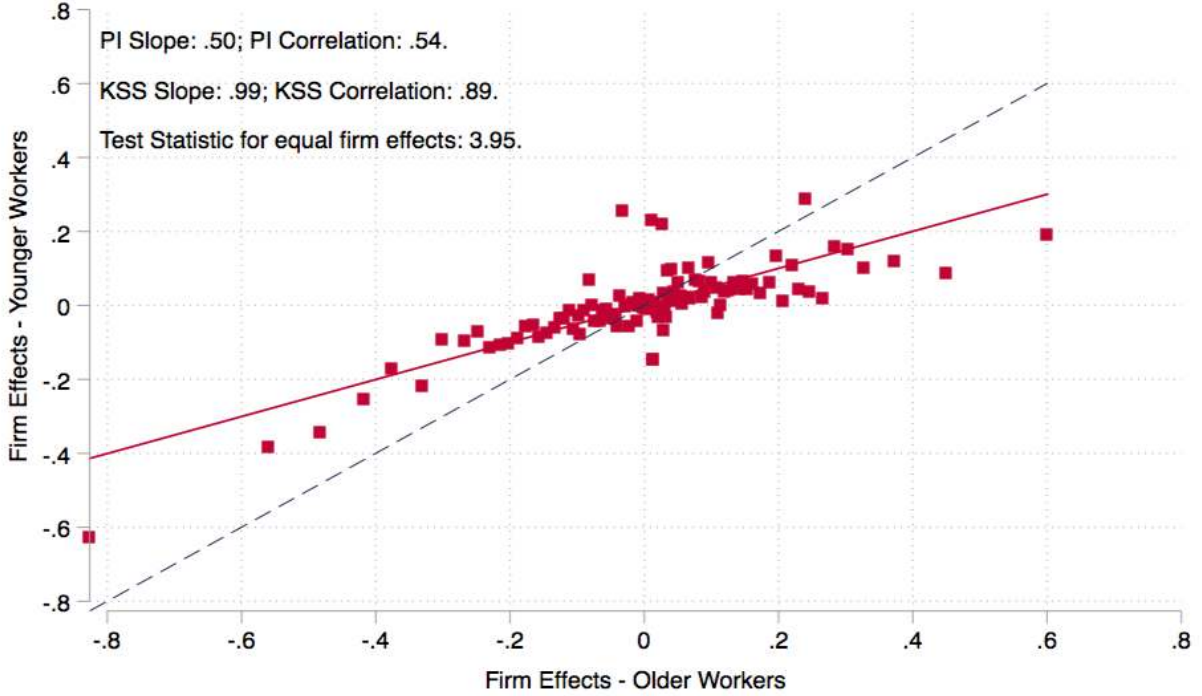
- Serafinelli, M. (2019). “good” firms, worker flows, and local productivity. *Journal of Labor Economics* 37(3), 747–792.
- Sherman, J. and W. J. Morrison (1950). Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics* 21(1), 124–127.
- Silver, D. W. (2016). Essays on labor economics and health care.
- Sølvsten, M. (2019). Robust estimation with many instruments. *Journal of Econometrics*.
- Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2017). Firming up inequality. Technical report, National Bureau of Economic Research.
- Sorkin, I. (2018). Ranking firms using revealed preference. *The quarterly journal of economics* 133(3), 1331–1393.
- Swamy, P. A. (1970). Efficient inference in a random coefficient regression model. *Econometrica*, 311–323.
- Theil, H. (1961). Economic forecasts and policy.
- Verdier, V. (2017). Estimation and inference for linear models with two-way fixed effects and sparsely matched data. *Review of Economics and Statistics* (0).
- Woodbury, M. A. (1949). The stability of out-input matrices. *Chicago, IL* 9.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.
- Wright, S. (1921). Correlation and causation. *Journal of agricultural research* 20(7), 557–585.
- Yen, J. Y. (1971). Finding the k shortest loopless paths in a network. *management Science* 17(11), 712–716.

Figure 1: Realized Mobility Network: Older workers



Note: This figure provides a visualization of the design matrix S_{xx} for the leave-two-out sample of older workers (see Table 1 for reference). The graph is plotted in the statistical software R using the *igraph* package and the large-scale graph layout (DrL) using the option to concentrate firms from the same blocks. High weight mobility refers to observations that have w_{i1}^2 or w_{i2}^2 above 1/500 and these observations form the bottlenecks between the three blocks.

Figure 2: Do Firm Effects Differ Across Age Groups?



Note: This figure plots the mean of the estimated firm effects for younger workers ($\hat{\psi}_j^Y$) by centiles of the estimated firm effects for older workers ($\hat{\psi}_j^O$) in the sample of 8,578 firms for which both sets of effects are leave-one-out identified. Both sets of firm effects are demeaned within this estimation sample. “PI slope” gives the coefficient from a person-year weighted projection of $\hat{\psi}_j^Y$ onto $\hat{\psi}_j^O$. “KSS slope” adjusts for attenuation bias by multiplying the PI slope by the ratio of the plug-in estimate of the person-year weighted variance of $\hat{\psi}_j^O$ to the KSS adjusted estimate of the same quantity. “PI correlation” gives the person-year weighted sample correlation between $\hat{\psi}_j^O$ and $\hat{\psi}_j^Y$ while “KSS correlation” adjusts this correlation for sampling error in both $\hat{\psi}_j^O$ and $\hat{\psi}_j^Y$ using leave out estimates of the relevant variances. “Test statistic” refers to the realization of $\hat{\theta}_{H_0}/\sqrt{\hat{\mathbb{V}}[\hat{\theta}_{H_0}]}$ where $\hat{\theta}_{H_0}$ is the quadratic form associated with the null hypothesis that the firm effects are equal across age groups, see Remark 6 and Appendix B.5 for details. From Theorem 2, $\hat{\theta}_{H_0}/\sqrt{\hat{\mathbb{V}}[\hat{\theta}_{H_0}]}$ converges to a $\mathcal{N}(0, 1)$ under the null hypothesis that $\psi_j^O = \psi_j^Y$ for all 8,578 firms.

Table 1: Summary Statistics

	<u>Pooled</u> [1]	<u>Younger Workers</u> [2]	<u>Older Workers</u> [3]
<u>Largest Connected Set</u>			
Number of Observations	1,859,459	1,011,111	643,020
Number of Movers	197,572	133,627	53,035
Number of Firms	73,933	62,848	26,606
Mean Log Daily Wage	4.7507	4.6741	4.8925
Variance Log Daily Wage	0.1985	0.1321	0.2722
<u>Leave One Out Sample</u>			
Number of Observations	1,319,972	661,528	425,208
Number of Movers	164,203	102,746	35,467
Number of Firms	42,489	33,151	10,733
Mean Log Daily Wage	4.8066	4.7275	4.9455
Variance Log Daily Wage	0.1843	0.1200	0.2591
Maximum Leverage (P_{ii})	0.9365	0.9437	0.9513
<u>Leave Two Out Sample</u>			
Number of Observations	1,126,964	514,284	301,948
Number of Movers	128,705	70,703	21,066
Number of Firms	24,424	16,263	3,756
Mean Log Daily Wage	4.8307	4.7528	4.9834
Variance Log Daily Wage	0.1898	0.1232	0.2760
Maximum Leverage (P_{ii})	0.8333	0.8481	0.8604

Note: Data in column 1 corresponds to VHW observations in the years 1999 and 2001. Column 2 restricts to workers born in the years 1965-1983. Column 3 considers workers born in the years 1937-1964. The largest connected set gives the largest sample in which all firms are connected by worker mobility. The leave out sample is the largest connected set such that every firm remains connected even after removing a worker from the sample, see Appendix B1 for details. The leave two out sample is the largest connected set such that every firm remains connected even after removing two workers from the sample, see Appendix B2 for details. Statistics on log daily wages are person-year weighted.

Table 2: Variance Decomposition

	<u>Pooled</u>		<u>Younger Workers</u>		<u>Older Workers</u>	
	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample
Variance of Log Wages	0.1843	0.1898	0.1200	0.1232	0.2591	0.2760
<u>Variance of Firm Effects</u>						
Plug in (PI)	0.0358	0.0316	0.0368	0.0314	0.0415	0.0304
Homoscedasticity Only (HO)	0.0295	0.0271	0.0270	0.0251	0.0350	0.0243
Leave Out (KSS)	0.0240	0.0238	0.0218	0.0221	0.0204	0.0180
<u>Variance of Person Effects</u>						
Plug in (PI)	0.1321	0.1341	0.0843	0.0827	0.2180	0.2406
Homoscedasticity Only (HO)	0.1173	0.1214	0.0647	0.0663	0.2046	0.2298
Leave Out (KSS)	0.1119	0.1179	0.0596	0.0634	0.1910	0.2221
<u>Covariance Firm, Person Effects</u>						
Plug in (PI)	0.0039	0.0077	-0.0058	-0.0008	-0.0032	-0.0006
Homoscedasticity Only (HO)	0.0097	0.0117	0.0030	0.0049	0.0040	0.0041
Leave Out (KSS)	0.0147	0.0149	0.0075	0.0075	0.0171	0.0115
<u>Correlation of Person, Firm Effects</u>						
Plug in (PI)	0.0565	0.1189	-0.1040	-0.0155	-0.0334	-0.0068
Homoscedasticity Only (HO)	0.1649	0.2033	0.0726	0.1205	0.0475	0.0542
Leave Out (KSS)	0.2830	0.2803	0.2092	0.2016	0.2744	0.1826
<u>Coefficient of Determination</u>						
Plug in (PI)	0.9546	0.9559	0.9183	0.9183	0.9774	0.9781
Homoscedasticity Only (HO)	0.9029	0.9078	0.8184	0.8256	0.9524	0.9551
Leave Out (KSS)	0.8976	0.9047	0.8091	0.8209	0.9489	0.9536

Note: See Table 1 for sample definitions. All variance components are person-year weighted. Wages have been pre-adjusted for a year fixed effect.

Table 3: Variance of Firm Effects under Different Leave-Out Strategies

	<u>Pooled</u>		<u>Younger Workers</u>		<u>Older Workers</u>	
	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample
<i>Descriptive Statistics</i>						
Number of Observations	5,163,446	4,758,713	2,632,596	2,290,223	2,016,202	1,652,729
Number of Movers	440,323	391,313	276,338	229,013	123,777	92,568
Number of Firms	86,646	60,896	71,673	46,084	32,809	16,396
Mean Log Daily Wage	4.7511	4.7662	4.6644	4.6799	4.8841	4.9158
Variance Log Daily Wage	0.1839	0.1854	0.1156	0.1162	0.2477	0.2579
Maximum Leverage (P_{ii})	0.9021	0.7713	0.9056	0.7775	0.9174	0.7639
<i>Variance of firm Effects</i>						
Variance of Firm Effects (PI)	0.0304	0.0278	0.0303	0.0273	0.0376	0.0336
Variance of Firm Effects (Leave Person-Year Out)	0.0296	0.0273	0.0302	0.0272	0.0314	0.0322
Variance of Firm Effects (Leave Match Out)	0.0243	0.0231	0.0221	0.0218	0.0265	0.0280
Variance of Firm Effects (Leave Worker Out)	0.0241	0.0233	0.0227	0.0218	0.0270	0.0268

Note: These samples consist of VHW data spanning the years 1996-2001 (T=6). "Pooled" considers all workers, "Younger workers" were born in the years 1965-1983, "Older workers" were born in the years 1937-1964. Variance of Firm Effects (PI) denotes the uncorrected plug-in estimate of the variance of firm effects. Variance of Firm Effects (Leave Person-Year Out) computes the leave-out bias correction by leaving a single person-year observation out. Variance of Firm Effects (Leave Match Out) computes the leave-out bias correction by leaving entire worker-firm matches out. Variance of Firm Effects (Leave Worker Out) computes the leave-out bias correction by leaving out each worker's entire employment history. See Appendix B3.1 for computational details. Wages have been pre-adjusted for a cubic in age and year fixed effects.

Table 4: Projecting Firm Effects on Covariates

	(1)	(2)
Older Worker	0.0272 (0.0009) [0.0003]	-0.0016 (0.0024) [0.0001]
Log Firm Size		0.0276 (0.0007) [0.0001]
Older Worker x Log Firm Size		0.0028 (0.0005) [0.0002]
Predicted Gap in Firm Effects Older vs. Younger Worker	0.0272 (0.0009) [0.0003]	0.0054 (0.0019) [0.0008]
Number of Observations	1,319,972	1,319,972

Note: This table reports projections of the firm effects estimated in the pooled leave one out sample defined in Table 1 onto worker and firm characteristics. In Column 1, the firm effects are projected onto a constant and a dummy for the workers being born in the years 1937-1964. In Column 2, we regress firm effects on a constant, a dummy for workers born in the years 1937-1964, log firm size and an interaction between these two covariates. The brackets show the standard error associated with a given linear projection coefficient. In round brackets we report the standard error based on Theorem 1. In square brackets we report the naive standard error that one would obtain when regressing firm effects on a given covariate and using standard Eicker-White heteroskedastic robust standard errors. "Predicted Gap in Firm Effects" reports the predicted difference in firm effects between older and younger workers according to either Column 1 or Column 2 evaluated at the median firm size (12 Workers). See text for details.

Table 5: Inference on the Variance of Firm Effects

	<u>Pooled</u>		<u>Younger Workers</u>		<u>Older Workers</u>	
	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample
<u>Variance of Firm Effects</u>						
Leave out estimate	0.0240 (0.0006)	0.0238 (0.0006)	0.0218 (0.0006)	0.0221 (0.0006)	0.0204 (0.0025)	0.0180 (0.0013)
Sum of Squared Eigenvalues	2.1141E-06	1.6236E-06	3.9720E-06	2.9863E-06	0.0001	0.0001
Percentage of movers with no split sample estimator	14.92%	0.00%	21.00%	0.00%	22.00%	0.00%
Percentage of mover pairs in \mathcal{B}	0.04%	0.05%	0.03%	0.05%	0.21%	0.46%
<u>Confidence Intervals</u>						
95% Confidence Intervals - Strong id (q=0)	[0.0228; 0.0251]	[0.0227; 0.0249]	[0.0207; 0.0230]	[0.0210; 0.0232]	[0.0155; 0.0254]	[0.0154; 0.0205]
95% Confidence Intervals - Weak id (q=1)	[0.0228; 0.0251]	[0.0227; 0.0249]	[0.0207; 0.0230]	[0.0209; 0.0234]	[0.0157; 0.0267]	[0.0158; 0.0221]
95% Confidence Intervals - Weak id (q=2)	[0.0228; 0.0251]	[0.0227; 0.0251]	[0.0207; 0.0231]	[0.0209; 0.0234]	[0.0158; 0.0272]	[0.0161; 0.0231]
Curvature - κ_1	0.0182	0.0416	0.0197	0.0689	0.3333	0.7845
Curvature - κ_2	0.0271	0.1458	0.0444	0.0678	0.2779	1.1088
<u>Diagnostics</u>						
Lindeberg Condition - q=1	0.1878	0.0865	0.2061	0.0371	0.0359	0.0503
Lindeberg Condition - q=2	0.0866	0.1604	0.0237	0.1639	0.0491	0.0479
<u>Eigenvalues Ratios</u>						
Eigenvalue Ratio - 1	0.0135	0.0233	0.0189	0.0211	0.3132	0.5883
Eigenvalue Ratio - 2	0.0131	0.0202	0.0058	0.0135	0.0387	0.0499
Eigenvalue Ratio - 3	0.0112	0.0131	0.0050	0.0087	0.0314	0.0187

Note: This table conducts inference on the variance of firm effects using the samples described in Table 1. The round brackets report standard error estimates constructed according to the procedure described in Section 4.2. Associated confidence intervals are computed under different assumptions on q , as described in Section 6.1. Curvature reports the maximal curvature, see Appendix C6.1 for further details. "Eigenvalue ratio - 1" gives the ratio of the squared largest eigenvalue of the matrix \tilde{A} to the sum of all its squared eigenvalues. "Eigenvalue Ratio - 2" and "Eigenvalue Ratio - 3" report the same ratio using the second and third largest eigenvalues respectively. "Percentage of movers with no split sample estimator" reports the percentage of movers for which it is impossible to find two independent unbiased estimators of their conditional mean. "Percentage of mover pairs in \mathcal{B} " gives the fraction of mover pairs where we use an unconditional variance estimate, see Section 4.2 for details.

Table 6: Montecarlo Results

	Pooled		Younger Workers		Older Workers	
	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample	Leave one out sample	Leave two out sample
True Variance of Firm effects	0.0240	0.0208	0.0219	0.0186	0.0222	0.0199
<i>Mean, Standard deviation across Simulations</i>						
Variance of Firm Effects - Leave Out (KSS)	0.0239 (0.00037)	0.0208 (0.00037)	0.0219 (0.00040)	0.0186 (0.00039)	0.0223 (0.00221)	0.0198 (0.00276)
Variance of Firm Effects - Homescedatic Only (HO)	0.0336 (0.00036)	0.0260 (0.00036)	0.0328 (0.00037)	0.0242 (0.00038)	0.0345 (0.00216)	0.0264 (0.00275)
Variance of Firm Effects - Plug in (PI)	0.0352 (0.00036)	0.0270 (0.00036)	0.0352 (0.00036)	0.0256 (0.00038)	0.0360 (0.00216)	0.0273 (0.00275)
Mean Estimated Standard error:	0.00050	0.00039	0.00058	0.00042	0.00254	0.00273
<i>Coverage Rate</i>						
Leave out - Strong id (q=0)	0.9890	0.9610	0.9970	0.9600	0.9720	0.8820
Leave out - Weak id (q=1)	0.9900	0.9640	0.9970	0.9630	0.9800	0.9590

Note: See Section 8.7 for description of data generating process (DGP). "True variance of firm effects" gives the variance of the firm effects in the DGP. KSS, PI, and HO report the average of the estimated variance of firm effects across simulations for the three different estimators. Round brackets report the standard deviation of estimator across simulations. "Leave Out - Strong Id" reports the coverage rate of confidence interval for the variance of firms effects based upon KSS standard error and a normal approximation. "Leave Out - Weak Id" reports the coverage rate of the confidence interval described in Section 6.1 under q=1. All results based upon 1,000 Monte Carlo draws.

Appendix A Data

This Appendix describes construction of the data used in the application of Section 8.

A.1 Veneto Workers History

Our data come from the Veneto Workers History (VWH) file, which provides social security based earnings records on annual job spells for all workers employed in the Italian region of Veneto at any point between the years 1975 and 2001. Each job-year spell in the VWH lists a start date, an end date, the number of days worked that year, and the total wage compensation received by the employee in that year. The earnings records are not top-coded. We also observe the gender of each worker and several geographic variables indicating the location of each employer. See [Card, Devicienti, and Maida \(2014\)](#) and [Serafinelli \(2019\)](#) for additional discussion and analysis of the VWH.

We consider data from the years 1984–2001 as prior to that information on days worked tend to be of low quality. To construct the person-year panel used in our analysis, we follow the sample selection procedures described in [Card, Heining, and Kline \(2013\)](#). First, we drop employment spells in which the worker’s age lies outside the range 18–64. The average worker in this sample has 1.21 jobs per year. To generate unique worker-firm assignments in each year, we restrict attention to spells associated with “dominant jobs” where the worker earned the most in each corresponding year. From this person-year file, we then exclude workers that (i) report a daily wage less than 5 real euros or have zero days worked (1.5% of remaining person-year observations) (ii) report a log daily wage change one year to the next that is greater than 1 in absolute value (6%) (iii) are employed in the public sector (10%) or (iv) have more than 10 jobs in any year or that have gender missing (0.1%).

Appendix B Computation

This Appendix describes the key computational aspects of the leave-out estimator $\hat{\theta}$, with an emphasis on the application to two-way fixed effects models with two time periods discussed in Example 4 and Section 8.

B.1 Leave-One-Out Connected Set

Existence of $\hat{\theta}$ requires $P_{ii} < 1$ (see Lemma 1) and the following describes an algorithm which prunes the data to ensure that $P_{ii} < 1$. In the two-way fixed effects model of Section 8.2, this condition requires that the bipartite network formed by worker-firm links remains connected when any one worker is removed. This boils down to finding workers that constitute cut vertices or *articulation points* in the corresponding bipartite network.

The algorithm below takes as input a connected bipartite network \mathcal{G} where workers and firms are vertices. Edges between two vertices correspond to the realization of a match between a worker and a firm (see Jochmans and Weidner, 2016; Bonhomme, 2017, for discussion). In practice, one typically starts with a \mathcal{G} corresponding to the *largest connected component* of a given bipartite network (see, e.g., Card et al., 2013). The output of the algorithm is a subset of \mathcal{G} where removal of any given worker does not break the connectivity of the associated graph.

The algorithm relies on existing functions that efficiently finds articulation points and largest connected components. In MATLAB such functions are available in the *Boost Graph Library* and in R they are available in the *igraph* package.

Algorithm 1 Leave-One-Out Connected Set

- 1: **function** PRUNINGNETWORK(\mathcal{G}) ▷ $\mathcal{G} \equiv$ Connected bipartite network of firms and workers
 - 2: Construct \mathcal{G}_1 from \mathcal{G} by deleting all workers that are articulation points in \mathcal{G}
 - 3: Let \mathcal{G} be the largest connected component of \mathcal{G}_1
 - 4: Return \mathcal{G}
 - 5: **end function**
-

The algorithm typically completes in less than a minute for datasets of the size considered in our application. Furthermore, the vast majority of firms removed using this algorithm are only associated with one mover.

B.2 Leave-Two-Out Connected Set

We also introduced a leave-two-out connected set, which is a subset of the original data such that removal of any *two* workers does not break the connectedness of the bipartite network formed by

worker-firm links. The following algorithm proceeds by applying the idea in Algorithm 1 to each of the networks constructed by dropping one worker. A crucial difference from Algorithm 1 is that *two* workers who do not break connectedness in the input network may break connectedness when other workers have been removed. For this reason, the algorithm runs in an iterative fashion until it fails to remove any additional workers.

Algorithm 2 Leave-Two-Out Connected Set

```

1: function PRUNINGNETWORK2( $\mathcal{G}$ )  ▷  $\mathcal{G} \equiv$  Leave-one-out connected bipartite network
   of firms and workers
2:    $a = 1$ 
3:   while  $a > 0$  do
4:      $\mathcal{G}^{del} = \emptyset$ 
5:     for  $g = 1, \dots, N$  do
6:       Construct  $\mathcal{G}_1$  from  $\mathcal{G}$  by deleting worker  $g$ 
7:       Add all workers that are articulation points in  $\mathcal{G}_1$  to  $\mathcal{G}^{del}$ 
8:     end for
9:      $a = |\mathcal{G}^{del}|$ 
10:    if  $a > 0$  then
11:      Construct  $\mathcal{G}_1$  from  $\mathcal{G}$  by deleting all workers in  $\mathcal{G}^{del}$ 
12:      Let  $\mathcal{G}_2$  be the largest connected component of  $\mathcal{G}_1$ 
13:      Let  $\mathcal{G}$  be the output of applying Algorithm 1 to  $\mathcal{G}_2$ 
14:    end if
15:  end while
16:  Return  $\mathcal{G}$ 
17: end function

```

B.3 Computing $\hat{\theta}$

Our proposed leave-out estimator is a function of the $2n$ quadratic forms

$$P_{ii} = x_i' S_{xx}^{-1} x_i \quad B_{ii} = x_i' S_{xx}^{-1} A S_{xx}^{-1} x_i \quad \text{for } i = 1, \dots, n.$$

The estimates reported in Section 8 of the paper rely on exact computation of these quantities. In our application, k is on the order of hundreds of thousands, making it infeasible to compute S_{xx}^{-1} directly. To circumvent this obstacle, we instead compute the k -dimensional vector $z_{i,exact} = S_{xx}^{-1} x_i$ separately for each $i = 1, \dots, n$. That is, we solve separately for each column of Z_{exact} in the system

$$\underset{k \times k}{S_{xx}} \underset{k \times n}{Z_{exact}} = \underset{k \times n}{X'}.$$

We then form $P_{ii} = x'_i z_{i,exact}$ and $B_{ii} = z'_{i,exact} A z_{i,exact}$. The solution $z_{i,exact}$ is computed via MATLAB's preconditioned conjugate gradient routine *pcg*. In computing this solution, we utilize the preconditioner developed by Koutis et al. (2011), which is optimized for diagonally dominant design matrices S_{xx} . These column-specific calculations are parallelized across different cores using MATLAB's *parfor* command.

B.3.1 Leaving a Cluster Out

Table 3 applies the leave-cluster-out estimator introduced in Remark 3 to estimate the variance of firm effects with more than two time periods and potential serial correlation. The estimator takes the form $\hat{\theta}_{cluster} = \sum_{i=1}^n y_i \tilde{x}'_i \hat{\beta}_{-c(i)}$ where $\hat{\beta}_{-c(i)}$ is the OLS estimator obtained after leaving out all observations in the cluster to which observation i belongs. A representation of $\hat{\theta}_{cluster}$ that is useful for computation takes the observations in the c -th cluster and collect their outcomes in y_c and their regressors in X_c . The leave-cluster-out estimator is then

$$\hat{\theta}_{cluster} = \hat{\beta}' A \hat{\beta} - \sum_{c=1}^C y'_c B_c (I - P_c)^{-1} (y_c - X_c \hat{\beta}),$$

where C denotes the total number of clusters, $P_c = X_c S_{xx}^{-1} X'_c$, and $B_c = X_c S_{xx}^{-1} A S_{xx}^{-1} X'_c$. Since the entries of P_c and B_c are of the form $P_{i\ell} = x'_i S_{xx}^{-1} x_\ell$ and $B_{i\ell} = x'_i S_{xx}^{-1} A S_{xx}^{-1} x_\ell$, computation can proceed in a similar fashion as described earlier for the leave-one-out estimator.

When defining the cluster as a worker-firm match, Table 3 applies $\hat{\theta}_{cluster}$ to the two-way fixed effects model in (6). When defining the cluster as a worker, the individual effects can not be estimated after leaving a cluster out. Table 3 therefore applies $\hat{\theta}_{cluster}$ after demeaning at the individual level. This transformation removes the individual effects so that the resulting model can be estimated after leaving a cluster out.

B.3.2 Johnson-Lindenstrauss Approximation

When n is on the order of hundreds of millions and k is on the order of tens of millions, the exact algorithm may no longer be tractable. The JLA simplifies computation of P_{ii} considerably by only requiring the solution of p systems of k linear equations. That is, one need only solve for the columns of Z_{JLA} in the system

$$\begin{matrix} S_{xx} & Z_{JLA} \\ k \times k & k \times p \end{matrix} = \begin{matrix} (R_P X)' \\ k \times p \end{matrix},$$

which reduces computation time dramatically when p is small relative to n .

To compute B_{ii} , it is necessary to solve linear systems involving both A_1 and A_2 , leading to $2p$

systems of equations when $A_1 \neq A_2$. However, for variance decompositions like the ones considered in Section 8.2, the same $2p$ systems can be reused for all three variance components, leading to a total of $3p$ systems of equations for the full variance decomposition. This is so because the three variance components use the matrices $A_\psi = A'_f A_f$, $A_{\alpha,\psi} = \frac{1}{2}(A'_d A_f + A'_f A_d)$, and $A_\alpha = A'_d A_d$ where

$$A'_f = \frac{1}{\sqrt{n}} \begin{bmatrix} 0 & 0 & 0 \\ f_1 - \bar{f} & \dots & f_n - \bar{f} \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad A'_d = \frac{1}{\sqrt{n}} \begin{bmatrix} d_1 - \bar{d} & \dots & d_n - \bar{d} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

Based on these insights, Algorithm 3 below takes as inputs X , A_f , A_d , and p , and returns \hat{P}_{ii} and three different \hat{B}_{ii} 's which are ultimately used to construct the corresponding variance component $\hat{\theta}_{JLA}$ as defined in Section 1.2.

Algorithm 3 Johnson-Lindenstrauss Approximation for Two-Way Fixed Effects Models

- 1: **function** JLA(X, A_f, A_d, p)
 - 2: Generate $R_B, R_P \in \mathbb{R}^{p \times n}$, where (R_B, R_P) are composed of mutually independent Rademacher entries
 - 3: Compute $(R_P X)', (R_B A_f)', (R_B A_d)' \in \mathbb{R}^{k \times p}$
 - 4: **for** $\kappa = 1, \dots, p$ **do**
 - 5: Let $r_{\kappa,0}, r_{\kappa,1}, r_{\kappa,2} \in \mathbb{R}^k$ be the κ -th columns of $(R_P X)', (R_B A_f)', (R_B A_d)'$
 - 6: Let $z_{\kappa,\ell} \in \mathbb{R}^k$ be the solution to $S_{xx} z = r_{\kappa,\ell}$ for $\ell = 0, 1, 2$
 - 7: **end for**
 - 8: Construct $Z_\ell = (z_{1,\ell}, \dots, z_{p,\ell}) \in \mathbb{R}^{k \times p}$ for $\ell = 0, 1, 2$
 - 9: Construct $\hat{P}_{ii} = \frac{1}{p} \|Z'_0 x_i\|^2$, $\hat{B}_{ii,\psi} = \frac{1}{p} \|Z'_1 x_i\|^2$, $\hat{B}_{ii,\alpha} = \frac{1}{p} \|Z'_2 x_i\|^2$, $\hat{B}_{ii,\alpha\psi} = \frac{1}{p} (Z'_1 x_i)' (Z'_2 x_i)$ for $i = 1, \dots, n$
 - 10: Return $\{\hat{P}_{ii}, \hat{B}_{ii,\psi}, \hat{B}_{ii,\alpha}, \hat{B}_{ii,\alpha\psi}\}_{i=1}^n$
 - 11: **end function**
-

B.3.3 Performance of the JLA

Figure B.1 evaluates the performance of the Johnson-Lindenstrauss approximation across 4 VWH samples that correspond to different (overlapping) time intervals (2000–2001; 1999–2001; 1998–2001; 1997–2001). The x -axis in Figure B.1 reports the total number of person and firm effects associated with a particular sample.

Figure B.1 shows that the computation time for exact computation of (B_{ii}, P_{ii}) increases rapidly as the number of parameters of the underlying AKM model grow; in the largest dataset considered – which involves more than a million worker and firm effects – exact computation takes about 8 hours. Computation of JLA complete in markedly shorter time: in the largest dataset considered computation time is less than 5 minutes when $p = 500$ and slightly over 6 minutes when $p = 2500$.

Notably, the JLA delivers estimates of the variance of firm effects almost identical to those computed via the exact method, with the quality of the approximation increasing for larger p . For instance, in the largest dataset, the exact estimate of variance of firm effects is 0.028883. By comparison, the JLA estimate equals 0.028765 when $p = 500$ and 0.0289022 when $p = 2500$.

In summary: for a sample with more than a million worker and firm effects, the JLA cuts computation time by a factor of 100 while introducing an approximation error of roughly 10^{-4} .

B.3.4 Scaling to Very Large Datasets


We now study how the JLA scales to much larger datasets of the dimension considered by [Card et al. \(2013\)](#) who fit models involving tens of millions of worker and firm effects to German social security records. To study the computational burden of a model of this scale, we rely on a synthetic dataset constructed from our original leave-one-out sample analyzed in Column 1 of Table 2, i.e., the pooled Veneto sample comprised of wage observations from the years 1999 and 2001. We scale the data by creating replicas of this base sample. To connect the replicas, we draw at random 10% of the movers and randomly exchange their period 1 firm assignments across replicas. By construction, this permutation maintains each (replicated) firm’s size while ensuring leave-one-out connectedness of the resulting network.

Wage observations are drawn from a variant of the DGP described in Section 8.7 adapted to the levels formulation of the model. Specifically, each worker’s wage is the sum of a rescaled person effect, a rescaled firm effect, and an error drawn independently in each period from a normal with variance $\frac{1}{2} \exp(\hat{\alpha}_0 + \hat{\alpha}_1 B_{gg} + \hat{\alpha}_2 P_{gg} + \hat{\alpha}_3 \ln L_{g2} + \hat{\alpha}_4 \ln L_{g1})$. As highlighted by Figure B.1, computing the exact estimator in these datasets would be extremely costly. Drawing from a stable DGP allows us to instead benchmark the JLA estimator against the true value of the variance of firm effects.


Figure B.2 displays the results. When setting $p = 250$, the JLA delivers a variance of firm effects remarkably close to the true variance of firm effects defined by our DGP. As expected, the distance between our approximation and the true variance component decreases with the sample size for a fixed p . Remarkably, we are able to compute the AKM variance decomposition in a dataset with approximately 15 million person and year effects in only 35 minutes. Increasing the number of simulated draws in the JLA to $p = 500$ delivers estimates of the variance of firm effects nearly indistinguishable from the true value. This is achieved in approximately one hour in the largest simulated dataset considered. The results of this exercise strongly suggest the leave-out estimator can be scaled to extremely large datasets involving the universe of administrative wage records in large countries such as Germany or the United States.

Figure B.1: Performance of the JLA Algorithm

(a) Computation Time




(b) Quality of the Approximation




Note: Both panels consider 4 different samples of increasing length. The four samples contain data from the years 2000–2001, 1999–2001, 1998–2001, and 1997–2001, respectively. The x -axis reports the number of person and firm effects in each sample. Panel (a) shows the time to compute the KSS estimate when relying on either exact computation of $\{B_{ii}, P_{ii}\}_{i=1}^n$ or the Johnson-Lindenstrauss approximation (JLA) of these numbers using a p of either 500 or 2500. Panel (b) shows the resulting estimates and the plug-in estimate. Computations performed on a 32 core machine with 256 GB of dedicated memory. Source: VWH dataset.

Figure B.2: Scaling to Very Large Datasets

(a) Computation Time



(b) Quality of the Approximation



Note: Both panels consider synthetic datasets created from the pooled Veneto data in column 1 of Table 2 with $T = 2$. It considers $\{1, 5, 10, 15, 20\}$ replicas of this sample while generating random links across replicas such that firm size and T are kept fixed. Outcomes are generated from a DGP of the sort considered in Table 6. The x -axis reports the number of person and firm effects in each sample. Panel (a) shows the time to compute the Johnson-Lindenstrauss approximation $\hat{\theta}_{JLA}$ using a p of either 250 or 500. Panel (b) shows the resulting estimates, the plug-in estimate, and the true value of the variance of firm effects for the DGP. Computations performed on a 32 core machine with 256 GB of dedicated memory. Source: VWH dataset.

B.4 Split Sample Estimators

Sections 4.2 and 5.2 proposed standard error estimators predicated on being able to construct independent split sample estimators $\widehat{x'_i\beta_{-i,1}}$ and $\widehat{x'_i\beta_{-i,2}}$. This section describes an algorithm for construction of these split sample estimators in the two-way fixed effects model of Example 4. We restrict attention to the case with $T_g = 2$ and consider the model in first differences: $\Delta y_g = \Delta f'_g\psi + \Delta \varepsilon_g$ for $g = 1, \dots, N$. When worker g moves from firm j to j' , we can estimate $\Delta f'_g\psi = \psi_{j'} - \psi_j$ without bias using OLS on any sub-sample where firms j and j' are connected, i.e., on any sample where there exist a path between firm j and j' . To construct two disjoint sub-samples where firms j and j' are connected we therefore use an algorithm to find disjoint paths between these firms and distribute them into two sub-samples which will be denoted \mathcal{S}_1 and \mathcal{S}_2 . Because it can be computationally prohibitive to characterize all possible paths, we use a version of Dijkstra's algorithm to find many short paths.¹⁰

Our algorithm is based on a network where firms are vertices and two firms are connected by an edge if one or more workers moved between them. This view of the network is the same as the one taken in Section 7, but different from the one used in Sections B.1 and B.2 where both firms and workers were viewed as vertices. We use the adjacency matrix \mathcal{A} to characterize the network in this section. To build the sub-samples \mathcal{S}_1 and \mathcal{S}_2 , the algorithm successively drops workers from the network, so $\mathcal{A}_{-\mathcal{S}}$ will denote the adjacency matrix after dropping all workers in the set \mathcal{S} .

Given a network characterized by \mathcal{A} and two connected firms j and j' in the network, we let $\dot{P}_{jj'}(\mathcal{A})$ denote the shortest path between them.¹¹ If j and j' are not connected $\dot{P}_{jj'}(\mathcal{A})$ is empty. Each edge in the path $\dot{P}_{jj'}(\mathcal{A})$ may have more than one worker associated with it. For each edge in $\dot{P}_{jj'}(\mathcal{A})$ the first step of the algorithm picks at random a single worker associated with that edge and places them in \mathcal{S}_1 , while later steps place all workers associated with the shortest path in one of \mathcal{S}_1 and \mathcal{S}_2 . This special first step ensures that the algorithm finds two independent unbiased estimators of $\Delta f'_g\psi$ whenever the network \mathcal{A} is leave-two-out connected.

For a given worker g with firm assignments $j = j(g, 1), j' = j(g, 2)$ and a leave-two-out connected network \mathcal{A} the algorithm returns the $\{P_{g\ell,1}, P_{g\ell,2}\}_{\ell=1}^N$ introduced in Section 4.2. Specifically, $\widehat{\Delta f'_g\psi_{-g,1}} = \sum_{\ell=1}^N P_{g\ell,1}\Delta y_\ell$ and $\widehat{\Delta f'_g\psi_{-g,2}} = \sum_{\ell=1}^N P_{g\ell,2}\Delta y_\ell$ are independent unbiased estimators of

¹⁰The algorithm presented below keeps running until it cannot find any additional paths. In our empirical implementation we stop the algorithm when it fails to find any new paths or as soon as one of the two sub-samples reach a size of at least 100 workers. We found that increasing this cap on the sub-sample size has virtually no effect on the estimated confidence intervals, but tends to increase computation time substantially.

¹¹Many statistical software packages provide functions that can find shortest paths. In R they are available in the *igraph* package while in MATLAB a package that builds on the work of Yen (1971) is available at <https://www.mathworks.com/matlabcentral/fileexchange/35397-k-shortest-paths-in-a-graph-represented-by-a-sparse-matrix-yen-s-algorithm?focused=3779015&tab=function>.

$\Delta f'_g \psi$ that are also independent of Δy_g . If \mathcal{A} is only leave-one-out connected then the algorithm may only find one path connecting j and j' . When this happens the algorithm sets $P_{g\ell,2} = 0$ for all ℓ as required in the formulation of the conservative standard errors proposed in Appendix C.5.1.

Algorithm 4 Split Sample Estimator for Inference

- 1: **function** SPLITSAMPLEESTIMATOR(g, j, j', \mathcal{A}) $\triangleright \mathcal{A} \equiv$ Leave-one-out connected network
 - 2: Let $\mathcal{S}_1 = \emptyset$ and $\mathcal{S}_2 = \emptyset$
 - 3: For each edge in $\dot{P}_{jj'}(\mathcal{A}_{-g})$, pick at random one worker from \mathcal{A}_{-g} who is associated with that edge and add that worker to \mathcal{S}_1
 - 4: Add to \mathcal{S}_2 all workers from $\mathcal{A}_{-\{g, \mathcal{S}_1\}}$ who are associated with an edge in $\dot{P}_{jj'}(\mathcal{A}_{-\{g, \mathcal{S}_1\}})$
 - 5: Add to \mathcal{S}_1 all workers from $\mathcal{A}_{-\{g, \mathcal{S}_1, \mathcal{S}_2\}}$ who are associated with an edge in $\dot{P}_{jj'}(\mathcal{A}_{-g})$
 - 6: Let $stop = 1\{\dot{P}_{jj'}(\mathcal{A}_{-\{g, \mathcal{S}_1, \mathcal{S}_2\}}) = \emptyset\}$ and $s = 1$
 - 7: **while** $stop < 1$ **do**
 - 8: Add to \mathcal{S}_s all workers from $\mathcal{A}_{-\{g, \mathcal{S}_1, \mathcal{S}_2\}}$ who are associated with an edge in $\dot{P}_{jj'}(\mathcal{A}_{-\{g, \mathcal{S}_1, \mathcal{S}_2\}})$
 - 9: Let $stop = 1\{\dot{P}_{jj'}(\mathcal{A}_{-\{g, \mathcal{S}_1, \mathcal{S}_2\}}) = \emptyset\}$ and update s to $1 + 1\{s = 1\}$
 - 10: **end while**
 - 11: For $s = 1, 2$ and $\ell = 1, \dots, N$, let $P_{g\ell, s} = 1\{\ell \in \mathcal{S}_s\} \Delta f'_\ell (\sum_{m \in \mathcal{S}_s} \Delta f_m \Delta f'_m)^\dagger \Delta f_g$
 - 12: Return $\{P_{g\ell, 1}, P_{g\ell, 2}\}_{\ell=1}^N$
 - 13: **end function**
-

In line 5, all workers associated with the shortest path in line 3 are added to \mathcal{S}_1 if they were not added to \mathcal{S}_2 in line 4. This step ensures that all workers associated with $\dot{P}_{jj'}(\mathcal{A}_{-g})$ are used in the predictions. In line 11, $P_{g\ell, s}$ is constructed as the weight observation ℓ receives in the prediction $\Delta f'_g \hat{\psi}_s$ where $\hat{\psi}_s$ is the OLS estimator of ψ based on the sub-sample \mathcal{S}_s .

B.5 Test of Equal Firm Effects

This section describes computation and interpretation of the test of the hypothesis that firm effects for “younger” workers are equal to firm effects for the “older” workers which applies Remark 6 of the main text.

The hypothesis of interest corresponds to a restricted and unrestricted model which when written in matrix notation are

$$\Delta y = \Delta F \psi + \Delta \varepsilon \tag{9}$$

$$\Delta y = \Delta F_O \psi^O + \Delta F_Y \psi^Y + \Delta F_3 \psi_3 + \Delta \varepsilon = X \beta + \Delta \varepsilon \tag{10}$$

where Δy and ΔF collects the first differences Δy_g and Δf_g across g . ΔF_O represents ΔF for “doubly connected” firms present in each age group’s leave-one-out connected set interacted with a dummy for whether the worker is “old”; ΔF_Y represents ΔF for doubly connected firms interacted with a dummy for young; ΔF_3 represents ΔF for firms that are associated with either younger movers or older movers but not both. Finally, we let $X = (\Delta F_O, \Delta F_Y, \Delta F_3)$, $\beta = (\psi^{O'}, \psi^{Y'}, \psi_3')$, and $\psi = (\psi^{O'}, \psi_3')$.

The hypothesis in question is $\psi^O - \psi^Y = 0$ or equivalently $R\beta = 0$ for $R = [I_r, -I_r, 0]$ and $r = |\mathcal{J}| = \dim(\psi^O)$. Thus we can create the numerator of our test statistic by applying Remark 6 to (10) yielding

$$\hat{\theta} = \hat{\beta}' A \hat{\beta} - \sum_{g=1}^N B_{gg} \hat{\sigma}_g^2 \quad (11)$$

where $A = \frac{1}{r} R' (R S_{xx}^{-1} R')^{-1} R$; B_{gg} and $\hat{\sigma}_g^2$ are defined as in Section 1.

Two insights help to simplify computation. First, since $\Delta F_O' \Delta F_Y = 0$, $\Delta F_O' \Delta F_3 = 0$ and $\Delta F_Y' \Delta F_3 = 0$, we can estimate equation (10) via two separate regressions, one on the leave-one-out connected set for younger workers and the other on the leave-one-out connected set for older workers. We normalize the firm effects so that the same firm is dropped in both leave-one-out samples.

Second, we note that $\hat{\beta}' A \hat{\beta} = y' B y$ where

$$B = X S_{xx}^{-1} A S_{xx}^{-1} X' = \frac{P_X - P_{\Delta F}}{r}, \quad (12)$$

$P_X = X S_{xx}^{-1} X'$, and $P_{\Delta F} = \Delta F (\Delta F' \Delta F)^{-1} \Delta F'$. Equation (12) therefore implies that B_{ii} in (11) is simply a scaled difference between two statistical leverages: the first one obtained in the unrestricted model (10), say $P_{X,gg}$, and the other on the restricted model of (9), say $P_{\Delta F,gg}$. Section B.3 describes how to efficiently compute these statistical leverages. To conduct inference on the quadratic form in (11) we apply the routine described in Section 4.2.

Appendix C Proofs

This Appendix contains all technical details and proofs that were left out of the paper. The material is primarily presented in the order it appears in the paper and under the same headings.

C.1 Unbiased Estimation of Variance Components

C.1.1 Estimator

Lemma C.1. *It follows from the Sherman-Morrison-Woodbury formula that the two representations of $\hat{\theta}$ given in (1) and (2) are numerically identical, i.e., that $\hat{\beta}' A \hat{\beta} - \sum_{i=1}^n B_{ii} \hat{\sigma}_i^2 = \sum_{i=1}^n y_i \tilde{x}'_i \hat{\beta}_{-i}$ whenever S_{xx} has full rank and $\max_i P_{ii} < 1$.*

Proof. The Sherman-Morrison-Woodbury formula states that if S_{xx} has full rank and $P_{ii} < 1$, then

$$S_{xx}^{-1} + \frac{S_{xx}^{-1} x_i x_i' S_{xx}^{-1}}{1 - x_i' S_{xx}^{-1} x_i} = (S_{xx} - x_i x_i')^{-1}.$$

Furthermore, we have that $\tilde{x}'_i S_{xx}^{-1} x_i = x_i S_{xx}^{-1} A S_{xx}^{-1} x_i = B_{ii}$ so

$$\begin{aligned} y_i \tilde{x}'_i \hat{\beta}_{-i} &= y_i \tilde{x}'_i (S_{xx} - x_i x_i')^{-1} \sum_{\ell \neq i} x_{\ell} y_{\ell} = y_i \tilde{x}'_i S_{xx}^{-1} \sum_{\ell \neq i} x_{\ell} y_{\ell} + \frac{y_i \tilde{x}'_i S_{xx}^{-1} x_i x_i' S_{xx}^{-1}}{1 - x_i' S_{xx}^{-1} x_i} \sum_{\ell \neq i} x_{\ell} y_{\ell} \\ &= y_i \tilde{x}'_i \hat{\beta} - B_{ii} y_i^2 + y_i B_{ii} x_i' \underbrace{\frac{S_{xx}^{-1}}{1 - x_i' S_{xx}^{-1} x_i} \sum_{\ell \neq i} x_{\ell} y_{\ell}}_{= x_i' \hat{\beta}_{-i}} = y_i \tilde{x}'_i \hat{\beta} - B_{ii} y_i (y_i - x_i' \hat{\beta}_{-i}) \end{aligned}$$

where the last expression equals $y_i \tilde{x}'_i \hat{\beta} - B_{ii} \hat{\sigma}_i^2$. This finishes the proof since $\hat{\beta}' A \hat{\beta} = \sum_{i=1}^n y_i \tilde{x}'_i \hat{\beta}$. In the above the Sherman-Morrison-Woodbury formula was also used to establish that

$$x_i' \hat{\beta}_{-i} = x_i' (S_{xx} - x_i x_i')^{-1} \sum_{\ell \neq i} x_{\ell} y_{\ell} = x_i' \frac{S_{xx}^{-1}}{1 - x_i' S_{xx}^{-1} x_i} \sum_{\ell \neq i} x_{\ell} y_{\ell},$$

and from this it follows that $y_i - x_i' \hat{\beta}_{-i} = \frac{y_i - x_i' \hat{\beta}}{1 - P_{ii}}$ as claimed in the paper. \square

C.1.2 Large Scale Computation

All discussions of the computational aspects are collected in Appendix B.

C.1.3 Relation To Existing Approaches

Next we verify that the bias of $\hat{\theta}_{\text{HO}}$ is a function of the covariation between σ_i^2 and (B_{ii}, P_{ii}) .

Lemma C.2. *The bias of $\hat{\theta}_{HO}$ is $\sigma_{nB_{ii},\sigma_i^2} + S_B \frac{n}{n-k} \sigma_{P_{ii},\sigma_i^2}$ where*

$$\sigma_{nB_{ii},\sigma_i^2} = \sum_{i=1}^n B_{ii}(\sigma_i^2 - \bar{\sigma}^2), \quad \bar{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \sigma_i^2, \quad S_B = \sum_{i=1}^n B_{ii}, \quad \sigma_{P_{ii},\sigma_i^2} = \frac{1}{n} \sum_{i=1}^n P_{ii}(\sigma_i^2 - \bar{\sigma}^2).$$

Proof. Since $\hat{\sigma}^2 = \frac{1}{n-k} \sum_{i=1}^n (y_i - x_i' \hat{\beta})^2 = \frac{1}{n-k} \sum_{i=1}^n \sum_{\ell=1}^n M_{i\ell} \varepsilon_i \varepsilon_\ell$ we get that

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{HO}] - \theta &= \sum_{i=1}^n B_{ii} \sigma_i^2 - \left(\sum_{i=1}^n B_{ii} \right) \frac{1}{n-k} \sum_{i=1}^n M_{ii} \sigma_i^2 \\ &= \sum_{i=1}^n B_{ii} (\sigma_i^2 - \bar{\sigma}^2) - S_B \frac{1}{n-k} \sum_{i=1}^n M_{ii} (\sigma_i^2 - \bar{\sigma}^2) \\ &= \sigma_{nB_{ii},\sigma_i^2} + S_B \frac{n}{n-k} \sigma_{P_{ii},\sigma_i^2}. \end{aligned} \quad \square$$

Comparison to Jackknife Estimators

This subsection compares the leave-out estimator $\hat{\theta}$ to estimators predicated on jackknife bias corrections. We start by introducing some of the high-level assumptions that are typically used to motivate jackknife estimators. We then consider some variants of Examples 2 and 3 where these high-level conditions fail to hold and establish that the jackknife estimators have first order biases while the leave-out estimator retains consistency.

High-level Conditions Jackknife bias corrections are typically motivated by the high-level assumption that the bias of a plug-in estimator $\hat{\theta}_{PI}$ shrinks with the sample size in a known way and that the bias of $\frac{1}{n} \sum_{i=1}^n \hat{\theta}_{PI,-i}$ depends on sample size in an identical way, i.e.,

$$\mathbb{E}[\hat{\theta}_{PI}] = \theta + \frac{D_1}{n} + \frac{D_2}{n^2}, \quad \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \hat{\theta}_{PI,-i} \right] = \theta + \frac{D_1}{n-1} + \frac{D_2}{(n-1)^2} \quad \text{for some } D_1, D_2. \quad (13)$$

Under (13), the jackknife estimator $\hat{\theta}_{JK} = n\hat{\theta}_{PI} - \frac{n-1}{n} \sum_{i=1}^n \hat{\theta}_{PI,-i}$ has a bias of $-\frac{D_2}{n(n-1)}$.

For some long panel settings the bias in $\hat{\theta}_{PI}$ is shrinking in the number of time periods T such that

$$\mathbb{E}[\hat{\theta}_{PI}] = \theta + \frac{\dot{D}_1}{T} + \frac{\dot{D}_2}{T^2} \quad \text{for some } \dot{D}_1, \dot{D}_2.$$

In such settings, it may be that the biases of $\frac{1}{T} \sum_{t=1}^T \hat{\theta}_{PI,-t}$ and $\frac{1}{2}(\hat{\theta}_{PI,1} + \hat{\theta}_{PI,2})$ depend on T in an identical way, i.e.,

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \hat{\theta}_{PI,-t} \right] = \theta + \frac{\dot{D}_1}{T-1} + \frac{\dot{D}_2}{(T-1)^2} \quad \text{and} \quad \mathbb{E} \left[\frac{1}{2}(\hat{\theta}_{PI,1} + \hat{\theta}_{PI,2}) \right] = \theta + \frac{2\dot{D}_1}{T} + \frac{4\dot{D}_2}{T^2}.$$

From here it follows that the panel jackknife estimator $\hat{\theta}_{\text{PJk}} = T\hat{\theta}_{\text{PI}} - \frac{T-1}{T} \sum_{t=1}^T \hat{\theta}_{\text{PI},-t}$ has a bias of $-\frac{\mathring{D}_2}{T(T-1)}$ and that the split panel jackknife estimator $\hat{\theta}_{\text{SPJk}} = 2\hat{\theta}_{\text{PI}} - \frac{1}{2}(\hat{\theta}_{\text{PI},1} + \hat{\theta}_{\text{PI},2})$ has a bias of $-\frac{2\mathring{D}_2}{T^2}$, both of which shrink faster to zero than $\frac{\mathring{D}_1}{T}$ if $T \rightarrow \infty$. Typical sufficient conditions for bias-representations of this kind to hold (to second order) are that (i) $T \rightarrow \infty$, (ii) the design is stationary over time, and (iii) that $\hat{\theta}_{\text{PI}}$ is asymptotically linear (see, e.g., [Hahn and Newey, 2004](#); [Dhaene and Jochmans, 2015](#)). Below we illustrate that jackknife corrections can be inconsistent in Examples 2 and 3 when (i) and/or (ii) do not hold. Finally we note that $\hat{\theta}_{\text{PI}}$ (a quadratic function) need not be asymptotically linear as is evident from the non-normal asymptotic distribution of $\hat{\theta}$ derived in Theorem 3 of this paper.

Examples of Jackknife Failure

Example 2 (Special case). Consider the model

$$y_{gt} = \alpha_g + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T \geq 2),$$

where $\sigma_{gt}^2 = \sigma^2$ and suppose the parameter of interest is $\theta = \frac{1}{N} \sum_{g=1}^N \alpha_g^2$. For T even, we have the following bias calculations:

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{\text{PI}}] &= \theta + \frac{\sigma^2}{T}, & \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n \hat{\theta}_{\text{PI},-i}\right] &= \theta + \frac{\sigma^2}{T} + \frac{\sigma^2}{n(T-1)}, \\ \mathbb{E}\left[\frac{1}{T} \sum_{t=1}^T \hat{\theta}_{\text{PI},-t}\right] &= \theta + \frac{\sigma^2}{T-1}, & \mathbb{E}\left[\frac{1}{2}(\hat{\theta}_{\text{PI},1} + \hat{\theta}_{\text{PI},2})\right] &= \theta + \frac{2\sigma^2}{T}. \end{aligned}$$

The jackknife estimator $\hat{\theta}_{\text{JK}}$ has a first order bias of $-\frac{\sigma^2}{T(T-1)}$, which when $T = 2$ is as large as that of $\hat{\theta}_{\text{PI}}$ but of opposite sign. By contrast, both of the panel jackknife estimators, $\hat{\theta}_{\text{PJk}}$ and the leave-out estimator are exactly unbiased and consistent as $n \rightarrow \infty$ when T is fixed.

This example shows that the jackknife estimator can fail when applied to a setting where the number of regressors is large relative to sample size. Here the number of regressors is N and the sample size is NT , yielding a ratio of $1/T$ and we see that $1/T \rightarrow 0$ is necessary for consistency of $\hat{\theta}_{\text{JK}}$. While the panel jackknife corrections appear to handle the presence of many regressors, this property disappears in the next example which adds the “random coefficients” of Example 3.

Example 3 (Special case). Consider the model

$$y_{gt} = \alpha_g + x_{gt}\delta_g + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T \geq 3)$$

where $\sigma_{gt}^2 = \sigma^2$ and $\theta = \frac{1}{N} \sum_{g=1}^N \delta_g^2$.

An analytically convenient example arises when the regressor design is “balanced” across groups as follows:

$$(x_{g1}, x_{g2}, \dots, x_{gT}) = (x_1, x_2, \dots, x_T),$$

where x_1, x_2, x_3 take distinct values and $\sum_{t=1}^T x_t = 0$. The leave-out estimator is unbiased and consistent for any $T \geq 3$, whereas for even $T \geq 4$ we have the following bias calculations:

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{\text{PI}}] &= \theta + \frac{\sigma^2}{\sum_{t=1}^T x_t^2}, \\ \mathbb{E}\left[\frac{1}{T} \sum_{t=1}^T \hat{\theta}_{\text{PI},-t}\right] &= \theta + \frac{\sigma^2}{T} \sum_{t=1}^T \frac{1}{\sum_{s \neq t} (x_s - \bar{x}_{-t})^2}, \\ \mathbb{E}\left[\frac{1}{2}(\hat{\theta}_{\text{PI},1} + \hat{\theta}_{\text{PI},2})\right] &= \theta + \frac{\sigma^2}{2 \sum_{t=1}^{T/2} (x_t - \bar{x}_1)^2} + \frac{\sigma^2}{2 \sum_{t=T/2+1}^T (x_t - \bar{x}_2)^2}, \end{aligned}$$

where $\bar{x}_{-t} = \frac{1}{T-1} \sum_{s \neq t} x_s$, $\bar{x}_1 = \frac{2}{T} \sum_{t=1}^{T/2} x_t$, and $\bar{x}_2 = \frac{2}{T} \sum_{t=T/2+1}^T x_t$.

The calculations above reveal that non-stationarity in either the level or variability of x_t over time can lead to a negative bias in panel jackknife approaches, e.g.,

$$\mathbb{E}\left[\hat{\theta}_{\text{SPJK}}\right] - \theta \leq \frac{2\sigma^2}{\sum_{t=1}^T x_t^2} - \frac{\sigma^2}{2 \sum_{t=1}^{T/2} x_t^2} - \frac{\sigma^2}{2 \sum_{t=T/2+1}^T x_t^2} \leq 0$$

where the first inequality is strict if $\bar{x}_1 \neq \bar{x}_2$ and the second if $\sum_{t=1}^{T/2} x_t^2 \neq \sum_{t=T/2+1}^T x_t^2$. In fact, the following example

$$(x_1, x_2, \dots, x_T) = (-1, 2, 0, \dots, 0, -1)$$

renders the panel jackknife corrections inconsistent for small or large T :

$$\mathbb{E}[\hat{\theta}_{\text{PJK}}] = \theta - \frac{7/5}{6} \sigma^2 + O\left(\frac{1}{T}\right) \quad \text{and} \quad \mathbb{E}[\hat{\theta}_{\text{SPJK}}] = \theta - \frac{8/5}{6} \sigma^2 + O\left(\frac{1}{T}\right).$$

Inconsistency results here from biases of first order that are negative and larger in magnitude than the original bias of $\hat{\theta}_{\text{PI}}$ (which is $\frac{\sigma^2}{6}$).

Computations For this special case of example 2 we have that $A = \frac{I_N}{N}$ and $S_{xx} = T I_N$ so that $\tilde{A} = \frac{I_N}{NT}$ and $\text{trace}(\tilde{A}^2) = \frac{1}{NT^2} = o(1)$ which implies consistency of $\hat{\theta}$. Similarly we have that the

bias of $\tilde{\theta}$ is

$$\frac{1}{n} \sum_{g=1}^N T_g \mathbb{V}[\hat{\alpha}_g] = \frac{1}{n} \sum_{g=1}^N \sigma^2 = \frac{\sigma^2}{T} \quad \text{where } \hat{\alpha}_g = \frac{1}{T_g} \sum_{t=1}^{T_g} y_{gt}.$$

The same types of calculations lead to the other biases reported in the paper.

For this special case of example 3 we have that $A = \begin{bmatrix} 0 & 0 \\ 0 & \frac{I_N}{N} \end{bmatrix}$ and $S_{xx} = \begin{bmatrix} TI_N & 0 \\ 0 & I_N \sum_{t=1}^T x_t^2 \end{bmatrix}$ which implies that $\text{trace}(\tilde{A}^2) = \frac{1}{N(\sum_{t=1}^T x_t^2)^2} = o(1)$ and therefore consistency of $\hat{\theta}$. Similarly we

have that the bias of $\tilde{\theta}$ is

$$\frac{1}{n} \sum_{g=1}^N T_g \mathbb{V}[\hat{\delta}_g] = \frac{\sigma^2}{\sum_{t=1}^T x_t^2} \quad \text{where } \hat{\delta}_g = \frac{\sum_{t=1}^{T_g} x_t y_{gt}}{\sum_{t=1}^{T_g} x_t^2}.$$

The same types of calculations lead to the other biases reported above. Now for the numerical example $(x_1, x_2, \dots, x_T) = (-1, 2, 0, \dots, 0, -1)$ we have $\sum_{t=1}^T x_t^2 = 6$, $\sum_{t=T/2+1}^T (x_t - \bar{x}_2)^2 = 1 - \frac{2}{T}$, $\sum_{t=1}^{T/2} (x_t - \bar{x}_1)^2 = 2 \sum_{t=1}^{T/2} x_t^2 - T \bar{x}_1^2 = 5 - \frac{2}{T}$, and

$$\sum_{s \neq t} (x_s - \bar{x}_{-t})^2 = \begin{cases} 2 - \frac{4}{T-1} & \text{if } t = 2, \\ 5 - \frac{1}{T-1} & \text{if } t \in \{1, T\}, \\ 6 & \text{otherwise,} \end{cases}$$

Thus

$$\begin{aligned} \mathbb{E}[\hat{\theta}_{\text{PJK}}] - \theta &= \frac{T\sigma^2}{\sum_{t=1}^T x_t^2} - \sigma^2 \frac{(T-1)}{T} \sum_{t=1}^T \frac{1}{\sum_{s \neq t} (x_s - \bar{x}_{-t})^2} \\ &= \sigma^2 \frac{T}{6} - \sigma^2 \frac{T-1}{T} \left(\frac{2}{5 - \frac{1}{T-1}} + \frac{1}{2 - \frac{4}{T-1}} + \frac{T-3}{6} \right) \\ &= \sigma^2 \left(\frac{2}{3} - \frac{4}{6T} - \frac{T-1}{T} \frac{2}{5 - \frac{1}{T-1}} - \frac{T-1}{T} \frac{1}{2 - \frac{4}{T-1}} \right) = -\frac{7}{30} \sigma^2 + O\left(\frac{1}{T}\right) \end{aligned}$$

$$\begin{aligned} \text{and } \mathbb{E}[\hat{\theta}_{\text{SPJK}}] - \theta &= \frac{2\sigma^2}{\sum_{t=1}^T x_t^2} - \frac{\sigma^2}{2 \sum_{t=1}^{T/2} (x_t - \bar{x}_1)^2} + \frac{\sigma^2}{2 \sum_{t=T/2+1}^T (x_t - \bar{x}_2)^2} \\ &= \sigma^2 \left(\frac{1}{3} - \frac{1}{10 - \frac{4}{T}} - \frac{1}{2 - \frac{4}{T}} \right) = -\frac{8}{30} \sigma^2 + O\left(\frac{1}{T}\right). \end{aligned}$$

C.1.4 Finite Sample Properties

Here we provide a restatement and proof of Lemmas 1 and 2 together with a characterization of the finite sample distribution of $\hat{\theta}$ which was excluded from the main text.

Lemma C.3. Recall that $\theta^* = \hat{\beta}' A \hat{\beta} - \sum_{i=1}^n B_{ii} \sigma_i^2$.

1. If $\max_i P_{ii} < 1$, then $\mathbb{E}[\hat{\theta}] = \theta$.
2. Unbiased estimators of $\theta = \beta' A \beta$ exist for all A if and only if $\max_i P_{ii} < 1$.
3. If $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, then $\theta^* = \sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right)$ and $\hat{b} \sim \mathcal{N}(b, \mathbb{V}[\hat{b}])$.
4. If $\max_i P_{ii} < 1$ and $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, then $\hat{\theta} = \sum_{\ell=1}^{r_C} \lambda_\ell(C) \left(\hat{y}_\ell^2 - V_{\ell\ell} \right)$ where $\hat{y} \sim \mathcal{N}(\mu, V)$, $\mu = Q_C' X \beta$, $V = Q_C' \Sigma Q_C$, $C = (C_{i\ell})_{i,\ell}$, $\Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$, and $C = Q_C D_C Q_C'$ is a spectral decomposition of C such that $D_C = \text{diag}(\lambda_1(C), \dots, \lambda_{r_C}(C))$ and r_C is the rank of C .

Proof. First note that $\hat{\beta}' A \hat{\beta} = \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell} y_i y_\ell$ and $\hat{\sigma}_i^2 = y_i (y_i - x_i' \hat{\beta}_{-i}) = y_i M_{ii}^{-1} \sum_{\ell=1}^n M_{i\ell} y_\ell$, so

$$\begin{aligned} \hat{\theta} &= \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell} y_i y_\ell - \sum_{i=1}^n M_{ii}^{-1} M_{i\ell} y_i y_\ell \\ &= \sum_{i=1}^n \sum_{\ell=1}^n \left(B_{i\ell} - 2^{-1} M_{i\ell} \left(B_{ii} M_{ii}^{-1} + B_{\ell\ell} M_{\ell\ell}^{-1} \right) \right) y_i y_\ell = \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell} y_i y_\ell. \end{aligned}$$

The errors are mean zero and uncorrelated across observations, so

$$\mathbb{E}[\hat{\theta}] = \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell} x_i' \beta x_\ell' \beta = \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell} x_i' \beta x_\ell' \beta - \sum_{i=1}^n M_{ii}^{-1} M_{i\ell} x_i' \beta x_\ell' \beta = \theta,$$

since $\sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell} x_i x_\ell' = A$ and $\sum_{\ell=1}^n M_{i\ell} x_\ell = 0$. This shows the first claim of the lemma.

It suffices to show that no unbiased estimator of $\beta' S_{xx} \beta$ exist when $\max_i P_{ii} = 1$. Any potential unbiased estimator must have the representation $y' D y + U$ where $\mathbb{E}[U] = 0$ and $D = (D_{i\ell})_{i,\ell}$ satisfies (i) $D_{ii} = 0$ for all i and (ii) $X' D X = S_{xx}$ for $X = (x_1, \dots, x_n)'$. (ii) implies that D must be $D = I + P \tilde{D} M + M \tilde{D} P + M \tilde{D} M$ for some \tilde{D} where $P = (P_{i\ell})_{i,\ell}$ and $M = (M_{i\ell})_{i,\ell}$. If there exist a i with $P_{ii} = 1$, then $\sum_{\ell=1}^n P_{i\ell}^2 = P_{ii}$ yields $M_{i\ell} = 0$ for all ℓ which implies that D_{ii} must equal 1 to satisfy (ii). However, this makes it impossible to satisfy (i). This shows the second claim.

Recall the spectral decomposition $\tilde{A} = Q D Q'$ and definition of $\hat{b} = Q' S_{xx}^{1/2} \hat{\beta}$ which satisfies that $\hat{b} \sim \mathcal{N}(b, \mathbb{V}[\hat{b}])$ when $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$. We have that $\theta^* = \sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right)$ since

$$\hat{\beta}' A \hat{\beta} = \hat{\beta}' S_{xx}^{1/2} \tilde{A} S_{xx}^{1/2} \hat{\beta} = \hat{b}' D \hat{b} = \sum_{\ell=1}^r \lambda_\ell \hat{b}_\ell^2,$$

and

$$\sum_{i=1}^n B_{ii}\sigma_i^2 = \text{trace}(B\Sigma) = \text{trace}(A\mathbb{V}[\hat{\beta}]) = \text{trace}(D\mathbb{V}[\hat{b}]) = \sum_{\ell=1}^r \lambda_\ell \mathbb{V}[\hat{b}_\ell].$$

where $B = (B_{i\ell})_{i,\ell}$. This shows the third claim.

The matrix C is well-defined as $\max_i P_{ii} < 1$. Define $\hat{y} = Q'_C(y_1, \dots, y_n)'$ which satisfies that $\hat{y} \sim \mathcal{N}(\mu, V)$ when $\varepsilon_i \sim \mathcal{N}(0, \sigma_i^2)$. Furthermore,

$$\hat{\theta} = y'Cy = \hat{y}'D_C\hat{y} = \sum_{\ell=1}^{r_C} \lambda_\ell(C)\hat{y}_\ell^2,$$

and $C_{ii} = 0$ for all i , so that $\sum_\ell \lambda_\ell(C)V_{\ell\ell} = \text{trace}(C\Sigma) = 0$. This shows the last claim. \square

C.1.5 Consistency

The next result provides a restatement and proof of Lemma 3.

Lemma C.4. *If Assumption 1 and one of the following conditions hold, then $\hat{\theta} - \theta \xrightarrow{P} 0$.*

(i) *A is positive semi-definite, $\theta = \beta'A\beta = O(1)$, and $\text{trace}(\tilde{A}^2) = \sum_{\ell=1}^r \lambda_\ell^2 = o(1)$.*

(ii) *$A = \frac{1}{2}(A'_1A_2 + A'_2A_1)$ where $\theta_1 = \beta'A'_1A_1\beta$, $\theta_2 = \beta'A'_2A_2\beta$ satisfy (i).*

Proof. Suppose that A is positive semi-definite. The difference between $\hat{\theta}$ and θ is

$$\hat{\theta} - \theta = 2 \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell}x'_\ell\beta\varepsilon_i + \sum_{i=1}^n \sum_{\ell \neq i} B_{i\ell}\varepsilon_i\varepsilon_\ell + \sum_{i=1}^n B_{ii}(\varepsilon_i^2 - \hat{\sigma}_i^2),$$

and each term has mean zero so we show that their variances are small in large samples. The variance of the first term is

$$4 \sum_{i=1}^n \left(\sum_{\ell=1}^n B_{i\ell}x'_\ell\beta \right)^2 \sigma_i^2 \leq 4 \max_i \sigma_i^2 \beta' X' B^2 X \beta = 4 \max_i \sigma_i^2 \beta' A S_{xx}^{-1} A \beta \leq 4 \max_i \sigma_i^2 \theta \lambda_1 = o(1)$$

where $B = (B_{i\ell})_{i,\ell}$, the last inequality follows from positive semi-definiteness of A , and the last equality follows from $\theta = O(1)$ and $\lambda_1 \leq \text{trace}(\tilde{A}^2)^{1/2} = o(1)$. The variance of the second term is

$$2 \sum_{i=1}^n \sum_{\ell \neq i} B_{i\ell}^2 \sigma_i^2 \sigma_\ell^2 \leq 2 \max_i \sigma_i^4 \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell}^2 = 2 \max_i \sigma_i^4 \text{trace}(\tilde{A}^2) = o(1).$$

Finally, the variance of the third term is

$$\begin{aligned} & \sum_{i=1}^n \left(\sum_{\ell=1}^n M_{i\ell}^{-1} B_{\ell\ell} M_{i\ell} x'_{\ell} \beta \right)^2 \sigma_i^2 + 2 \sum_{i=1}^n \sum_{\ell \neq i} M_{ii}^{-2} B_{ii}^2 M_{i\ell}^2 \sigma_i^2 \sigma_{\ell}^2 \\ & \leq \frac{1}{c^2} \max_i \sigma_i^2 \max_i (x'_i \beta)^2 \sum_{i=1}^n B_{ii}^2 + \frac{2}{c} \max_i \sigma_i^4 \sum_{i=1}^n B_{ii}^2 = o(1) \end{aligned}$$

where $\min_i M_{ii} \geq c > 0$ and $\sum_{i=1}^n B_{ii}^2 \leq \text{trace}(\tilde{A}^2) = o(1)$. This shows the first claim of the lemma.

When A is non-definite, we write $A = \frac{1}{2} (A'_1 A_2 + A'_2 A_1)$ and note that

$$\beta' A S_{xx}^{-1} A \beta \leq \frac{1}{2} \left(\Theta_1 \lambda_{\max}(\tilde{A}_2) + \Theta_2 \lambda_{\max}(\tilde{A}_1) \right) \quad \text{and} \quad \text{trace}(\tilde{A}^2) \leq \text{trace}(\tilde{A}_1^2)^{1/2} \text{trace}(\tilde{A}_2^2)^{1/2}$$

where $\tilde{A}_{\ell} = S_{xx}^{-1/2} A'_{\ell} A_{\ell} S_{xx}^{-1/2}$ for $\ell = 1, 2$ and $\lambda_{\max}(\tilde{A}_2)$ is the largest eigenvalue of \tilde{A}_2 . Thus consistency of $\hat{\theta}$ follows from $\Theta_1 = O(1)$, $\Theta_2 = O(1)$, $\text{trace}(\tilde{A}_1^2) = o(1)$, and $\text{trace}(\tilde{A}_2^2) = o(1)$. \square

The next result provides a restatement and proof of Lemma 4.

Lemma C.5. *If Assumption 1, $n/p^4 = o(1)$, $\mathbb{V}[\hat{\theta}]^{-1} = O(n)$, and one of the following conditions hold, then $\mathbb{V}[\hat{\theta}]^{-1/2} (\hat{\theta}_{JLA} - \hat{\theta} - B_p) = o_p(1)$ where $|B_p| \leq \frac{1}{p} \sum_{i=1}^n P_{ii}^2 |B_{ii}| \sigma_i^2$.*

(i) A is positive semi-definite and $\mathbb{E}[\hat{\beta}' A \hat{\beta}] - \theta = O(1)$.

(ii) $A = \frac{1}{2} (A'_1 A_2 + A'_2 A_1)$ where $\theta_1 = \beta' A'_1 A_1 \beta$, $\theta_2 = \beta' A'_2 A_2 \beta$ satisfy (i) and $\frac{\mathbb{V}[\hat{\theta}_1] \mathbb{V}[\hat{\theta}_2]}{n \mathbb{V}[\hat{\theta}]^2} = O(1)$.

Proof. Define $B_p = \frac{1}{p} \sum_{i=1}^n B_{ii} \sigma_i^2 \frac{P_{i\ell}^4 - P_{ii}^2 (1 - P_{ii})^2}{(1 - P_{ii})^2}$. Letting $(\hat{\theta}_{JLA} - \hat{\theta})_2$ be a second order approximation of $\hat{\theta}_{JLA} - \hat{\theta}$, we first show that $\mathbb{E}[(\hat{\theta}_{JLA} - \hat{\theta})_2] = B_p$ and $\frac{\mathbb{V}[(\hat{\theta}_{JLA} - \hat{\theta})_2]}{\mathbb{V}[\hat{\theta}]} = O(\frac{1}{p})$. Then we finish the proof of the first claim by showing that the approximation error is ignorable. The bias bound follows immediately from the equality $\sum_{\ell \neq i} P_{i\ell}^2 = P_{ii}(1 - P_{ii})$ which leads to $0 \leq \sum_{\ell \neq i} P_{i\ell}^4 \leq P_{ii}^2 (1 - P_{ii})^2$.

We have $\hat{\theta}_{JLA} - \hat{\theta} = (\hat{\theta}_{JLA} - \hat{\theta})_2 + AE_2$ where

$$(\hat{\theta}_{JLA} - \hat{\theta})_2 = \sum_{i=1}^n \hat{\sigma}_i^2 \left(B_{ii} - \hat{B}_{ii} - \hat{B}_{ii} \hat{a}_i - \hat{B}_{ii} \left(\hat{a}_i^2 - \frac{1}{p} \frac{3P_{ii}^3 + P_{ii}^2}{1 - P_{ii}} \right) \right)$$

for $\hat{a}_i = \frac{\hat{P}_{ii} - P_{ii}}{1 - P_{ii}}$ and approximation error

$$AE_2 = \sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \left(\frac{1}{p} \frac{3\hat{P}_{ii}^2 + \hat{P}_{ii}^2 - (3P_{ii}^2 + P_{ii}^2)(1 + \hat{a}_i)^2}{(1 + \hat{a}_i)^2 (1 - P_{ii})} - \frac{\hat{a}_i^3}{1 + \hat{a}_i} \right).$$

For the mean calculation involving $(\hat{\theta}_{JLA} - \hat{\theta})_2$ we use independence between \hat{B}_{ii} , \hat{P}_{ii} , and $\hat{\sigma}_i^2$, unbiasedness of \hat{B}_{ii} , \hat{P}_{ii} , and $\hat{\sigma}_i^2$, and the variance formula

$$\mathbb{V}[\hat{a}_i] = \frac{2P_{ii}^2 - \sum_{\ell=1}^n P_{i\ell}^4}{p(1-P_{ii})^2} = \frac{1}{p} \frac{3P_{ii}^3 + P_{ii}^2}{1-P_{ii}} + \frac{P_{ii}^2(1-P_{ii})^2 - 2\sum_{\ell \neq i}^n P_{i\ell}^4}{p(1-P_{ii})^2}.$$

Taken together this implies that

$$\mathbb{E} \left[(\hat{\theta}_{JLA} - \hat{\theta})_2 \right] = - \sum_{i=1}^n B_{ii} \sigma_i^2 \left(\mathbb{V}[\hat{a}_i] - \frac{1}{p} \frac{3P_{ii}^3 + P_{ii}^2}{1-P_{ii}} \right) = B_p.$$

For the variance calculation we proceed term by term. We have for $y = (y_1, \dots, y_n)'$ that

$$\begin{aligned} \mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 (B_{ii} - \hat{B}_{ii}) \right] &= \mathbb{E} \left[\mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \mid y \right] \right] \leq \frac{2}{p} \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell}^2 \mathbb{E} \left[\hat{\sigma}_i^2 \hat{\sigma}_\ell^2 \right] = O \left(\frac{\text{trace}(\tilde{A}^2)}{p} \right), \\ \mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \hat{a}_i \right] &= \mathbb{E} \left[\mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \hat{a}_i \mid y, R_B \right] \right] \leq \frac{2}{p} \sum_{i=1}^n \sum_{\ell=1}^n P_{i\ell}^2 \frac{\mathbb{E}[\hat{B}_{ii} \hat{B}_{\ell\ell}] \mathbb{E}[\hat{\sigma}_i^2 \hat{\sigma}_\ell^2]}{(1-P_{ii})(1-P_{\ell\ell})} \\ &= O \left(\frac{\text{trace}(\tilde{A}^2)}{p} + \frac{\text{trace}(\tilde{A}_1^2)^{1/2} \text{trace}(\tilde{A}_2^2)^{1/2}}{p^2} \right) \end{aligned}$$

where $\tilde{A}_\ell = S_{xx}^{-1/2} A'_\ell A_\ell S_{xx}^{-1/2}$ for $\ell = 1, 2$,

$$\begin{aligned} \mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} (\hat{a}_i^2 - \mathbb{V}[\hat{a}_i]) \right] &= \sum_{i=1}^n \sum_{\ell=1}^n \mathbb{E} \left[\hat{B}_{ii} \hat{B}_{\ell\ell} \right] \mathbb{E} \left[\hat{\sigma}_i^2 \hat{\sigma}_\ell^2 \right] \text{Cov}(\hat{a}_i^2, \hat{a}_\ell^2) \\ &= O \left(\frac{\text{trace}(\tilde{A}^2)}{p^2} + \frac{\text{trace}(\tilde{A}_1^2)^{1/2} \text{trace}(\tilde{A}_2^2)^{1/2}}{p^3} \right) \\ \mathbb{V} \left[\sum_{i=1}^n \hat{\sigma}_i^2 (\hat{B}_{ii} - B_{ii}) \frac{2\sum_{\ell \neq i}^n P_{i\ell}^4 - P_{ii}^2(1-P_{ii})^2}{p(1-P_{ii})^2} \right] &= O \left(\frac{\text{trace}(\tilde{A}^2)}{p^3} \right) \\ \mathbb{V} \left[\sum_{i=1}^n B_{ii} (\hat{\sigma}_i^2 - \sigma_i^2) \frac{2\sum_{\ell \neq i}^n P_{i\ell}^4 - P_{ii}^2(1-P_{ii})^2}{p(1-P_{ii})^2} \right] &= O \left(\frac{\mathbb{V}[\hat{\theta}]}{p^2} \right) \end{aligned}$$

From this it follows that $\mathbb{V}[\hat{\theta}]^{-1/2} \left((\hat{\theta}_{JLA} - \hat{\theta})_2 - B_p \right) = o_p(1)$ since $\text{trace}(\tilde{A}^2) = O(\mathbb{V}[\hat{\theta}])$ and $\frac{\mathbb{V}[\hat{\theta}_1] \mathbb{V}[\hat{\theta}_2]}{p^4 \mathbb{V}[\hat{\theta}]^2} = o(1)$.

We now treat the approximation error while utilizing that $\mathbb{E}[\hat{a}_i^3] = O\left(\frac{1}{p}\right)$, $\mathbb{E}[\hat{a}_i^4] = O\left(\frac{1}{p^2}\right)$, and $\max_i |\hat{a}_i| = o_p(\log(n)/\sqrt{p})$ which follows from (Achlioptas, 2003, Theorem 1.1 and its proof).

Proceeding term by term, we list the conclusions

$$\begin{aligned} \sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \hat{a}_i^3 + \sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \hat{a}_i^4 &= O_p \left(\frac{\mathbb{E}[\hat{\Theta}_{1,\text{PI}} - \Theta_1] + \mathbb{E}[\hat{\Theta}_{2,\text{PI}} - \Theta_2]}{p^2} \right) \\ \sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \frac{\hat{a}_i^5}{1 + \hat{a}_i} &= O_p \left(\frac{\log(n)}{\sqrt{p}} \frac{\mathbb{E}[\hat{\Theta}_{1,\text{PI}} - \Theta_1] + \mathbb{E}[\hat{\Theta}_{2,\text{PI}} - \Theta_2]}{p^2} \right) \\ \frac{1}{p} \sum_{i=1}^n \hat{\sigma}_i^2 \hat{B}_{ii} \frac{3\hat{P}_{ii}^2 + \hat{P}_{ii}^2 - (3P_{ii}^2 + P_{ii}^2)(1 + \hat{a}_i)^2}{(1 + \hat{a}_i)^2(1 - P_{ii})} &= O_p \left(\left(1 + \frac{\log(n)}{\sqrt{p}}\right) \frac{\mathbb{E}[\hat{\Theta}_{1,\text{PI}} - \Theta_1] + \mathbb{E}[\hat{\Theta}_{2,\text{PI}} - \Theta_2]}{p^2} \right) \end{aligned}$$

which finishes the proof. \square

C.2 Examples

All mathematical discussions of the examples are collected in Appendix C.7.

C.3 Quadratic Forms of Fixed Rank

The next result provides a restatement and proof of Theorem 1.

Theorem C.1. *If Assumption 1 holds, r is fixed, and $\max_i w_i' w_i = o(1)$, then*

1. $\mathbb{V}[\hat{b}]^{-1/2}(\hat{b} - b) \xrightarrow{d} \mathcal{N}(0, I_r)$ where $b = Q' S_{xx}^{1/2} \beta$,
2. $\mathbb{V}[\hat{b}]^{-1} \hat{\mathbb{V}}[\hat{b}] \xrightarrow{p} I_r$,
3. $\hat{\theta} = \sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right) + o_p(\mathbb{V}[\hat{\theta}]^{1/2})$,

Proof. The proof has two steps: First, we write $\hat{\theta}$ as $\sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right)$ plus an approximation error which is of smaller order than $\mathbb{V}[\hat{\theta}]$. This argument establishes the last two claims of the lemma. Second, we use Lyapounov's CLT to show that $\hat{b} \in \mathbb{R}^r$ is jointly asymptotically normal.

Decomposition and Approximation From the proof of Lemma 2 it follows that

$$\hat{\theta} = \sum_{\ell=1}^r \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right) + \sum_{i=1}^n B_{ii} (\sigma_i^2 - \hat{\sigma}_i^2)$$

where we now show that the mean zero random variable $\sum_{i=1}^n B_{ii} (\sigma_i^2 - \hat{\sigma}_i^2)$ is $o_p(\mathbb{V}[\hat{\theta}]^{1/2})$.

We have

$$\sum_{i=1}^n B_{ii} (\hat{\sigma}_i^2 - \sigma_i^2) = \sum_{i=1}^n B_{ii} \sum_{\ell=1}^n M_{ii}^{-1} x_i' \beta M_{i\ell} \varepsilon_\ell + \sum_{i=1}^n B_{ii} (\varepsilon_i^2 - \sigma_i^2) + \sum_{i=1}^n B_{ii} \sum_{\ell \neq i} M_{ii}^{-1} M_{i\ell} \varepsilon_i \varepsilon_\ell.$$

The variances of these three terms are

$$\begin{aligned} \sum_{\ell=1}^n \sigma_{\ell}^2 \left(\sum_{i=1}^n M_{i\ell} B_{ii} M_{ii}^{-1} x'_i \beta \right)^2 &\leq \max_i \sigma_i^2 \sum_{i=1}^n B_{ii}^2 M_{ii}^{-2} (x'_i \beta)^2 \leq \max_i \sigma_i^2 \max_i (x'_i \beta)^2 M_{ii}^{-2} \times \sum_{i=1}^n B_{ii}^2, \\ &\sum_{i=1}^n B_{ii}^2 \mathbb{V}[\varepsilon_i^2] \leq \max_i \mathbb{E}[\varepsilon_i^4] \times \sum_{i=1}^n B_{ii}^2, \\ \sum_{i=1}^n \sum_{\ell \neq i} \left(B_{ii}^2 M_{ii}^{-2} + B_{ii} M_{ii}^{-1} B_{\ell\ell} M_{\ell\ell}^{-1} \right) M_{i\ell}^2 \sigma_i^2 \sigma_{\ell}^2 &\leq 2 \max_i \sigma_i^4 M_{ii}^{-2} \times \sum_{i=1}^n B_{ii}^2. \end{aligned}$$

Furthermore, we have that

$$\mathbb{V}[\hat{\theta}]^{-1} \sum_{i=1}^n B_{ii}^2 \leq \max_i w'_i w_i \mathbb{V}[\hat{\theta}]^{-1} \sum_{l=1}^r \lambda_l^2(\tilde{A}) \leq \max_i w'_i w_i \max_i \sigma_i^{-4} = o(1),$$

so each of the three variances are of smaller order than $\mathbb{V}[\hat{\theta}]$.

For the second claim it suffices to show that $\delta(v) := \frac{\hat{\mathbb{V}}[v'\hat{b}] - \mathbb{V}[v'\hat{b}]}{\mathbb{V}[v'\hat{b}]} = o_p(1)$ for all nonrandom $v \in \mathbb{R}^r$ with $v'v = 1$. Let $v \in \mathbb{R}^r$ be nonrandom with $v'v = 1$. As above we have that $\delta(v) = \sum_{i=1}^n w_i(v) (\hat{\sigma}_i^2 - \sigma_i^2)$ is a mean zero variable which is $o_p(1)$ if $\sum_{i=1}^n w_i(v)^4 = o(1)$ where $w_i(v) = \frac{(v'w_i)^2}{\sum_{i=1}^n \sigma_i^2 (v'w_i)^2}$. But this follows from

$$\sum_{i=1}^n w_i(v)^4 \leq \max_i \sigma_i^{-4} \max_i w'_i w_i = o(1)$$

where the inequality is implied by $\max_i w'_i w_i = o(1)$, $v'v = 1$, and $\sum_{i=1}^n w_i w'_i = I_r$.

Asymptotic Normality Next we show that all linear combinations of \hat{b} are asymptotically normal. Let $v \in \mathbb{R}^r$ be a non-random vector with $v'v = 1$. Lyapunov's CLT implies that $\mathbb{V}[v'\hat{b}]^{-1/2} v'(\hat{b} - b) \xrightarrow{d} N(0, 1)$ if

$$\mathbb{V}[v'\hat{b}]^{-2} \sum_{i=1}^n \mathbb{E}[\varepsilon_i^4] (v'Q'S_{xx}^{-1/2}x_i)^4 = \mathbb{V}[v'\tilde{\beta}]^{-2} \sum_{i=1}^n \mathbb{E}[\varepsilon_i^4] (v'w_i)^4 = o(1). \quad (14)$$

We have that $\max_i w'_i w_i = o(1)$ implies (14) since $\max_i (v'w_i)^2 \leq \max_i w'_i w_i$ and

$$\sum_{i=1}^n (v'w_i)^2 = 1, \quad \mathbb{V}[v'\tilde{\beta}]^{-1} \leq \max_i \sigma_i^{-2} = O(1), \quad \max_i \mathbb{E}[\varepsilon_i^4] = O(1),$$

by definition of w_i and Assumption 1. □

C.4 Quadratic Forms of Growing Rank

This appendix provides restatements and proofs of Theorems 2 and 3. The proofs relies on an auxiliary lemma which extends a central limit theorem given in Solvsten (2019).

C.4.1 A Central Limit Theorem

The proofs of Theorem 2 and Theorem 3 is based on the following lemma. Let $\{v_{n,i}\}_{i,n}$ be a triangular array of row-wise independent random variables with $\mathbb{E}[v_{n,i}] = 0$ and $\mathbb{V}[v_{n,i}] = \sigma_{n,i}^2$, let $\{\dot{w}_{n,i}\}_{i,n}$ be a triangular array of non-random weights that satisfy $\sum_{i=1}^n \dot{w}_{n,i}^2 \sigma_{n,i}^2 = 1$ for all n , and let $(W_n)_n$ be a sequence of symmetric non-random matrices in $\mathbb{R}^{n \times n}$ with zeroes on the diagonal that satisfy $2 \sum_{i=1}^n \sum_{\ell \neq i} W_{n,i\ell}^2 \sigma_{n,i}^2 \sigma_{n,\ell}^2 = 1$. For simplicity, we drop the subscript n on $v_{n,i}$, $\sigma_{n,i}^2$, $\dot{w}_{n,i}$ and W_n . Define

$$\mathcal{S}_n = \sum_{i=1}^n \dot{w}_i v_i \quad \text{and} \quad \mathcal{U}_n = \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell} v_i v_\ell.$$

Lemma C.6. *If $\max_i \mathbb{E}[v_i^4] + \sigma_i^{-2} = O(1)$,*

$$(i) \max_i \dot{w}_i^2 = o(1), \quad (ii) \lambda_{\max}(W^2) = o(1),$$

then $(\mathcal{S}_n, \mathcal{U}_n)' \xrightarrow{d} \mathcal{N}(0, I_2)$.

This lemma extends the main result of Appendix A2 in Solvsten (2019) to allow for $\{v_i\}_i$ to be an array of non-identically distributed variables and presents the conclusion in a way that is tailored to the application in this paper. The proof requires no substantially new ideas compared to Solvsten (2019), but we give it at the end of the next section for completeness.

C.4.2 Limit Distributions

Theorem C.2. *If*

$$(i) \mathbb{V}[\hat{\theta}]^{-1} \max_i \left((\tilde{x}'_i \beta)^2 + (\tilde{x}_i' \beta)^2 \right) = o(1), \quad (ii) \frac{\lambda_1^2}{\sum_{\ell=1}^r \lambda_\ell^2} = o(1),$$

and Assumption 1 holds, then $\mathbb{V}[\hat{\theta}]^{-1/2}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, 1)$.

Proof. The proof involves two steps: First, we decompose $\hat{\theta}$ into a weighted sum of two terms of the type described in Lemma C.6. Second, we use Lemma C.6 to show joint asymptotic normality of the two terms. The conclusion that $\hat{\theta}$ is asymptotically normal is immediate from there.

Decomposition The difference between $\hat{\theta}$ and θ is

$$\hat{\theta} - \theta = \sum_{i=1}^n (2\tilde{x}'_i\beta - \check{x}'_i\beta) \varepsilon_i + \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell} \varepsilon_i \varepsilon_\ell,$$

where these two terms are uncorrelated and have variances

$$V_S = \sum_{i=1}^n (2\tilde{x}'_i\beta - \check{x}'_i\beta)^2 \sigma_i^2 \quad \text{and} \quad V_U = 2 \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell}^2 \sigma_i^2 \sigma_\ell^2.$$

Thus we write $\mathbb{V}[\hat{\theta}]^{-1/2}(\hat{\theta} - \theta) = \omega_1 \mathcal{S}_n + \omega_2 \mathcal{U}_n$ where

$$\begin{aligned} \mathcal{S}_n &= V_S^{-1/2} \sum_{i=1}^n (2\tilde{x}'_i\beta - \check{x}'_i\beta) \varepsilon_i, & \mathcal{U}_n &= V_U^{-1/2} \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell} \varepsilon_i \varepsilon_\ell, \\ \omega_1 &= \sqrt{V_S / \mathbb{V}[\hat{\theta}]}, & \omega_2 &= \sqrt{V_U / \mathbb{V}[\hat{\theta}]}. \end{aligned}$$

Asymptotic Normality We will argue along converging subsequences. Move to a subsequence where ω_1 converges. If the limit is zero, then $\mathbb{V}[\hat{\theta}]^{-1/2}(\hat{\theta} - \theta) = \omega_2 \mathcal{U}_n + o_p(1)$ and so it follows from Result C.2 below and Theorem 2(ii) that $\hat{\theta}$ is asymptotically normal. Thus we consider the case where the limit of ω_1 is nonzero.

In the notation of Lemma C.6 we have

$$\dot{w}_i = \frac{(2\tilde{x}'_i\beta - \check{x}'_i\beta)}{V_S^{1/2}} \quad \text{and} \quad W_{i\ell} = \frac{C_{i\ell}}{V_U^{1/2}}.$$

For Lemma C.6(i) we have

$$\max_i \dot{w}_i^2 \leq 4\omega_1^{-1} \max_i \frac{(\tilde{x}'_i\beta)^2 + (\check{x}'_i\beta)^2}{\mathbb{V}[\hat{\theta}]} = o(1),$$

where the last equality follows from Theorem 2(i) and the nonzero limit of ω_1 .

For Lemma C.6(ii) we show instead that $\text{trace}(W^4) = o(1)$. It can be shown that for all n , $\text{trace}(C^4) \leq c_U \cdot \text{trace}(B^4) = c_U \cdot \text{trace}(\tilde{A}^4) \leq c_U \lambda_1^2 \cdot \text{trace}(\tilde{A}^2)$ and $V_U \geq c_L \min_i \sigma_i^4 \cdot \text{trace}(\tilde{A})$, where the finite and nonzero constants c_U and c_L do not depend on n (but depend on $\min_i M_{ii}$ which is bounded away from zero). Thus, Assumption 1 implies that

$$\text{trace}(W^4) \leq \frac{c_U \lambda_1^2 \cdot \text{trace}(\tilde{A}^2)}{(c_L \min_i \sigma_i^4 \cdot \text{trace}(\tilde{A}^2))^2} = O\left(\frac{\lambda_1^2}{\text{trace}(\tilde{A}^2)}\right) = o(1)$$

where the last equality follows from Theorem 2(ii). □

Theorem C.3. If $\max_i \mathbf{w}'_{iq} \mathbf{w}_{iq} = o(1)$, $\mathbb{V}[\hat{\theta}_q]^{-1} \max_i \left((\tilde{x}'_{iq} \beta)^2 + (\hat{x}'_{iq} \beta)^2 \right) = o(1)$, and Assumptions 1 and 2 holds, then

1. $\mathbb{V}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)']^{-1/2} \left((\hat{\mathbf{b}}'_q, \hat{\theta}_q)' - \mathbb{E}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] \right) \xrightarrow{d} \mathcal{N}(0, I_{q+1})$
2. $\hat{\theta} = \sum_{\ell=1}^q \lambda_\ell \left(\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell] \right) + \hat{\theta}_q + o_p(\mathbb{V}[\hat{\theta}]^{1/2})$

for

$$\mathbb{V}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] = \sum_{i=1}^n \begin{bmatrix} \mathbf{w}_{iq} \mathbf{w}'_{iq} \sigma_i^2 & 2\mathbf{w}_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} x'_\ell \beta \right) \sigma_i^2 \\ 2\mathbf{w}'_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} x'_\ell \beta \right) \sigma_i^2 & 4 \left(\sum_{\ell \neq i} C_{i\ell q} x'_\ell \beta \right)^2 \sigma_i^2 + 2 \sum_{\ell \neq i} C_{i\ell q}^2 \sigma_i^2 \sigma_\ell^2 \end{bmatrix},$$

$$C_{i\ell q} = B_{i\ell q} - 2^{-1} M_{i\ell} \left(M_{ii}^{-1} B_{iiq} + M_{\ell\ell}^{-1} B_{\ell\ell q} \right), \quad B_{i\ell q} = x'_i S_{xx}^{-1/2} \tilde{A}_q S_{xx}^{-1/2} x_\ell, \quad \tilde{A}_q = \sum_{\ell=q+1}^r \lambda_\ell q_\ell q'_\ell, \\ \tilde{x}_{iq} = \sum_{\ell=1}^n B_{i\ell q} x_\ell, \quad \text{and} \quad \check{x}_{iq} = \sum_{\ell=1}^n M_{i\ell} M_{\ell\ell}^{-1} B_{\ell\ell q} x_\ell.$$

Proof. The proof involves two steps: First, we write $\hat{\theta}$ as the sum of (1a) a quadratic function applied to $\hat{\mathbf{b}}_q$, (1b) an approximation error which is of smaller order than $\mathbb{V}[\hat{\theta}]$, and (2) a weighted sum of two terms, \mathcal{S}_n and \mathcal{U}_n , of the type described in Lemma C.6. Second, we use Lemma C.6 to show that $(\hat{\mathbf{b}}'_q, \mathcal{S}_n, \mathcal{U}_n)' \in \mathbb{R}^{q+2}$ is jointly asymptotically normal.

Decomposition and Approximation We have that

$$\hat{\theta} = \sum_{\ell=1}^q \lambda_\ell (\hat{b}_\ell^2 - \mathbb{V}[\hat{b}_\ell]) + \hat{\theta}_q + o_p(\mathbb{V}[\hat{\theta}]^{1/2}) \quad \text{for} \quad \hat{\theta}_q = \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell q} y_i y_\ell$$

since

$$\hat{\beta}' A \hat{\beta} = \sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \sum_{i=1}^n \sum_{\ell=1}^n B_{i\ell q} y_i y_\ell$$

and

$$\begin{aligned} \sum_{i=1}^n B_{ii} \hat{\sigma}_i^2 &= \sum_{i=1}^n B_{ii1} \sigma_i^2 + \sum_{i=1}^n B_{iiq} \hat{\sigma}_i^2 + \sum_{i=1}^n B_{ii,-q} (\hat{\sigma}_i^2 - \sigma_i^2) \\ &= \sum_{\ell=1}^q \lambda_\ell \mathbb{V}[\hat{b}_\ell] + \sum_{i=1}^n B_{iiq} \hat{\sigma}_i^2 + o_p(\mathbb{V}[\hat{\theta}]^{1/2}) \end{aligned}$$

where $B_{ii,-q} = B_{ii} - B_{iiq}$ and it follows from $\max_i \mathbf{w}'_{iq} \mathbf{w}_{iq} = o(1)$ and the calculations in the proof of Theorem 1 that the mean zero random variable $\sum_{i=1}^n B_{ii,-q} (\hat{\sigma}_i^2 - \sigma_i^2)$ is $o_p(\mathbb{V}[\hat{\theta}]^{1/2})$.

We will further center and rescale $\hat{\theta}_q$ by writing

$$\mathbb{V}[\hat{\theta}_q]^{-1/2} \left(\hat{\theta}_q - \mathbb{E}[\hat{\theta}_q] \right) = \omega_1 \mathcal{S}_n + \omega_2 \mathcal{U}_n$$

where

$$\begin{aligned}\mathcal{S}_n &= V_S^{-1/2} \sum_{i=1}^n (2\tilde{x}'_{iq}\beta - \tilde{x}'_{iq}\beta) \varepsilon_i, & \mathcal{U}_n &= V_U^{-1/2} \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell q} \varepsilon_i \varepsilon_\ell, \\ V_S &= \sum_{i=1}^n (2\tilde{x}'_{iq}\beta - \tilde{x}'_{iq}\beta)^2 \sigma_i^2, & V_U &= 2 \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell q}^2 \sigma_i^2 \sigma_\ell^2, \\ \omega_1 &= \sqrt{V_S / \mathbb{V}[\hat{\theta}_q]}, & \omega_2 &= \sqrt{V_U / \mathbb{V}[\hat{\theta}_q]},\end{aligned}$$

and \mathcal{U}_n is uncorrelated with both \mathcal{S}_n and $\hat{\mathbf{b}}_q$.

Asymptotic Normality As in the proof of Theorem 2, we will argue along converging subsequences and therefore move to a subsequence where ω_1 converges. If the limit is zero, then the conclusion of the theorem follows from Lemma C.6 applied to $(\mathbb{V}[v'\hat{\mathbf{b}}_q]^{-1/2}(v'\hat{\mathbf{b}}_q - \mathbb{E}[v'\hat{\mathbf{b}}_q]), \mathcal{U}_n)'$ for $v \in \mathbb{R}^q$ with $v'v = 1$. Thus we consider the case where the limit of ω_1 is nonzero.

Next we use Lemma C.6 to show that

$$\left(\frac{v'\hat{\mathbf{b}}_q - \mathbb{E}[v'\hat{\mathbf{b}}_q] + u\mathcal{S}_n}{\mathbb{V}[\hat{\mathbf{b}}_q + u\mathcal{S}_n]^{1/2}}, \mathcal{U}_n \right)' \xrightarrow{d} \mathcal{N}(0, I_2)$$

for any non-random $(v', u)' \in \mathbb{R}^{q+1}$ with $v'v + u^2 = 1$. In the notation of Lemma C.6 we have

$$\dot{w}_i = \frac{v'w_{iq} + uV_S^{-1/2} (2\tilde{x}'_{iq}\beta - \tilde{x}'_{iq}\beta)}{\mathbb{V}[\hat{\mathbf{b}}_q + u\mathcal{S}_n]^{1/2}} \quad \text{and} \quad W_{i\ell} = \frac{C_{i\ell q}}{V_U^{1/2}}.$$

A simple calculation shows that $\mathbb{V}[v'\hat{\mathbf{b}}_q + u\mathcal{S}_n] \geq \min_i \sigma_i^2 \gg 0$, so $\max_i \dot{w}_i^2 = o(1)$ follows from Theorem 3(i), Theorem 3(ii), and ω_1 being bounded away from zero.

Similarly, we have as in the proof of Theorem 2 that

$$\text{trace}(C_q^4) \leq c \text{trace}(B_q^4) \leq c \lambda_{q+1}^2 \sum_{\ell=q+1}^r \lambda_\ell^2 \quad \text{and} \quad V_U^2 \geq \omega_2^{-4} \min_i \sigma_i^8 \text{trace}(\tilde{A}^2)^2$$

for $C_q = (C_{i\ell q})_{i,\ell}$ and $B_q = (B_{i\ell q})_{i,\ell}$, so Assumptions 1 and 2 yield $\text{trace}(W^4) = o(1)$. \square

C.4.3 Proof of a Central Limit Theorem

The proof of Lemma C.6 uses the notation and verifies the conditions of Lemmas A2.1 and A2.2 in S¸olvsten (2019) referred to as SS2.1 and SS2.2, respectively. First, we show marginal convergence in distribution of \mathcal{S}_n and \mathcal{U}_n . Then, we show joint convergence in distribution of \mathcal{S}_n and \mathcal{U}_n . Let $V_n = (v_1, \dots, v_n)$ where $\{v_i\}_i$ are as in the setup of Lemma C.6.

Before starting we note that $\max_i \sigma_i^{-2} = O(1)$ and $2 \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell}^2 \sigma_i^2 \sigma_\ell^2 = 1$ implies that

$\text{trace}(W^2) = \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell}^2 = O(1)$ and therefore that

$$\lambda_{\max}(W^2) = o(1) \Leftrightarrow \text{trace}(W^4) = o(1).$$

Marginal Distributions

Result C.1. $\max_i \mathbb{E}[v_i^4] + \sigma_i^{-2} = O(1)$, $\sum_{i=1}^n \dot{w}_i^2 \sigma_i^2 = 1$, and Lemma C.6(i) implies that $\mathcal{S}_n \xrightarrow{d} \mathcal{N}(0, 1)$.

In the notation of SS2.1 we have,

$$\Delta_i^0 \mathcal{S}_n = \dot{w}_i v_i \quad \text{and} \quad E[T_n | V_n] = 1 + \frac{1}{2} \sum_{i=1}^n \dot{w}_i^2 (v_i^2 - \sigma_i^2),$$

and it follows from $\max_i \mathbb{E}[v_i^4] + \sigma_i^{-2} = O(1)$, $\sum_{i=1}^n \dot{w}_i^2 \sigma_i^2 = 1$, and Lemma C.6(i) that

$$E[T_n | V_n] \xrightarrow{\mathcal{L}^1} 1, \quad \sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{S}_n)^2] = 1, \quad \sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{S}_n)^4] \leq \max_i \frac{\mathbb{E}[v_i^4]}{\sigma_i^2} \dot{w}_i^2 = o(1),$$

so Result C.1 follows from SS2.1.

Result C.2. $\max_i \mathbb{E}[v_i^4] + \sigma_i^{-2} = O(1)$, $2 \sum_{i=1}^n \sum_{\ell \neq i} W_{n,i\ell}^2 \sigma_{n,i}^2 \sigma_{n,\ell}^2 = 1$, and Lemma C.6(ii) implies that $\mathcal{U}_n \xrightarrow{d} \mathcal{N}(0, 1)$.

In the notation of SS2.1 we have,

$$\Delta_i^0 \mathcal{U}_n = 2v_i \sum_{\ell \neq i} W_{i\ell} v_\ell \quad \text{and} \quad E[T_n | V_n] = \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i} (v_i + \sigma_i^2) W_{i\ell} W_{ik} v_\ell v_k,$$

and

$$\sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{U}_n)^2] = 2, \quad \sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{U}_n)^4] \leq 2^5 \max_i \mathbb{E}[v_i^4]^2 \max_i \sigma_i^{-4} \max_i \sum_{\ell \neq i} W_{i\ell}^2,$$

where $\max_i \sum_{\ell \neq i} W_{i\ell}^2 \leq \sqrt{\text{trace}(W^4)} = o(1)$. Now, split $E[T_n | V_n] - 1$ into three terms

$$\begin{aligned} a_n &= \sum_{i=1}^n \sum_{\ell \neq i} \sigma_i^2 W_{i\ell}^2 (v_\ell + v_\ell^2 - \sigma_\ell^2) \\ b_n &= 2 \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i, \ell} \sigma_k^2 W_{\ell k} W_{ik} v_i v_\ell + \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell}^2 v_i (v_\ell^2 - \sigma_\ell^2) \\ c_n &= \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i, \ell} W_{i\ell} W_{ik} (v_i^2 - \sigma_i^2) v_\ell v_k. \end{aligned}$$

Interlude: Convergence in \mathcal{L}^1

a_n , b_n , and c_n are a linear sum, a quadratic sum, and a cubic sum. We will need to treat similar sums later, so we record some simple sufficient conditions for their convergence. For brevity, let $\sum_{i \neq \ell}^n = \sum_{i=1}^n \sum_{\ell \neq i}$, and $\sum_{i \neq \ell \neq k}^n = \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i, \ell}$, etc. We use the notation $u_i = (v_{i1}, v_{i2}, v_{i3}, v_{i4}) \in \mathbb{R}^4$ to denote independent random vectors in order that the result applies to combinations of v_i and $v_i^2 - \sigma_i^2$ as in a_n , b_n , and c_n above. For the inferential results we will also treat quartic sums, so we provide the sufficient conditions here.

Result C.3. *Let $S_{n1} = \sum_{i=1}^n \omega_i v_{i1}$, $S_{n2} = \sum_{i \neq \ell}^n \omega_{i\ell} v_{i1} v_{\ell 2}$, $S_{n3} = \sum_{i \neq \ell \neq k}^n \omega_{i\ell k} v_{i1} v_{\ell 2} v_{k3}$, and $S_{n4} = \sum_{i \neq \ell \neq k \neq m}^n \omega_{i\ell k m} v_{i1} v_{\ell 2} v_{k3} v_{m4}$ where the weights ω_i , $\omega_{i\ell}$, $\omega_{i\ell k}$, and $\omega_{i\ell k m}$ are non-random. Suppose that $\mathbb{E}[u_i] = 0$, $\max_i \mathbb{E}[u_i' u_i] = O(1)$.*

1. If $\sum_{i=1}^n \omega_i^2 = o(1)$, then $S_{n1} \xrightarrow{\mathcal{L}^1} 0$.
2. If $\sum_{i \neq \ell}^n \omega_{i\ell}^2 = o(1)$, then $S_{n2} \xrightarrow{\mathcal{L}^1} 0$.
3. If $\sum_{i \neq \ell \neq k}^n \omega_{i\ell k}^2 = o(1)$, then $S_{n3} \xrightarrow{\mathcal{L}^1} 0$.
4. If $\sum_{i \neq \ell \neq k \neq m}^n \omega_{i\ell k m}^2 = o(1)$, then $S_{n4} \xrightarrow{\mathcal{L}^1} 0$.

Consider S_{n3} , the other results follows from the same line of reasoning. In the notation of SS2.2 we have,

$$\Delta_i^0 S_{n3} = v_{i1} \sum_{\ell \neq i} \sum_{k \neq i, \ell} \omega_{i\ell k} v_{\ell 2} v_{k3} + v_{i2} \sum_{\ell \neq i} \sum_{k \neq i, \ell} \omega_{i\ell k} v_{\ell 1} v_{k3} + v_{i3} \sum_{\ell \neq i} \sum_{k \neq i, \ell} \omega_{i\ell k} v_{\ell 1} v_{k2}.$$

Focusing on the first term we have,

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \left[\left(v_{i1} \sum_{\ell \neq i} \sum_{k \neq i, \ell} \omega_{i\ell k} v_{\ell 2} v_{k 3} \right)^2 \right] &\leq \max_i \mathbb{E}[u'_i u_i]^3 \sum_{i \neq \ell \neq k}^n \left(\omega_{i\ell k}^2 + \omega_{i\ell k} \omega_{ik\ell} \right) \\ &\leq 2 \max_i \mathbb{E}[u'_i u_i]^3 \sum_{i \neq \ell \neq k}^n \omega_{i\ell k}^2, \end{aligned}$$

so the results follows from SS2.2, $\sum_{i \neq \ell \neq k}^n \omega_{i\ell k}^2 = o(1)$, and the observation that the last bound also applies to the other two terms in $\Delta_i^0 \mathcal{S}_{n3}$.

Marginal Distributions, Continued

To see how $a_n \xrightarrow{\mathcal{L}^1} 0$, $b_n \xrightarrow{\mathcal{L}^1} 0$ and $c_n \xrightarrow{\mathcal{L}^1} 0$ follows from Result C.3, let $\bar{W}_{i\ell} = \sum_{k=1}^n W_{ik} W_{k\ell}$ and note that $\text{trace}(W^4) = \sum_{i=1}^n \sum_{\ell=1}^n \bar{W}_{i\ell}^2$. We have

$$\begin{aligned} \sum_{i=1}^n \left(\sum_{\ell \neq i} \sigma_\ell^2 W_{i\ell}^2 \right)^2 &\leq \max_i \sigma_i^4 \sum_{i=1}^n \bar{W}_{ii}^2. \\ \sum_{i=1}^n \sum_{\ell \neq i} \left(\sum_{k \neq i, \ell} \sigma_k^2 W_{\ell k} W_{ik} \right)^2 &\leq \max_i \sigma_i^4 \sum_{i=1}^n \sum_{\ell=1}^n \bar{W}_{i\ell}^2 \\ \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell}^4 &= O \left(\max_{i, \ell} W_{i\ell}^2 \right) \\ \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i, \ell} W_{i\ell}^2 W_{ik}^2 &= O \left(\max_i \sum_{\ell \neq i} W_{i\ell}^2 \right), \end{aligned}$$

all of which are $o(1)$ as $\text{trace}(W^4) = o(1)$.

Joint Distribution

Let $(u_1, u_2)' \in R^2$ be given and non-random with $u_1^2 + u_2^2 = 1$. Define $\mathcal{W}_n = u_1 \mathcal{S}_n + u_2 \mathcal{U}_n$. Lemma C.6 follows if we show that $\mathcal{W}_n \xrightarrow{d} \mathcal{N}(0, 1)$. In the notation of SS2.1 we have,

$$\Delta_i^0 \mathcal{W}_n = u_1 \dot{w}_i v_i + u_2 2v_i \sum_{\ell \neq i} W_{i\ell} v_\ell$$

and

$$\begin{aligned}\mathbb{E}[T_n | V_n] &= u_1^2 \left(1 + \frac{1}{2} \sum_{i=1}^n \dot{w}_i^2 (v_i^2 - \sigma_i^2) \right) + u_2^2 \sum_{i=1}^n \sum_{\ell \neq i} \sum_{k \neq i} (v_i + \sigma_i^2) W_{i\ell} W_{ik} v_\ell v_k \\ &\quad + u_1 u_2 3 \sum_{i=1}^n \sum_{\ell \neq i} (v_i^2 + \sigma_i^2) \dot{w}_i W_{i\ell} v_j.\end{aligned}$$

The proofs of Result C.1 and Result C.2 showed that

$$\sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{W}_n)^2] = O(1), \quad \sum_{i=1}^n \mathbb{E}[(\Delta_i^0 \mathcal{W}_n)^4] = o(1)$$

and that the first two terms of $\mathbb{E}[T_n | V_n]$ converge to $u_1^2 + u_2^2 = 1$. Thus the lemma follows if we show that the “conditional covariance”

$$3 \sum_{i=1}^n \sum_{\ell \neq i} (v_i^2 + \sigma_i^2) \dot{w}_i W_{i\ell} v_j$$

converges to 0 in \mathcal{L}^1 . This conditional covariance involves a linear and a quadratic sum so

$$\begin{aligned}\sum_{i=1}^n \left(\sum_{\ell \neq i} \sigma_\ell^2 w_\ell W_{i\ell} \right)^2 &\leq \max_i \sigma_i^4 \max_\ell \lambda_\ell^2(W) \sum_{i=1}^n \dot{w}_i^2 = O(\max_\ell \lambda_\ell^2(W)) \\ \sum_{i=1}^n \sum_{\ell \neq i} \dot{w}_i^2 W_{i\ell}^2 &\leq \sum_{i=1}^n \sum_{\ell \neq i} W_{i\ell}^2 \max_i \dot{w}_i^2 = O(\max_i \dot{w}_i^2)\end{aligned}$$

ends the proof.

C.5 Asymptotic Variance Estimation

This appendix provides restatements and proofs of Lemmas 5 and 6 which establish consistency of the proposed standard error estimators that rely on sample splitting. Furthermore, it gives adjustments to those standard errors that guarantee existence whenever two independent unbiased estimators of $x'_i \beta$ cannot be formed. However, these adjustments may provide a somewhat conservative assessment of the uncertainty in $\hat{\theta}$ as further investigated in the simulations of Section 8.7.

Lemma C.7. *For $s = 1, 2$, suppose that $\widehat{x'_i \beta}_{-i,s} = \sum_{\ell \neq i}^n P_{i\ell,s} y_\ell$ satisfies $\sum_{\ell \neq i}^n P_{i\ell,s} x'_\ell \beta = x'_i \beta$, $P_{i\ell,1} P_{i\ell,2} = 0$ for all ℓ , and $\lambda_{\max}(P_s P'_s) = O(1)$.*

1. *If the conditions of Theorem 2 hold and $|\mathcal{B}| = O(1)$, then $\frac{\hat{\theta} - \theta}{\widehat{\mathbb{V}}[\hat{\theta}]^{1/2}} \xrightarrow{d} \mathcal{N}(0, 1)$.*

2. If the conditions of Theorem 2 hold, then $\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \left[\hat{\theta} \pm z_\alpha \hat{\mathbb{V}}[\hat{\theta}]^{1/2} \right] \right) \geq 1 - \alpha$.

Proof. The proof continues in two steps: First, we show that $\hat{\mathbb{V}}[\hat{\theta}]$ has a positive bias which is of smaller order than $\mathbb{V}[\hat{\theta}]$ when $|\mathcal{B}| = O(1)$. Second, we show that $\hat{\mathbb{V}}[\hat{\theta}] - \mathbb{E}[\hat{\mathbb{V}}[\hat{\theta}]] = o_p(\mathbb{V}[\hat{\theta}])$. When combined with Theorem 2, these conclusions imply the two claims of the lemma.

Bias of $\hat{\mathbb{V}}[\hat{\theta}]$ For the first term in $\hat{\mathbb{V}}[\hat{\theta}]$, a simple calculation shows that

$$\begin{aligned} \mathbb{E} \left[4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} y_\ell \right)^2 \tilde{\sigma}_i^2 \right] &= 4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} x'_\ell \beta \right)^2 \sigma_i^2 + 4 \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell}^2 \sigma_i^2 \sigma_\ell^2 \\ &\quad + 4 \sum_{i=1}^n \sum_{\ell \neq i} \sum_{m=1}^n C_{mi} C_{m\ell} (P_{mi,1} P_{m\ell,2} + P_{mi,2} P_{m\ell,1}) \sigma_i^2 \sigma_\ell^2 \\ &= \mathbb{V}[\hat{\theta}] + 2 \sum_{i=1}^n \sum_{\ell \neq i} \tilde{C}_{i\ell} \sigma_i^2 \sigma_\ell^2. \end{aligned}$$

For the second term in $\hat{\mathbb{V}}[\hat{\theta}]$, we note that if $P_{ik,-\ell} P_{\ell k,-i} = 0$ for all k , then independence between error terms yield $\mathbb{E}[\widehat{\sigma_i^2 \sigma_\ell^2}] = \mathbb{E}[\hat{\sigma}_{i,-\ell}^2] \mathbb{E}[\hat{\sigma}_{\ell,-i}^2] = \sigma_i^2 \sigma_\ell^2$. Otherwise if $P_{i\ell,1} + P_{i\ell,2} = 0$, then

$$\begin{aligned} \mathbb{E} \left[\widehat{\sigma_i^2 \sigma_\ell^2} \right] &= \mathbb{E} \left[\left(\varepsilon_i - \sum_{j \neq i} P_{ij,1} \varepsilon_j \right) \left(\varepsilon_i - \sum_{k \neq i} P_{ik,2} \varepsilon_k \right) (x'_\ell \beta + \varepsilon_\ell) \left(\varepsilon_\ell - \sum_{m \neq \ell} P_{\ell m,-i} \varepsilon_m \right) \right] \\ &= \sigma_i^2 \sigma_\ell^2 + x'_\ell \beta \mathbb{E} \left[\left(\varepsilon_i - \sum_{j \neq i} P_{ij,1} \varepsilon_j \right) \left(\varepsilon_i - \sum_{k \neq i} P_{ik,2} \varepsilon_k \right) \sum_{m \neq \ell} P_{\ell m,-i} \varepsilon_m \right] \end{aligned}$$

where the second term is zero since $P_{\ell i,-i} = 0$ and $P_{ij,1} P_{ij,2} = 0$ for all j . The same argument applies with the roles of i and ℓ reversed when $P_{\ell i,1} + P_{\ell i,2} = 0$.

Finally, when $(i, \ell) \in \mathcal{B}$ we have

$$\mathbb{E} \left[\widehat{\sigma_i^2 \sigma_\ell^2} \right] = \left(\sigma_i^2 \left(\sigma_\ell^2 + ((x_\ell - \bar{x})' \beta)^2 \right) + O\left(\frac{1}{n}\right) \right) 1_{\{\tilde{C}_{i\ell} < 0\}}$$

where the remainder is uniform in (i, ℓ) and stems from the use of \bar{y} as an estimator of $\bar{x}'\beta$. Thus for sufficiently large n , $\mathbb{E}[\tilde{C}_{i\ell} \widehat{\sigma_i^2 \sigma_\ell^2}]$ is smaller than $\tilde{C}_{i\ell} \sigma_i^2 \sigma_\ell^2$ leading to a positive bias in $\hat{\mathbb{V}}[\hat{\theta}]$. This bias is

$$\sum_{(i,\ell) \in \mathcal{B}} \tilde{C}_{i\ell} \sigma_i^2 \left(\sigma_\ell^2 1_{\{\tilde{C}_{i\ell} > 0\}} + ((x_\ell - \bar{x})' \beta)^2 1_{\{\tilde{C}_{i\ell} < 0\}} \right) + O\left(\frac{1}{n} \mathbb{V}[\hat{\theta}]\right)$$

which is ignorable when $|\mathcal{B}| = O(1)$.

Variability of $\hat{\mathbb{V}}[\hat{\theta}]$ Now, $\hat{\mathbb{V}}[\hat{\theta}] - \mathbb{E}[\hat{\mathbb{V}}[\hat{\theta}]]$ involves a number of terms all of which are linear, quadratic,

cubic, or quartic sums. Result C.3 provides sufficient conditions for their convergence in \mathcal{L}^1 and therefore in probability. We have already treated versions of linear, quadratic, and cubic terms carefully in the proof of Lemma C.6. Thus, we report here the calculations for the quartic terms (details for the remaining terms can be provided upon request) as they also highlight the role of the high-level condition $\lambda_{\max}(P_s P_s') = O(1)$ for $s = 1, 2$.

The quartic term in $4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} y_\ell \right)^2 \tilde{\sigma}_i^2$ is $\sum_{i \neq \ell \neq m \neq k} \omega_{i\ell mk} \varepsilon_i \varepsilon_\ell \varepsilon_m \varepsilon_k$ where

$$\omega_{i\ell mk} = \sum_{j=1}^n C_{ji} C_{j\ell} M_{jm,1} M_{jk,2} \quad \text{and} \quad M_{i\ell, s} = \begin{cases} 1, & \text{if } i = \ell, \\ -P_{i\ell, s}, & \text{if } i \neq \ell. \end{cases}$$

Letting \odot denote Hadamard (element-wise) product and $M_s = I_n - P_s$, we have

$$\begin{aligned} \sum_{i \neq \ell \neq m \neq k} \omega_{i\ell mk}^2 &\leq \sum_{i, \ell, m, k} \omega_{i\ell mk}^2 = \sum_{j, j'} (C^2)_{jj'}^2 (M_1 M_1')_{jj'} (M_2 M_2')_{jj'} \\ &= \text{trace} \left((C^2 \odot C^2) (M_1 M_1' \odot M_2 M_2') \right) \\ &\leq \lambda_{\max} (M_1 M_1' \odot M_2 M_2') \text{trace} (C^2 \odot C^2) = O \left(\text{trace} (C^4) \right) = o \left(\mathbb{V}[\hat{\theta}]^2 \right) \end{aligned}$$

where $\lambda_{\max} (M_1 M_1' \odot M_2 M_2') = O(1)$ follows from $\lambda_{\max}(P_s P_s') = O(1)$ and we established the last equality in the proof of Theorem 2. The quartic term involved in $2 \sum_{i=1}^n \sum_{\ell \neq i} \tilde{C}_{i\ell} \widehat{\sigma}_i^2 \widehat{\sigma}_\ell^2$ has variability of the same order as $\sum_{i \neq \ell \neq m \neq k} \omega_{i\ell mk} \varepsilon_i \varepsilon_\ell \varepsilon_m \varepsilon_k$ where

$$\omega_{i\ell mk} = \tilde{C}_{i\ell} M_{im,1} M_{lk,1} + \sum_{j=1}^n \tilde{C}_{ij} M_{im,1} M_{jk,1} M_{j\ell,2}.$$

Letting $\tilde{C} = (\tilde{C}_{i\ell})_{i,\ell}$, we find that

$$\begin{aligned} \sum_{i \neq \ell \neq m \neq k} \omega_{i\ell mk}^2 &\leq 2 \sum_{i,\ell} \tilde{C}_{i\ell}^2 (M_1 M_1')_{ii} (M_2 M_2')_{\ell\ell} + 2 \sum_{j,j'} \sum_i \tilde{C}_{ij} \tilde{C}_{ij'} (M_1 M_1')_{ii} (M_1 M_1')_{jj'} (M_2 M_2')_{jj'} \\ &= O \left(\sum_{i,\ell} \tilde{C}_{i\ell}^2 + \text{trace} \left((\tilde{C}^2 \odot M_1 M_1') (M_1 M_1' \odot M_2 M_2') \right) \right) \\ &= O \left(\text{trace} (\tilde{C}^2) \right). \end{aligned}$$

We have $\tilde{C} = C \odot C + 2(C \odot P_1)'(C \odot P_2) + 2(C \odot P_2)'(C \odot P_1)$, from which we obtain that

$$\text{trace}(\tilde{C}^2) = O \left(\left(\max_{i,\ell} C_{i\ell}^2 + \lambda_{\max}(C^2) \right) \text{trace}(C^2) \right) = o \left(\mathbb{V}[\hat{\theta}]^2 \right)$$

where we established the last equality in the proof of Theorem 2. \square

Section 5.2 proposed standard errors for the case of $q > 0$, but left a few details to the appendix since the definitions were completely analogous to the previous lemma. Those definitions are $\tilde{C}_{ilq} = C_{ilq}^2 + 2 \sum_{m=1}^n C_{miq} C_{mlq} (P_{mi,1} P_{ml,2} + P_{mi,2} P_{ml,1})$ where C_{ilq} was introduced in the proof of Theorem 3 and is of the form $C_{ilq} = B_{ilq} - 2^{-1} M_{il} \left(M_{ii}^{-1} B_{iiq} + M_{\ell\ell}^{-1} B_{\ell\ell q} \right)$ for $B_{ilq} = B_{il} - \sum_{s=1}^q \lambda_s w_{is} w_{\ell s}$.

Furthermore, the proposed standard error estimator relies on

$$\widetilde{\sigma_i^2 \sigma_\ell^2} = \begin{cases} \hat{\sigma}_{i,-\ell}^2 \cdot \hat{\sigma}_{\ell,-i}^2, & \text{if } P_{ik,-\ell} P_{\ell k,-i} = 0 \text{ for all } k, \\ \tilde{\sigma}_i^2 \cdot \tilde{\sigma}_{\ell,-i}^2, & \text{else if } P_{i\ell,1} + P_{i\ell,2} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot \tilde{\sigma}_\ell^2, & \text{else if } P_{\ell i,1} + P_{\ell i,2} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot (y_\ell - \bar{y})^2 \cdot 1_{\{\tilde{C}_{ilq} < 0\}}, & \text{otherwise.} \end{cases}$$

Lemma C.8. For $s = 1, 2$, suppose that $\widehat{x'_i \beta}_{-i,s}$ satisfies $\sum_{\ell \neq i} P_{i\ell,s} x'_\ell \beta = x'_i \beta$, $P_{i\ell,1} P_{i\ell,2} = 0$ for all ℓ , and $\lambda_{\max}(P_s P'_s) = O(1)$ where $P_s = (P_{i\ell,s})_{i,\ell}$.

1. If the conditions of Theorem 3 hold and $|\mathcal{B}| = O(1)$, then $\Sigma_q^{-1} \hat{\Sigma}_q \xrightarrow{p} I_{q+1}$.
2. If the conditions of Theorem 3 hold, then $\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \hat{C}_{\alpha,q}^\theta \right) \geq 1 - \alpha$.

The following provides a proof of the first claim of this lemma, while we postpone a proof of the second claim to the end of Appendix C.6.

Proof. The statements $\mathbb{V}[\hat{\mathbf{b}}_q]^{-1} \hat{\mathbb{V}}[\hat{\mathbf{b}}_q] \xrightarrow{p} I_q$ and $\mathbb{V}[\hat{\theta}_q]^{-1} \hat{\mathbb{V}}[\hat{\theta}_q] \xrightarrow{p} 1$ follow by applying the arguments in Theorem C.1 and Lemma C.7. Thus we focus on the remaining claim that

$$\delta(v) := \frac{\hat{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q] - \mathcal{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q]}{\mathbb{V}[v' \hat{\mathbf{b}}_q]^{1/2} \mathbb{V}[\hat{\theta}_q]^{1/2}} \xrightarrow{p} 0 \quad \text{where} \quad \hat{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q] = 2 \sum_{i=1}^n v' \mathbf{w}_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right) \tilde{\sigma}_i^2$$

for all non-random $v \in \mathbb{R}^q$ with $v'v = 1$.

Unbiasedness of $\hat{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q]$ Since $\tilde{\sigma}_i^2$ is unbiased for σ_i^2 , it follows that

$$\mathbb{E} \left[\hat{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q] \right] = 2 \sum_{i=1}^n v' \mathbf{w}_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} x'_\ell \beta \right) \sigma_i^2 + 2 \sum_{i=1}^n v' \mathbf{w}_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} \mathbb{E}[\varepsilon_\ell \tilde{\sigma}_i^2] \right) = \mathcal{C}[v' \hat{\mathbf{b}}_q, \hat{\theta}_q]$$

as split sampling ensures that $\mathbb{E}[\varepsilon_\ell \tilde{\sigma}_i^2] = 0$ for $\ell \neq i$.

Variability of $\hat{C}[v'\hat{\mathbf{b}}_q, \hat{\theta}_q]$ Now, $\hat{C}[v'\hat{\mathbf{b}}_q, \hat{\theta}_q] - C[v'\hat{\mathbf{b}}_q, \hat{\theta}_q]$ is composed of the following linear, quadratic, and quartic sums:

$$\begin{aligned} & \sum_{i=1}^n v' \mathbf{w}_{iq} \left[\left(\varepsilon_i^2 - \sigma_i^2 \right) \sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta + \sigma_i^2 \sum_{\ell \neq i} C_{i\ell q} \varepsilon_{\ell} + \sum_{\ell \neq i} C_{i\ell q} \sigma_{\ell}^2 \sum_{k \neq \ell} (M_{i\ell,1} M_{ik,2} + M_{i\ell,2} M_{ik,1}) \varepsilon_k \right] \\ & \sum_{i=1}^n v' \mathbf{w}_{iq} \left[\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \sum_m \sum_{k \neq m} M_{im,1} M_{ik,2} \varepsilon_m \varepsilon_k + \sum_{\ell \neq i} C_{i\ell q} \varepsilon_{\ell} \left(\varepsilon_i^2 - \sigma_i^2 \right) \right. \\ & \quad \left. + \sum_{\ell \neq i} C_{i\ell q} \sum_{k \neq \ell} (M_{i\ell,1} M_{ik,2} + M_{i\ell,2} M_{ik,1}) \varepsilon_k \left(\varepsilon_{\ell}^2 - \sigma_{\ell}^2 \right) \right] \\ & \sum_{i=1}^n v' \mathbf{w}_{iq} \sum_{\ell \neq i} C_{i\ell q} \sum_{m \neq \ell} \sum_{k \neq m, \ell} M_{im,1} M_{ik,2} \varepsilon_{\ell} \varepsilon_m \varepsilon_k \end{aligned}$$

These seven terms are $o_p(\mathbb{V}[v'\hat{\mathbf{b}}_q]^{1/2} \mathbb{V}[\hat{\theta}_q]^{1/2})$ by Result C.3 as outlined in the following.

$$\begin{aligned} & \sum_{i=1}^n (v' \mathbf{w}_{iq})^2 \left(\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \right)^2 = O(\max_i w'_{iq} \mathbf{w}_{iq} \mathbb{V}[\hat{\theta}_q]) = o(\mathbb{V}[v'\hat{\mathbf{b}}_q] \mathbb{V}[\hat{\theta}_q]) \\ & \sum_{\ell=1}^n \left(\sum_{i=1}^n v' \mathbf{w}_{iq} C_{i\ell q} \right)^2 = O(\lambda_{\max}(C_q^2) \mathbb{V}[v'\hat{\mathbf{b}}_q]) = O(\lambda_{q+1}^2 \mathbb{V}[v'\hat{\mathbf{b}}_q]) = o(\mathbb{V}[v'\hat{\mathbf{b}}_q] \mathbb{V}[\hat{\theta}_q]) \\ & \sum_{k=1}^n \left(\sum_{i=1}^n v' \mathbf{w}_{iq} \sum_{\ell} C_{i\ell q} M_{i\ell,1} M_{ik,2} \right)^2 = O(\max_i w'_{iq} \mathbf{w}_{iq} \text{trace}(C_q M_1 \odot C_q M_1)) = o(\mathbb{V}[v'\hat{\mathbf{b}}_q] \mathbb{V}[\hat{\theta}_q]) \\ & \sum_{m=1}^n \sum_{k=1}^n \left(\sum_{i=1}^n v' \mathbf{w}_{iq} \sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta M_{im,1} M_{ik,2} \right)^2 = O \left(\sum_{i=1}^n (v' \mathbf{w}_{iq})^2 \left(\sum_{\ell \neq i} C_{i\ell q} x'_{\ell} \beta \right)^2 \right) \\ & \sum_{i=1}^n \sum_{\ell \neq i} C_{i\ell q}^2 (v' \mathbf{w}_{iq})^2 = O(\max_i w'_{iq} \mathbf{w}_{iq} \mathbb{V}[\hat{\theta}_q]) \\ & \sum_{k=1}^n \sum_{\ell=1}^n \left(\sum_{i=1}^n v' \mathbf{w}_{iq} C_{i\ell q} M_{i\ell,1} M_{ik,2} \right)^2 = O \left(\mathbb{V}[v'\hat{\mathbf{b}}_q] \lambda_{\max}((C_q \odot M_1)(C_q \odot M_1)') \right) = o(\mathbb{V}[v'\hat{\mathbf{b}}_q] \mathbb{V}[\hat{\theta}_q]) \\ & \sum_{\ell=1}^n \sum_{m=1}^n \sum_{k=1}^n \left(\sum_{i=1}^n v' \mathbf{w}_{iq} C_{i\ell q} M_{im,1} M_{ik,2} \right)^2 = O \left(\mathbb{V}[v'\hat{\mathbf{b}}_q] \lambda_{\max}(C_q^2) \right) \quad \square \end{aligned}$$

C.5.1 Conservative Variance Estimation

The standard error estimators considered in the preceding two lemmas relied on existence of the independent and unbiased estimators $\widehat{x'_i \beta}_{-i,1}$ and $\widehat{x'_i \beta}_{-i,2}$. This part of the appendix creates an adjustment for observations where these estimators do not exist. The adjustment ensures that one

can obtain valid inference as stated in the lemma at the end of the subsection.

For observations where it is not possible to create $\widehat{x'_i\beta_{-i,1}}$ and $\widehat{x'_i\beta_{-i,2}}$, we construct $\widehat{x'_i\beta_{-i,1}}$ to satisfy the requirements in Lemma 6 and set $P_{i\ell,2} = 0$ for all ℓ so that $\widehat{x'_i\beta_{-i,2}} = 0$. Then we define $\mathcal{Q}_i = 1_{\{\max_\ell P_{i\ell,2}^2 = 0\}}$ as an indicator that $\widehat{x'_i\beta_{-i,2}}$ could not be constructed as an unbiased estimator.

Based on this we let

$$\widehat{\mathbb{V}}_2[\widehat{\theta}] = 4 \sum_{i=1}^n \left(\sum_{\ell \neq i} C_{i\ell} y_\ell \right)^2 \tilde{\sigma}_{i,2}^2 - 2 \sum_{i=1}^n \sum_{\ell \neq i} \tilde{C}_{i\ell} \widehat{\sigma}_i^2 \widehat{\sigma}_{\ell,2}^2$$

where $\tilde{\sigma}_{i,2}^2 = (1 - \mathcal{Q}_i) \tilde{\sigma}_i^2 + \mathcal{Q}_i (y_i - \bar{y})^2$ and

$$\widehat{\sigma}_i^2 \widehat{\sigma}_{\ell,2}^2 = \begin{cases} \hat{\sigma}_{i,-\ell}^2 \cdot \hat{\sigma}_{\ell,-i}^2, & \text{if } P_{ik,-\ell} P_{lk,-i} = 0 \text{ for all } k \text{ and } \mathcal{Q}_{i\ell} = \mathcal{Q}_{\ell i} = 0 \\ \tilde{\sigma}_i^2 \cdot \tilde{\sigma}_{\ell,-i}^2, & \text{else if } P_{i\ell,1} + P_{i\ell,2} = 0 \text{ and } \mathcal{Q}_i = \mathcal{Q}_{\ell i} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot \tilde{\sigma}_\ell^2, & \text{else if } P_{\ell i,1} + P_{\ell i,2} = 0 \text{ and } \mathcal{Q}_\ell = \mathcal{Q}_{i\ell} = 0, \\ \hat{\sigma}_{i,-\ell}^2 \cdot (y_\ell - \bar{y})^2 \cdot 1_{\{\tilde{C}_{i\ell} < 0\}}, & \text{else if } \mathcal{Q}_{i\ell} = 0, \\ (y_i - \bar{y})^2 \cdot \hat{\sigma}_{\ell,-i}^2 \cdot 1_{\{\tilde{C}_{i\ell} < 0\}}, & \text{else if } \mathcal{Q}_{\ell i} = 0, \\ (y_i - \bar{y})^2 \cdot (y_\ell - \bar{y})^2 \cdot 1_{\{\tilde{C}_{i\ell} < 0\}}, & \text{otherwise} \end{cases}$$

where we let $\mathcal{Q}_{i\ell} = 1_{\{P_{i\ell,1} \neq 0 \neq \mathcal{Q}_i\}}$. The definition of $\widehat{\mathbb{V}}_2[\widehat{\theta}]$ is such that $\widehat{\mathbb{V}}_2[\widehat{\theta}] = \widehat{\mathbb{V}}[\widehat{\theta}]$ when two independent unbiased estimators of $x'_i\beta$ can be formed for all observations, i.e., when $\mathcal{Q}_i = 0$ for all i .

Similarly, we let

$$\widehat{\Sigma}_{q,2} = \sum_{i=1}^n \begin{bmatrix} w_{iq} w'_{iq} \hat{\sigma}_{i,2}^2 & 2w_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right) \tilde{\sigma}_{i,2}^2 \\ 2w'_{iq} \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right) \tilde{\sigma}_{i,2}^2 & 4 \left(\sum_{\ell \neq i} C_{i\ell q} y_\ell \right)^2 \tilde{\sigma}_i^2 - 2 \sum_{\ell \neq i} \tilde{C}_{i\ell q} \widehat{\sigma}_i^2 \widehat{\sigma}_{\ell,2}^2 \end{bmatrix}$$

where $\hat{\sigma}_{i,2}^2 = (1 - \mathcal{Q}_i) \hat{\sigma}_i^2 + \mathcal{Q}_i (y_i - \bar{y})^2$ and $\widehat{\sigma}_i^2 \widehat{\sigma}_{\ell,2}^2$ is defined as $\widehat{\sigma}_i^2 \widehat{\sigma}_{\ell,2}^2$ but using $\tilde{C}_{i\ell q}$ instead of $\tilde{C}_{i\ell}$.

The following lemma shows that these estimators of the asymptotic variance leads to valid inference when coupled with the confidence intervals proposed in Sections 4 and 6.

Lemma C.9. *Suppose that $\sum_{\ell \neq i} P_{i\ell,1} x'_\ell \beta = x'_i \beta$, either $\sum_{\ell \neq i} P_{i\ell,2} x'_\ell \beta = x'_i \beta$ or $\max_\ell P_{i\ell,2}^2 = 0$, $P_{i\ell,1} P_{i\ell,2} = 0$ for all ℓ , and $\lambda_{\max}(P_s P'_s) = O(1)$ where $P_s = (P_{i\ell, s})_{i,\ell}$.*

1. *If the conditions of Theorem 2 hold, then $\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \left[\widehat{\theta} \pm z_\alpha \widehat{\mathbb{V}}_2[\widehat{\theta}]^{1/2} \right] \right) \geq 1 - \alpha$.*
2. *If the conditions of Theorem 3 hold, then $\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in C_\alpha^\theta(\widehat{\Sigma}_{q,2}) \right) \geq 1 - \alpha$.*

The following provides a proof of the first claim of this lemma, while we postpone a proof of the second claim to the end of Appendix C.6.

Proof. As in the proof of Lemma 5 it suffices to show that $\hat{\mathbb{V}}_2[\hat{\theta}]$ has a positive bias in large samples and that $\hat{\mathbb{V}}_2[\hat{\theta}] - \mathbb{E}[\hat{\mathbb{V}}_2[\hat{\theta}]]$ is $o_p(\mathbb{V}[\hat{\theta}])$. The second claim involves no new arguments relative to the proof of Lemma 5 and is therefore omitted. Thus we briefly report the positive bias in $\hat{\mathbb{V}}_2[\hat{\theta}]$.

We have that

$$\begin{aligned} \mathbb{E} [\hat{\mathbb{V}}_2[\hat{\theta}]] &= \mathbb{V}[\hat{\theta}] + 4 \sum_{i:Q_i=1} \left(\sum_{\ell \neq i} C_{i\ell} x'_\ell \beta \right)^2 ((x_i - \bar{x})' \beta)^2 \\ &\quad + 2 \sum_{(i,\ell) \in \mathcal{B}_1} \tilde{C}_{i\ell} \sigma_i^2 \left(\sigma_\ell^2 1_{\{\tilde{C}_{i\ell} > 0\}} + ((x_\ell - \bar{x})' \beta)^2 1_{\{\tilde{C}_{i\ell} < 0\}} \right) \\ &\quad + 2 \sum_{(i,\ell) \in \mathcal{B}_2} \tilde{C}_{i\ell} \sigma_\ell^2 \left(\sigma_i^2 1_{\{\tilde{C}_{i\ell} > 0\}} + ((x_i - \bar{x})' \beta)^2 1_{\{\tilde{C}_{i\ell} < 0\}} \right) \\ &\quad + 2 \sum_{(i,\ell) \in \mathcal{B}_3} \tilde{C}_{i\ell} \left(\sigma_i^2 \sigma_\ell^2 1_{\{\tilde{C}_{i\ell} > 0\}} + \left(2\sigma_i^2 ((x_\ell - \bar{x})' \beta)^2 + ((x_i - \bar{x})' \beta (x_\ell - \bar{x})' \beta)^2 \right) 1_{\{\tilde{C}_{i\ell} < 0\}} \right) \\ &\quad + O\left(\frac{1}{n} \mathbb{V}[\hat{\theta}]\right) \end{aligned}$$

where the remainder stems from estimation of \bar{y} and $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ refers to pairs of observations that fall in each of the three last cases in the definition of $\sigma_i^2 \sigma_\ell^2$. \square

C.6 Inference with Nuisance Parameters

This Appendix starts by defining curvature and accompanying critical value for a given curvature as introduced in Section 6. Then it derives the closed form representation of $C_\alpha^\theta(\tilde{\Sigma}_1)$ for any variance matrix $\tilde{\Sigma}_1 \in \mathbb{R}^{2 \times 2}$ where for general q we have

$$C_\alpha^\theta(\tilde{\Sigma}_q) = \left[\min_{(\hat{b}_1, \dots, \hat{b}_q, \hat{\theta}_q)' \in \mathbb{E}_\alpha(\tilde{\Sigma}_q)} \sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \hat{\theta}_q, \max_{(\hat{b}_1, \dots, \hat{b}_q, \hat{\theta}_q)' \in \mathbb{E}_\alpha(\tilde{\Sigma}_q)} \sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \hat{\theta}_q \right]$$

and

$$\mathbb{E}_\alpha(\tilde{\Sigma}_q) = \left\{ (\mathbf{b}'_q, \theta_q)' \in \mathbb{R}^{q+1} : \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}' \tilde{\Sigma}_q^{-1} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix} \leq z_{\alpha, \kappa(\tilde{\Sigma}_q)}^2 \right\}.$$

Finally, it proofs validity of $\hat{C}_{\alpha,q}^\theta = C_\alpha^\theta(\hat{\Sigma}_q)$ and $C_\alpha^\theta(\hat{\Sigma}_{q,2})$ for any fixed q . As for $\hat{\Sigma}_q$ and $\hat{\Sigma}_{q,2}$, we partition $\tilde{\Sigma}_q$ into $\tilde{\Sigma}_q = \begin{bmatrix} \tilde{\mathbb{V}}[\hat{\mathbf{b}}_q] & \tilde{\mathbb{C}}[\hat{\mathbf{b}}_q, \hat{\theta}_q]' \\ \tilde{\mathbb{C}}[\hat{\mathbf{b}}_q, \hat{\theta}_q] & \tilde{\mathbb{V}}[\hat{\theta}_q] \end{bmatrix}$ with $\tilde{\mathbb{V}}[\hat{\theta}_q] \in \mathbb{R}$. In Section 6, $\hat{C}_{\alpha,q}^\theta = C_\alpha^\theta(\hat{\Sigma}_q)$, $\hat{\mathbb{E}}_{\alpha,q} = \mathbb{E}_\alpha(\hat{\Sigma}_q)$, and $\hat{\kappa}_q = \kappa(\hat{\Sigma}_q)$.

C.6.1 Preliminaries

Critical value function For a given curvature $\kappa > 0$ and confidence level $1 - \alpha$, the critical value function $z_{\alpha, \kappa}$ is the $(1 - \alpha)$ 'th quantile of

$$\rho(\chi_q, \chi_1, \kappa) = \sqrt{\chi_q^2 + \left(\chi_1 + \frac{1}{\kappa}\right)^2} - \frac{1}{\kappa}$$

where χ_q^2 and χ_1^2 are independently distributed variates from the χ -squared distribution with q and 1 degrees of freedom, respectively. $\rho(\chi_q, \chi_1, \kappa)$ is the Euclidean distance from (χ_q, χ_1) to the circle with center $(0, -\frac{1}{\kappa})$ and radius $\frac{1}{\kappa}$. The critical value function at $\kappa = 0$ is the limit of $z_{\alpha, \kappa}$ as $\kappa \downarrow 0$, which is the $(1 - \alpha)$ 'th quantile of a central χ_1^2 random variable. See [Andrews and Mikusheva \(2016\)](#) for additional details.

Curvature The confidence interval $C_\alpha^\theta(\tilde{\Sigma}_q)$ inverts hypotheses of the type $H_0 : \theta = c$ versus $H_1 : \theta \neq c$ based on the value of the test statistic

$$\min_{\mathbf{b}_q, \theta_q : g(\mathbf{b}_q, \theta_q, c) = 0} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}' \tilde{\Sigma}_q^{-1} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}$$

where $g(\mathbf{b}_q, \theta_q, c) = \sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \theta_q - c$ and $\mathbf{b}_q = (\hat{b}_1, \dots, \hat{b}_q)'$. This testing problem depends on the manifold $S = \{x = \tilde{\Sigma}_q^{-1/2}(\mathbf{b}_q, \theta_q)' : g(\mathbf{b}_q, \theta_q, c) = 0\}$ for which we need an upper bound on the maximal curvature. We derive this upper bound using the parameterization $\mathbf{x}(\dot{y}) = \tilde{\Sigma}_q^{-1/2}(\dot{y}_1, \dots, \dot{y}_q, c - \sum_{\ell=1}^q \lambda_\ell \dot{y}_\ell^2)'$ which maps from \mathbb{R}^q to S , is a homeomorphism, and has a Jacobian of full rank:

$$d\mathbf{x}(\dot{y}) = \tilde{\Sigma}_q^{-1/2} \begin{bmatrix} \text{diag}(1, \dots, 1) \\ -2\lambda_1 \dot{y}_1, \dots, -2\lambda_q \dot{y}_q \end{bmatrix}$$

The maximal curvature of S , $\kappa(\tilde{\Sigma}_q)$, is then given as $\kappa(\tilde{\Sigma}_q) = \max_{\dot{y} \in \mathbb{R}^q} \kappa_{\dot{y}}$ where

$$\kappa_{\dot{y}} = \sup_{u \in \mathbb{R}^q} \frac{\|(I - P_{\dot{y}})V(u \odot u)\|}{\|d\mathbf{x}(\dot{y})u\|^2}, \quad V = \tilde{\Sigma}_q^{-1/2} \begin{bmatrix} 0 \\ -2\lambda_1, \dots, -2\lambda_q \end{bmatrix},$$

and $P_{\dot{y}} = d\mathbf{x}(\dot{y})(d\mathbf{x}(\dot{y})'d\mathbf{x}(\dot{y}))^{-1}d\mathbf{x}(\dot{y})'$. See [Andrews and Mikusheva \(2016\)](#) for additional details.

Curvature when $q = 1$ In this case the maximization over u drops out and we have

$$\kappa(\tilde{\Sigma}_1) = \max_{\dot{y} \in \mathbb{R}} \frac{\sqrt{V'V - \frac{(v'V)^2}{v'v}}}{v'v} \quad \text{where } v = \tilde{\Sigma}_1^{-1/2}(1, -2\lambda_1 \dot{y})'$$

and $V = \tilde{\Sigma}_1^{-1/2}(0, -2\lambda_1)$. The value $\dot{y}^* = -\frac{\tilde{\rho}\tilde{\mathbb{V}}[\hat{\theta}_q]}{2\lambda_1\tilde{\mathbb{V}}[\hat{b}_1]}$ for $\tilde{\rho} = \frac{\tilde{\mathbb{C}}[\hat{b}_1, \hat{\theta}_q]}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}\tilde{\mathbb{V}}[\hat{\theta}_q]^{1/2}}$ is both a minimizer of $v'v$ and $(v'V)^2$, so we obtain that $\kappa(\tilde{\Sigma}_1) = \frac{2|\lambda_1|\tilde{\mathbb{V}}[\hat{b}_1]}{\tilde{\mathbb{V}}[\hat{\theta}_q]^{1/2}(1-\tilde{\rho}^2)^{1/2}}$.

Curvature when $q > 1$ In this case we first maximize over \dot{y} and then over u . For a fixed u we want to find

$$\max_{\dot{y} \in \mathbb{R}^q} \frac{\sqrt{V_u'V_u - V_u'P_{\dot{y}}V_u}}{v_{u,\dot{y}}'v_{u,\dot{y}}}, \quad \text{where } V_u = \tilde{\Sigma}_q^{-1/2}(0, -2\sum_{\ell=1}^q \lambda_\ell u_\ell^2), \quad v_{u,\dot{y}} = \tilde{\Sigma}_q^{-1/2}(u', -2u'D_q\dot{y})',$$

and $D_q = \text{diag}(\lambda_1, \dots, \lambda_q)$. The value for \dot{y} that solves $-2D_q\dot{y} = \tilde{\mathbb{V}}[\hat{b}_q]^{-1}\tilde{\mathbb{C}}[\hat{b}_q, \hat{\theta}_q]$ sets $P_{\dot{y}}V_u = 0$ and minimizes $v_{u,\dot{y}}'v_{u,\dot{y}}$. Thus we obtain

$$\kappa(\tilde{\Sigma}_q) = \frac{2 \max_{u \in \mathbb{R}^q} \frac{|u'D_q u|}{u'\tilde{\mathbb{V}}[\hat{b}_q]^{-1}u}}{\left(\tilde{\mathbb{V}}[\hat{\theta}_q] - \tilde{\mathbb{C}}[\hat{b}_q, \hat{\theta}_q]' \tilde{\mathbb{V}}[\hat{b}_q]^{-1} \tilde{\mathbb{C}}[\hat{b}_q, \hat{\theta}_q]\right)^{1/2}} = \frac{2|\dot{\lambda}_1(\tilde{\mathbb{V}}[\hat{b}_q]^{1/2}D_q\tilde{\mathbb{V}}[\hat{b}_q]^{1/2})|}{\left(\tilde{\mathbb{V}}[\hat{\theta}_q] - \tilde{\mathbb{C}}[\hat{b}_q, \hat{\theta}_q]' \tilde{\mathbb{V}}[\hat{b}_q]^{-1} \tilde{\mathbb{C}}[\hat{b}_q, \hat{\theta}_q]\right)^{1/2}}$$

where $\dot{\lambda}_1(\cdot)$ is the eigenvalue of largest magnitude. This formula simplifies to the one derived above when $q = 1$.

C.6.2 Closed Form Representation of $C_\alpha^\theta(\tilde{\Sigma}_1)$

An implicit representation of $C_\alpha^\theta(\tilde{\Sigma}_1)$ is

$$C_\alpha^\theta(\tilde{\Sigma}_1) = \left[\lambda_1 b_{1,-}^2 + \theta_{1,-}, \lambda_1 b_{1,+}^2 + \theta_{1,+} \right]$$

where $b_{1,\pm}$ and $\theta_{1,\pm}$ are solutions to

$$b_{1,\pm} = \hat{b}_1 \pm z_{\alpha, \kappa(\tilde{\Sigma}_1)} \left(\tilde{\mathbb{V}}[\hat{b}_1](1 - \tilde{a}(b_{1,\pm})) \right)^{1/2}, \quad (15)$$

$$\theta_{1,\pm} = \hat{\theta}_1 - \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} (\hat{b}_1 - b_{1,\pm}) \pm z_{\alpha, \kappa(\tilde{\Sigma}_1)} \left(\tilde{\mathbb{V}}[\hat{\theta}_1](1 - \tilde{\rho}^2) \tilde{a}(b_{1,\pm}) \right)^{1/2} \quad (16)$$

for $\tilde{a}(\hat{b}_1) = \left(1 + \left(\frac{\text{sgn}(\lambda_1)\kappa(\tilde{\Sigma}_1)\hat{b}_1}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} + \frac{\tilde{\rho}}{\sqrt{1-\tilde{\rho}^2}} \right)^2 \right)^{-1}$.

This construction is fairly intuitive. When $\tilde{\rho} = 0$, the interval has endpoints that combine

$$\lambda_1 \left(\hat{b}_1 \pm z_{\alpha, \kappa(\tilde{\Sigma}_1)} \left(\tilde{\mathbb{V}}[\hat{b}_1](1 - \tilde{a}(b_{1,\pm})) \right)^{1/2} \right)^2 \quad \text{and} \quad \hat{\theta}_q \pm z_{\alpha, \kappa(\tilde{\Sigma}_1)} \left(\tilde{\mathbb{V}}[\hat{\theta}_q] a(b_{1,\pm}) \right)^{1/2}$$

where $a(\hat{b}_1)$ estimates the fraction of $\mathbb{V}[\hat{\theta}]$ that stems from $\hat{\theta}_1$ when $\mathbb{E}[\hat{b}_1] = \hat{b}_1$. When $\tilde{\rho}$ is non-zero, $C_\alpha^\theta(\tilde{\Sigma}_1)$ involves an additional rotation of $(\hat{b}_1, \hat{\theta}_1)'$. This representation of $C_\alpha^\theta(\tilde{\Sigma}_1)$ is however not

unique as (15),(16) can have multiple solutions. Thus we derive the representation above together with an additional side condition that ensures uniqueness and represents $b_{1,\pm}$ and $\theta_{1,\pm}$ as solutions to a fourth order polynomial.

Derivation The upper end of $C_\alpha^\theta(\tilde{\Sigma}_1)$ is found by noting that maximization over a linear function in θ_1 implies that the constraint must bind at the maximum. Thus we can reformulate the bivariate problem as a univariate problem

$$\max_{(\hat{b}_1, \hat{\theta}_1) \in E_\alpha(\tilde{\Sigma}_1)} \lambda_1 \hat{b}_1^2 + \hat{\theta}_1 = \max_{\hat{b}_1} \lambda_1 \hat{b}_1^2 + \hat{\theta}_1 - \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} (\hat{b}_1 - \hat{b}_1) + \sqrt{\tilde{\mathbb{V}}[\hat{\theta}_1](1 - \tilde{\rho}^2)} \left(z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2 - \frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]} \right)$$

where we are implicitly enforcing the constraint on \hat{b}_1 that the term under the square-root is non-negative. Thus we will find a global maximum in \hat{b}_1 and note that it satisfies this constraint. The first order condition for a maximum is

$$2\lambda_1 \hat{b}_1 + \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} + \frac{\hat{b}_1 - \hat{b}_1}{\tilde{\mathbb{V}}[\hat{b}_1]} \sqrt{\frac{\tilde{\mathbb{V}}[\hat{\theta}_1](1 - \tilde{\rho}^2)}{z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2 - \frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]}}} = 0$$

which after a rearrangement and squaring of both sides yields $\frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]} = (1 - a(\hat{b}_1))z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2$. This in turn leads to the representation of $b_{1,\pm}$ given in (15). All solutions to this equation satisfies the implicit non-negativity constraint since any solution \hat{b} satisfies

$$z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2 - \frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]} = a(\hat{b}_1)z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2 > 0.$$

A slightly different arrangement of the first order condition reveals the equivalent quartic condition

$$\frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]} \left(1 + \left(\frac{\text{sgn}(\lambda_1)\kappa(\tilde{\Sigma}_1)\hat{b}_1}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} + \frac{\tilde{\rho}}{\sqrt{1 - \tilde{\rho}^2}} \right)^2 \right) = \left(\frac{\text{sgn}(\lambda_1)\kappa(\tilde{\Sigma}_1)\hat{b}_1}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} + \frac{\tilde{\rho}}{\sqrt{1 - \tilde{\rho}^2}} \right)^2 z_{\alpha, \kappa(\tilde{\Sigma}_1)}^2 \quad (17)$$

which has at most four solutions that are given on closed form. Thus the solution $b_{1,+}$ can be found as the maximizer of

$$\lambda_1 \hat{b}_1^2 + \hat{\theta}_1 - \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} (\hat{b}_1 - \hat{b}_1) + z_{\alpha, \kappa(\tilde{\Sigma}_1)} \left(\tilde{\mathbb{V}}[\hat{\theta}_1] a(\hat{b}_1) \right)^{1/2}$$

among the at most four solutions to (17). More importantly, the maximum is the upper end of

$C_\alpha^\theta(\tilde{\Sigma}_1)$. Now, for the minimization problem we instead have

$$\min_{(\hat{b}_1, \hat{\theta}_1) \in E_\alpha(\tilde{\Sigma}_1)} \lambda_1 \hat{b}_1^2 + \hat{\theta}_1 = \min_{\hat{b}_1} \lambda_1 \hat{b}_1^2 + \hat{\theta}_1 - \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} (\hat{b}_1 - \hat{b}_1) - \sqrt{\tilde{\mathbb{V}}[\hat{\theta}_1] (1 - \tilde{\rho}^2)} \left(z_{\alpha, \kappa}^2(\tilde{\Sigma}_1) - \frac{(\hat{b}_1 - \hat{b}_1)^2}{\tilde{\mathbb{V}}[\hat{b}_1]} \right)$$

which when rearranging and squaring the first order condition again leads to (17) as a necessary condition for a minimum. Thus $b_{1,-}$ and the lower end of $C_\alpha^\theta(\tilde{\Sigma}_1)$ can be found by minimizing

$$\lambda_1 \hat{b}_1^2 + \hat{\theta}_1 - \tilde{\rho} \frac{\tilde{\mathbb{V}}[\hat{\theta}_1]^{1/2}}{\tilde{\mathbb{V}}[\hat{b}_1]^{1/2}} (\hat{b}_1 - \hat{b}_1) - z_{\alpha, \kappa}(\tilde{\Sigma}_1) \left(\tilde{\mathbb{V}}[\hat{\theta}_1] a(\hat{b}_1) \right)^{1/2}$$

over the at most four solutions to (17).

C.6.3 Asymptotic Validity

Lemma C.10. *If $\Sigma_q^{-1} \hat{\Sigma}_q \xrightarrow{p} I_{q+1}$ and the conditions of Theorem 3 hold, then*

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \hat{C}_{\alpha, q}^\theta \right) \geq 1 - \alpha.$$

Proof. The following two conditions are the inputs to the proof of Theorem 2 in Andrews and Mikusheva (2016), from which it follows that

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left(\theta \in \hat{C}_{\alpha, q}^\theta \right) = \liminf_{n \rightarrow \infty} \mathbb{P} \left(\min_{(\hat{\mathbf{b}}'_q, \hat{\theta}_q)' : g(\hat{\mathbf{b}}_q, \theta_q, \theta) = 0} \left(\hat{\mathbf{b}}_q - \mathbf{b}_q \right)' \hat{\Sigma}_q^{-1} \left(\hat{\mathbf{b}}_q - \mathbf{b}_q \right) \leq z_{\alpha, \hat{\kappa}_q}^2 \right) \geq 1 - \alpha$$

where $g(\mathbf{b}_q, \theta_q, \theta) = \sum_{\ell=1}^q \lambda_\ell \hat{b}_\ell^2 + \theta_q - \theta$ and $\mathbf{b}_q = (\hat{b}_1, \dots, \hat{b}_q)'$.

Condition (i) requires that $\hat{\Sigma}_q^{-1/2} \left((\hat{\mathbf{b}}'_q, \hat{\theta}_q)' - \mathbb{E}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] \right) \xrightarrow{d} \mathcal{N}(0, I_{q+1})$, which follows from Theorem 3 and $\Sigma_q^{-1} \hat{\Sigma}_q \xrightarrow{p} I_{q+1}$.

Condition (ii) is satisfied if the conditions of Lemma 1 in Andrews and Mikusheva (2016) are satisfied. To verify this, take the manifold

$$\tilde{S} = \left\{ \dot{x} \in \mathbb{R}^{q+1} : \tilde{g}(\dot{x}) = 0 \right\}$$

for

$$\tilde{g}(\dot{x}) = \dot{x}' \hat{\Sigma}_q^{1/2} \begin{bmatrix} D_q & 0 \\ 0 & 0 \end{bmatrix} \hat{\Sigma}_q^{1/2} \dot{x} + (2\mathbb{E}[\hat{\mathbf{b}}_q'], 1) \begin{bmatrix} D_q & 0 \\ 0 & 1 \end{bmatrix} \hat{\Sigma}_q^{1/2} \dot{x}.$$

The curvature of \tilde{S} is $\hat{\kappa}$, $\tilde{g}(0) = 0$, and \tilde{g} is continuously differentiable with a Jacobian of rank 1. These are the conditions of Lemma 1 in Andrews and Mikusheva (2016). \square

Proof of the second claims in Lemmas C.8 and C.9. The proof contains two main parts. One part

is to establish that the biases of $\hat{\Sigma}_q$ and $\hat{\Sigma}_{q,2}$ are positive semidefinite in large samples, and that $\mathbb{E}[\hat{\Sigma}_q]^{-1}\hat{\Sigma}_q - I_{q+1}$ and $\mathbb{E}[\hat{\Sigma}_{q,2}]^{-1}\hat{\Sigma}_{q,2} - I_{q+1}$ are $o_p(1)$. These arguments are analogues to those presented in the proofs of Lemmas C.8 and C.9 and are therefore only sketched. The other part is to show that this positive semidefinite asymptotic bias in the variance estimator does not alter the validity of the confidence interval based on it. We only cover $\hat{\Sigma}_{q,2}$ as that estimator simplifies to $\hat{\Sigma}_q$ when the design is sufficiently well-behaved.

Validity First, we let QDQ' be the spectral decomposition of $\mathbb{E}[\hat{\Sigma}_{q,2}]^{-1/2}\Sigma_q\mathbb{E}[\hat{\Sigma}_{q,2}]^{-1/2}$. Here, $\text{QQ}' = \text{Q}'\text{Q} = I_{q+1}$ and all diagonal entries in the diagonal matrix D belongs to $(0, 1]$ in large samples. Now,

$$\mathbb{P}\left(\theta \in C_\alpha^\theta(\hat{\Sigma}_{q,2})\right) = \mathbb{P}\left(\min_{(\hat{\mathbf{b}}'_q, \hat{\theta}_q)': g(\hat{\mathbf{b}}_q, \hat{\theta}_q, \theta)=0} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}' \mathbb{E}[\hat{\Sigma}_{q,2}]^{-1} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix} \leq z_{\alpha, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}])}^2\right) + o(1)$$

where the minimum distance statistic above satisfies

$$\min_{(\hat{\mathbf{b}}'_q, \hat{\theta}_q)': g(\hat{\mathbf{b}}_q, \hat{\theta}_q, \theta)=0} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix}' \mathbb{E}[\hat{\Sigma}_{q,2}]^{-1} \begin{pmatrix} \hat{\mathbf{b}}_q - \mathbf{b}_q \\ \hat{\theta}_q - \theta_q \end{pmatrix} = \min_{x \in S_2} (\xi - x)'(\xi - x)$$

where $S_2 = \{x : x = \text{Q}'\mathbb{E}[\hat{\Sigma}_{q,2}]^{-1/2} \left((\hat{\mathbf{b}}'_q, \hat{\theta}_q)' - \mathbb{E}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] \right), g(\hat{\mathbf{b}}_q, \hat{\theta}_q, \theta) = 0\}$ and the random vector $\xi = \text{Q}'\mathbb{E}[\hat{\Sigma}_{q,2}]^{-1/2} \left((\hat{\mathbf{b}}'_q, \hat{\theta}_q)' - \mathbb{E}[(\hat{\mathbf{b}}'_q, \hat{\theta}_q)'] \right)$ has the property that $\text{D}^{-1/2}\xi \xrightarrow{d} \mathcal{N}(0, I_{q+1})$. From the geometric consideration in Andrews and Mikusheva (2016) it follows that S_2 has curvature of $\kappa(\mathbb{E}[\hat{\Sigma}_{q,2}])$ since curvature is invariant to rotations. Furthermore,

$$\begin{aligned} \min_{x \in S_2} (\xi - x)'(\xi - x) &\leq \rho^2 \left(\|\xi_{-1}\|, |\xi_1|, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}]) \right) \\ &\leq \rho^2 \left(\|(\text{D}^{-1/2}\xi)_{-1}\|, |(\text{D}^{-1/2}\xi)_1|, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}]) \right) \end{aligned}$$

where $\xi = (\xi_1, \xi'_{-1})'$ and $\text{D}^{-1/2}\xi = ((\text{D}^{-1/2}\xi)_1, (\text{D}^{-1/2}\xi)'_{-1})$ and the first inequality follows from the proof of Theorem 1 in Andrews and Mikusheva (2016). Thus

$$\begin{aligned} \liminf_{n \rightarrow \infty} \mathbb{P}\left(\theta \in C_\alpha^\theta(\hat{\Sigma}_{q,2})\right) &= \liminf_{n \rightarrow \infty} \mathbb{P}\left(\min_{x \in S_2} (\xi - x)'(\xi - x) \leq z_{\alpha, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}])}^2\right) \\ &\geq \liminf_{n \rightarrow \infty} \mathbb{P}\left(\rho^2 \left(\chi_q, \chi_1, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}]) \right) \leq z_{\alpha, \kappa(\mathbb{E}[\hat{\Sigma}_{q,2}])}^2\right) = 1 - \alpha \end{aligned}$$

since $(\|\xi_{-1}\|, |\xi_1|) \xrightarrow{d} (\chi_q, \chi_1)$.

Bias and variability in $\hat{\Sigma}_{q,2}$ We finish by reporting the positive semidefinite bias in $\hat{\Sigma}_{q,2}$. We

have that

$$\mathbb{E} \left[\hat{\Sigma}_{q,2} \right] = \Sigma_q + \sum_{i:Q_i=1} \sigma_i^2 \begin{pmatrix} \mathbf{w}_{iq} \\ 2 \sum_{\ell \neq i} C_{i\ell} x'_\ell \beta \end{pmatrix} \begin{pmatrix} \mathbf{w}_{iq} \\ 2 \sum_{\ell \neq i} C_{i\ell} x'_\ell \beta \end{pmatrix}' + \begin{bmatrix} 0 & 0 \\ 0 & \mathcal{B} \end{bmatrix} + O \left(\frac{1}{n} \mathbb{V}[\hat{\theta}] \right)$$

where

$$\begin{aligned} \mathcal{B} &= 2 \sum_{(i,\ell) \in \mathcal{B}_1} \tilde{C}_{i\ell q} \sigma_i^2 \left(\sigma_\ell^2 1_{\{\tilde{C}_{i\ell q} > 0\}} + ((x_\ell - \bar{x})' \beta)^2 1_{\{\tilde{C}_{i\ell q} < 0\}} \right) \\ &+ 2 \sum_{(i,\ell) \in \mathcal{B}_2} \tilde{C}_{i\ell q} \sigma_\ell^2 \left(\sigma_i^2 1_{\{\tilde{C}_{i\ell q} > 0\}} + ((x_i - \bar{x})' \beta)^2 1_{\{\tilde{C}_{i\ell q} < 0\}} \right) \\ &+ 2 \sum_{(i,\ell) \in \mathcal{B}_3} \tilde{C}_{i\ell q} \left(\sigma_i^2 \sigma_\ell^2 1_{\{\tilde{C}_{i\ell q} > 0\}} + \left(2\sigma_i^2 ((x_\ell - \bar{x})' \beta)^2 + ((x_i - \bar{x})' \beta (x_\ell - \bar{x})' \beta)^2 \right) 1_{\{\tilde{C}_{i\ell q} < 0\}} \right) \end{aligned}$$

for \mathcal{B}_1 , \mathcal{B}_2 , \mathcal{B}_3 referring to pairs of observations that fall in each of the three last cases in the definition of $\sigma_i^2 \sigma_\ell^2$. \square

C.7 Verifying Conditions

Example 1. The only non-immediate conclusions are that:

$$\begin{aligned} \mathbb{V}[\hat{\theta}]^{-1} \max_i (\tilde{x}'_i \beta)^2 &= O \left(\frac{\max_i (x'_i \beta)^2 / n^2}{\min_i \sigma_i^2 \text{trace}(\tilde{A}^2)} \right) = O \left(\frac{\max_i (x'_i \beta)^2}{r} \right) \\ \mathbb{V}[\hat{\theta}]^{-1} \max_i (\tilde{x}'_i \beta)^2 &= O \left(\frac{\max_{i,j} M_{jj}^{-2} (P_{jj} - \frac{1}{n})^2 (x'_j \beta)^2 (\sum_{\ell=1}^n |M_{i\ell}|)^2 / n^2}{\min_i \sigma_i^2 \text{trace}(\tilde{A}^2)} \right) \\ &= O \left(\frac{\max_{i,j} (x'_j \beta)^2 (\sum_{\ell=1}^n |M_{i\ell}|)^2}{r} \right). \end{aligned}$$

Example 2. We first derive the representations of $\hat{\sigma}_\alpha^2$ given in section 2. When there are no common regressors, the representation in (5) follows from $B_{ii} = \frac{1}{nT_{g(i)}} \left(1 - \frac{T_{g(i)}}{n} \right)$ and

$$\hat{\sigma}_g^2 = \frac{1}{T_g} \sum_{t=1}^{T_g} y_{gt} \left(y_{gt} - \frac{1}{T_g - 1} \sum_{s \neq t} y_{gs} \right) = \frac{1}{T_g} \sum_{i:g(i)=g} \hat{\sigma}_i^2$$

which yields that

$$\sum_{i=1}^n B_{ii} \hat{\sigma}_i^2 = \frac{1}{n} \sum_{g=1}^N \left(1 - \frac{T_g}{n} \right) \hat{\sigma}_g^2.$$

With common regressors, it follows from the formula for block inversion of matrices that

$$\begin{aligned}\tilde{X}' &= AS_{xx}^{-1} \begin{bmatrix} D' \\ X' \end{bmatrix} = \frac{1}{n} \begin{bmatrix} (D' - \bar{d}\mathbf{1}'_n) \left(I - X (X'(I - P_D)X')^{-1} X'(I - P_D) \right) \\ 0 \end{bmatrix} \\ &= \frac{1}{n} \begin{bmatrix} D' - \bar{d}\mathbf{1}'_n - \hat{\Gamma}' X'(I - P_D) \\ 0 \end{bmatrix}\end{aligned}$$

where $D = (d_1, \dots, d_n)'$, $X = (x_{g(1)t(1)}, \dots, x_{g(n)t(n)})'$, $P_D = DS_{dd}^{-1}D'$, $\mathbf{1}_n = (1, \dots, 1)'$, and $S_{dd} = D'D$. Thus it follows that

$$\tilde{x}_i = \frac{1}{n} \begin{pmatrix} d_i - \bar{d} - \hat{\Gamma}'(x_{g(i)t(i)} - \bar{x}_{g(i)}) \\ 0 \end{pmatrix}.$$

The no common regressors claims are immediate. With common regressors we have

$$P_{i\ell} = T_{g(i)}^{-1} \mathbf{1}_{\{g(i)=g(\ell)\}} + n^{-1}(x_{g(i)t(i)} - \bar{x}_{g(i)})' W^{-1} (x_{g(\ell)t(\ell)} - \bar{x}_{g(\ell)}) = T_{g(i)}^{-1} \mathbf{1}_{\{i=\ell\}} + O(n^{-1})$$

where $W = \frac{1}{n} \sum_{g=1}^N \sum_{t=1}^T (x_{gt} - \bar{x}_g)(x_{gt} - \bar{x}_g)'$ so $P_{ii} \leq C < 1$ in large samples. The eigenvalues of \tilde{A} are equal to the eigenvalues of

$$\frac{1}{n} \left(I_N - nS_{dd}^{-1/2} \bar{d}\bar{d}' S_{dd}^{-1/2} \right) \left(I_N + \frac{1}{n} S_{dd}^{1/2} D' X W^{-1} X' D S_{dd}^{-1/2} \right)$$

which in turn satisfies that $\frac{c_1}{n} \leq \lambda_\ell \leq \frac{c_2}{n}$ for $\ell = 1, \dots, N-1$ and $c_2 \geq c_1 > 0$ not depending on n . $w'_i w_i = O(P_{ii})$ so Theorem 1 applies when N is fixed and $\min_g T_g \rightarrow \infty$. Finally,

$$\begin{aligned}\max_i \mathbb{V}[\hat{\theta}]^{-1} (\hat{x}'_i \beta)^2 &= O \left(\frac{\max_{g,t} \alpha_g^2 + \|x_{gt}\|^2 \frac{1}{n} \sum_{i=1}^n \|x_{g(i)t(i)}\|^2 \sigma_\alpha^2}{N} \right) \\ \max_i \mathbb{V}[\hat{\theta}]^{-1} (\hat{x}'_i \beta)^2 &= O \left(\frac{\max_{i,j} (x'_{ij} \beta)^2 (\sum_{\ell=1}^n |M_{i\ell}|)^2}{N} \right)\end{aligned}$$

and $\sum_{\ell=1}^n |M_{i\ell}| = O(1)$ so Theorem 2 applies when $N \rightarrow \infty$.

We finish this example with a setup where an unbalanced panel leads to a bias and inconsistency in $\hat{\theta}_{\text{HO}}$. Consider

$$y_{gt} = \alpha_g + \varepsilon_{gt} \quad (g = 1, \dots, N, t = 1, \dots, T_g)$$

where N is even, $(T_g = 2, \mathbb{E}[\varepsilon_{gt}^2] = 2\sigma^2)$ for $g \leq N/2$ and $(T_g = 3, \mathbb{E}[\varepsilon_{gt}^2] = \sigma^2)$ for $g > N/2$, and the

estimand is,

$$\theta = \frac{1}{n} \sum_{g=1}^N T_g \alpha_g^2 \quad \text{where } n = \sum_{g=1}^N T_g = \frac{5N}{2}.$$

Here we have that $\tilde{A} = I_N/n$ and $\text{trace}(\tilde{A}^2) = N/n^2 = o(1)$ as $n \rightarrow \infty$ so the leave-out estimator is consistent. Furthermore,

$$nB_{ii} = P_{ii} = \begin{cases} \frac{1}{2}, & \text{if } i \leq N, \\ \frac{1}{3}, & \text{otherwise,} \end{cases} \quad \sigma_i^2 = \begin{cases} 2\sigma^2, & \text{if } i \leq N, \\ \sigma^2, & \text{otherwise,} \end{cases}$$

so

$$\begin{aligned} \mathbb{E}[\tilde{\theta}] - \theta &= \sum_{i=1}^n B_{ii} \sigma_i^2 = \frac{\sigma^2}{n} \left(N + \frac{N}{2} \right) = \frac{3\sigma^2}{5}, \\ \mathbb{E}[\hat{\theta}_{\text{HO}}] - \theta &= \sigma_{nB_{ii}, \sigma_i^2} + S_B \frac{n}{n-N} \sigma_{P_{ii}, \sigma_i^2} = \frac{2\sigma^2}{50} + \frac{2}{3} \times \frac{2\sigma^2}{50} = \frac{\sigma^2}{15}. \end{aligned}$$

Example 3. \tilde{A} is diagonal with N diagonal entries of $\frac{1}{n} \frac{T_g}{S_{zz,g}}$, so $\lambda_g = \frac{1}{n} \frac{T_g}{S_{zz,g}}$ for $g = 1, \dots, N$. $\text{trace}(\tilde{A}^2) \leq \frac{\lambda_1}{\min_g S_{zz,g}} \frac{1}{n} \sum_{g=1}^N T_g = O(\lambda_1)$. $\max_i w'_i w_i = \max_{g,t} \frac{(z_{gt} - \bar{z}_g)^2}{S_{zz,g}} = o(1)$ when $\min_g S_{zz,g} \rightarrow \infty$. Furthermore, $\mathbb{V}[\hat{\theta}]^{-1} = O(\frac{n^2}{N})$, so

$$\mathbb{V}[\hat{\theta}]^{-1} \max_i (\tilde{x}'_i \beta)^2 = O\left(\max_{g,t} \frac{z_{gt}^2 \delta_g^2}{N S_{zz,g}} \right) = o(1),$$

and $M_{i\ell} = 0$ if $g(i) \neq g(\ell)$ so

$$\mathbb{V}[\hat{\theta}]^{-1} \max_i (\tilde{x}'_i \beta)^2 = O\left(\max_g \left(\frac{n \sum_{i:g(i)=g} B_{ii}}{\sqrt{N}} \right)^2 \right) = O\left(\max_g \left(\frac{T_g}{\sqrt{N} S_{xx,g}} \right)^2 \right) = o(1)$$

both under the condition that $N \rightarrow \infty$ and $\frac{\sqrt{N} S_{xx,1}}{T_1} \rightarrow \infty$. Used above:

$$\begin{aligned} P_{i\ell} &= T_{g(i)}^{-1} \mathbf{1}_{\{g(i)=g(\ell)\}} + \frac{(z_{g(i)t(i)} - \bar{z}_{g(i)})(z_{g(i)t(\ell)} - \bar{z}_{g(i)})}{S_{zz,g(i)}} \mathbf{1}_{\{g(i)=g(\ell)\}} \\ B_{ii} &= \frac{1}{n} \frac{z_{g(i)t(i)} - \bar{z}_{g(i)}}{S_{zz,g(i)}} \frac{T_{g(i)}}{S_{zz,g(i)}}. \end{aligned}$$

Finally,

$$\begin{aligned}\max_i w'_{iq} w_{iq} &= \max_t \frac{(z_{1t} - \bar{z}_1)^2}{S_{zz,1}} = o(1) \\ \mathbb{V}[\hat{\theta}_q]^{-1} \max_i (\tilde{x}'_{iq} \beta)^2 &= O\left(\max_{g \geq 2, t} \frac{z_{gt}^2 \delta_g^2}{N S_{zz,g}}\right) = o(1), \\ \mathbb{V}[\hat{\theta}_q]^{-1} \max_i (\tilde{x}'_{iq} \beta)^2 &= O\left(\max_{g \geq 2} \left(\frac{T_g}{\sqrt{N} S_{xx,g}}\right)^2\right) = o(1)\end{aligned}$$

under the conditions that $\frac{\sqrt{N}}{T_2} S_{zz,2} \rightarrow \infty$ and $S_{zz,1} \rightarrow \infty$. Thus, Theorem 3 applies when $\frac{\sqrt{N}}{T_1} S_{zz,1} = O(1)$.

Example 4. Let $\dot{f}_i = (\mathbf{1}_{\{j(g,t)=0\}}, f'_i)' = (\mathbf{1}_{\{j(g,t)=0\}}, \mathbf{1}_{\{j(g,t)=1\}}, \dots, \mathbf{1}_{\{j(g,t)=J\}})'$ and define the following partial design matrices with and without dropping ψ_0 from the model:

$$S_{ff} = \sum_{i=1}^n f_i f'_i, \quad S_{\dot{f}\dot{f}} = \sum_{i=1}^n \dot{f}_i \dot{f}'_i, \quad S_{\Delta f \Delta f} = \sum_{g=1}^N \Delta f_g \Delta f'_g, \quad S_{\Delta \dot{f} \Delta \dot{f}} = \sum_{g=1}^N \Delta \dot{f}_g \Delta \dot{f}'_g,$$

where $\Delta \dot{f}_g = \dot{f}_{i(g,2)} - \dot{f}_{i(g,1)}$. Letting \dot{D} be a diagonal matrix that holds the diagonal of $S_{\Delta \dot{f} \Delta \dot{f}}$ we have that

$$E = \dot{D} S_{\dot{f}\dot{f}}^{-1} \quad \text{and} \quad \mathcal{L} = \dot{D}^{-1/2} S_{\Delta \dot{f} \Delta \dot{f}} \dot{D}^{-1/2}.$$

$S_{\Delta \dot{f} \Delta \dot{f}}$ is rank deficient with $S_{\Delta \dot{f} \Delta \dot{f}} \mathbf{1}_{J+1} = 0$ from which it follows that the non-zero eigenvalues of $E^{1/2} \mathcal{L} E^{1/2}$ (which are the non-zero eigenvalues of $S_{\dot{f}\dot{f}}^{-1} S_{\Delta \dot{f} \Delta \dot{f}}$) are also the eigenvalues of $S_{\Delta f \Delta f} (S_{ff}^{-1} + \frac{\mathbf{1}_J \mathbf{1}'_J}{S_{ff,11}})$. Finally, from the Woodbury formula we have that A_{ff} is invertible with

$$A_{ff}^{-1} = n(S_{ff} - n \bar{f} \bar{f}')^{-1} = n \left(S_{ff}^{-1} + n \frac{S_{ff}^{-1} \bar{f} \bar{f}' S_{ff}^{-1}}{1 - n \bar{f}' S_{ff}^{-1} \bar{f}} \right) = n \left(S_{ff}^{-1} + \frac{\mathbf{1}_J \mathbf{1}'_J}{S_{ff,11}} \right),$$

so

$$\lambda_\ell = \lambda_\ell(A_{ff} S_{\Delta f \Delta f}^{-1}) = \frac{1}{\lambda_{J+1-\ell}(S_{\Delta f \Delta f} A_{ff}^{-1})} = \frac{1}{n \lambda_{J+1-\ell}(E^{1/2} \mathcal{L} E^{1/2})}.$$

With $E_{jj} = 1$ for all j , we have that

$$\frac{\lambda_1^2}{\sum_{\ell=1}^J \lambda_\ell^2} = \frac{\dot{\lambda}_J^{-2}}{\sum_{\ell=1}^J \dot{\lambda}_\ell^{-2}} \leq \frac{4}{(\sqrt{J} \dot{\lambda}_J)^2}$$

since $\dot{\lambda}_\ell \leq 2$ (Chung, 1997, Lemma 1.7). An algebraic definition of Cheeger's constant \mathcal{C} is

$$\mathcal{C} = \min_{X \subseteq \{0, \dots, J\} : \sum_{j \in X} \dot{D}_{jj} \leq \frac{1}{2} \sum_{j=0}^J \dot{D}_{jj}} \frac{-\sum_{j \in X} \sum_{k \notin X} S_{\Delta f \Delta f, jk}}{\sum_{j \in X} \dot{D}_{jj}}$$

and it follows from the Cheeger inequality $\dot{\lambda}_J \geq 1 - \sqrt{1 - \mathcal{C}^2}$ (Chung, 1997, Theorem 2.3) that $\sqrt{J} \dot{\lambda}_J \rightarrow \infty$ if $\sqrt{J} \mathcal{C} \rightarrow \infty$.

For the stochastic block model we consider J odd and order the firms so that the first $(J+1)/2$ firms belongs to the first block, and the remaining firms belong to the second block. We assume that $\Delta \dot{f}_g$ is generated *i.i.d.* across g according to

$$\Delta \dot{f} = W(1 - D) + BD$$

where (W, B, D) are mutually independent, $P(D = 1) = 1 - P(D = 0) = p_b \leq \frac{1}{2}$, W is uniformly distributed on $\{v \in \mathbb{R}^{J+1} : v' \mathbf{1}_{J+1} = 0, v'v = 2, \max_j v_j = 1, v'c = 0\}$, and B is uniformly distributed on $\{v \in \mathbb{R}^{J+1} : v' \mathbf{1}_{J+1} = 0, v'v = 2, \max_j v_j = 1, (v'c)^2 = 4\}$ for $c = (\mathbf{1}'_{(J+1)/2}, -\mathbf{1}'_{(J+1)/2})'$. In this model $E_{jj} = 1$ for all j . The following lemma characterizes the large sample behavior of $S_{\Delta f \Delta f}$ and \mathcal{L} . Based on this lemma it is relatively straightforward (but tedious) to verify the high-level conditions imposed in the paper.

Lemma C.11. *Suppose that $\frac{\log(J)}{np_b} + \frac{J \log(J)}{n} \rightarrow 0$ as $n \rightarrow \infty$ and $J \rightarrow \infty$. Then*

$$\left\| \underline{\mathcal{L}}^\dagger \frac{J+1}{n} S_{\Delta f \Delta f} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| = o_p(1) \quad \text{and} \quad \left\| \underline{\mathcal{L}}^\dagger \mathcal{L} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| = o_p(1)$$

where $\underline{\mathcal{L}} = I_{J+1} - \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} - (1 - 2p_b) \frac{cc'}{J+1}$ and $\|\cdot\|$ returns the largest singular value of its argument. Additionally, $\max_\ell \dot{\lambda}_\ell^{-1} |\dot{\lambda}_\ell - \dot{\lambda}_\ell| = o_p(1)$ where $\dot{\lambda}_1 \geq \dots \geq \dot{\lambda}_J$ are the non-zero eigenvalues of $\underline{\mathcal{L}}^\dagger$.

Proof. First note that

$$\frac{J+1}{n} \mathbb{E}[S_{\Delta f \Delta f}] - \underline{\mathcal{L}} = \frac{2+2p_b}{J-1} \left(I_{J+1} - \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} - \frac{cc'}{J+1} \right) + \frac{4p_b}{J-1} \frac{cc'}{J+1},$$

and $\underline{\mathcal{L}}^\dagger = I_{J+1} - \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} - \left(1 - \frac{1}{2p_b}\right) \frac{cc'}{J+1}$, so

$$\begin{aligned} \left\| \underline{\mathcal{L}}^\dagger \frac{J+1}{n} \mathbb{E}[S_{\Delta f \Delta f}] - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| &= \left\| \frac{2+2p_b}{J-1} \left(I_{J+1} - \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} - \frac{cc'}{J+1} \right) + \frac{2}{J-1} \frac{cc'}{J+1} \right\| \\ &= \frac{2+2p_b}{J-1} \end{aligned}$$

Therefore, we can instead show that $\|S\| = o_p(1)$ for the zero mean random matrix

$$S = (\underline{\mathcal{L}}^\dagger)^{1/2} \frac{J+1}{n} \left(S_{\Delta f \Delta f} - \mathbb{E}[S_{\Delta f \Delta f}] \right) (\underline{\mathcal{L}}^\dagger)^{1/2} = \sum_{g=1}^N s_g s_g' - \mathbb{E}[s_g s_g']$$

where $s_g = \sqrt{\frac{J+1}{n}} \Delta f_g - \frac{\sqrt{2p_b-1}}{\sqrt{2p_b n}} \Delta f_g' c \frac{c}{\sqrt{J+1}}$. Now since

$$s_g' s_g = O\left(\frac{J}{n} + \frac{1}{np_b}\right) \quad \text{and} \quad \left\| \sum_{g=1}^N \mathbb{E}[s_g s_g' s_g s_g'] \right\| = O\left(\frac{J}{n} + \frac{1}{np_b}\right)$$

it follows from (Oliveira, 2009, Corollary 7.1) that $\mathbb{P}(\|S\| \geq t) \leq 2(J+1)e^{-\frac{t^2(\frac{J}{n} + \frac{1}{np_b})}{c(8+4t)}}$ for some constant c not depending on n . Letting $t \propto \sqrt{\frac{\log(J/\delta_n)}{np_b} + \frac{J \log(J/\delta_n)}{n}}$ for δ_n that approaches zero slowly enough that $\frac{\log(J/\delta_n)}{np_b} + \frac{J \log(J/\delta_n)}{n} \rightarrow 0$ yields the conclusion that $\|S\| = o_p(1)$.

Since $\mathcal{L} = \dot{D}^{-1/2} S_{\Delta f \Delta f} \dot{D}^{-1/2}$ the second conclusion follows from the first if $\|\frac{J+1}{n} \dot{D} - I_{J+1}\| = o_p(1)$. We have $\frac{J+1}{n} \mathbb{E}[\dot{D}] = I_{J+1}$ and $\frac{J+1}{n} \dot{D}_{jj} = \frac{J+1}{n} \sum_{g=1}^N (\Delta f_g' e_j)^2$ where e_j is the j -th basis vector in \mathbb{R}^{J+1} and $\mathbb{P}((\Delta f_g' e_j)^2 = 1) = 1 - \mathbb{P}((\Delta f_g' e_j)^2 = 0) = \frac{2}{J+1}$. Thus it follows from $\mathbb{V}(\frac{J+1}{n} \dot{D}_{jj}) \leq 2 \frac{J+1}{n}$ and standard exponential inequalities that $\|\frac{J+1}{n} \dot{D} - I_{J+1}\| = \max_j |\frac{J+1}{n} \dot{D}_{jj} - 1| = o_p(1)$ since $\frac{J \log(J)}{n} \rightarrow 0$.

Finally, we note that $\left\| \underline{\mathcal{L}}^\dagger \mathcal{L} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}_{J+1}'}{J+1} \right\| \leq \epsilon$ implies

$$v' \underline{\mathcal{L}} v (1 - \epsilon) \leq v' \mathcal{L} v \leq v' \underline{\mathcal{L}} v (1 + \epsilon)$$

which together with the Courant-Fischer min-max principle yields $(1 - \epsilon) \leq \frac{\lambda_j}{\underline{\lambda}_j} \leq (1 + \epsilon)$. \square

Next, we will verify the high-level conditions of the paper in a model that uses $\frac{n}{J+1} \underline{\mathcal{L}}$ in place of $S_{\Delta f \Delta f}$ and $\frac{1}{n} \underline{\mathcal{L}}^\dagger$ in place of \tilde{A} and $\frac{n}{J+1} I_{J+1}$ in place of \dot{D} . Using an underscore to denote objects from this model we have

$$\begin{aligned} \max_g \underline{P}_{gg} &= \max_g \frac{J+1}{n} \Delta f_g' \underline{\mathcal{L}}^\dagger \Delta f_g = 2 \frac{J+1}{n} + 2 \frac{(1-2p_b)}{np_b} = o(1), \\ \text{trace}(\underline{\tilde{A}}^2) &= \frac{\text{trace}((\underline{\mathcal{L}}^\dagger)^2)}{n^2} = \frac{J-1}{n^2} + \frac{1}{4(np_b)^2} = o(1), \\ \frac{\lambda_1^2}{\sum_{\ell=1}^J \lambda_\ell^2} &= \frac{1}{\lambda_J^2 \text{trace}((\underline{\mathcal{L}}^\dagger)^2)} = \frac{1}{(J-1)4p_b^2 + 1} \end{aligned}$$

which is $o(1)$ if and only if $\sqrt{J}p_b \rightarrow \infty$, and $\frac{\lambda_2^2}{\sum_{\ell=1}^J \lambda_\ell^2} \leq \frac{1}{J}$. Furthermore,

$$\begin{aligned} \max_g \underline{w}_{g1}^2 &= \max_g \left(\frac{c'(\underline{\mathcal{L}}^\dagger)^{1/2} \Delta \dot{f}_g}{\sqrt{n}} \right)^2 = \left(\frac{2}{\sqrt{2p_b n}} \right)^2 = \frac{2}{np_n} = o(1), \\ \max_g (\tilde{x}'_g \beta)^2 &= \max_g \left(\frac{1}{n} \psi' \underline{\mathcal{L}}^\dagger \Delta \dot{f}_g \right)^2 \leq \frac{2}{n^2} \left[\max_g (\Delta \dot{f}'_g \psi)^2 + \left(1 - \frac{1}{2p_b} \right)^2 (\bar{\psi}_{cl,1} - \bar{\psi}_{cl,2})^2 \right] \\ &= O \left(\frac{1}{n^2} + \frac{1}{(np_b)^2} \right) \end{aligned}$$

which is $o(\mathbb{V}[\hat{\theta}])$ if $\sqrt{J}p_b \rightarrow \infty$ as $\text{trace}(\tilde{\underline{A}}^2) = O(\mathbb{V}[\hat{\theta}])$ and

$$\max_g (\tilde{x}'_{g1} \beta)^2 = \max_g \left(\frac{1}{n} \psi' \Delta \dot{f}_g \right)^2 = O \left(\frac{1}{n^2} \right) = o(\mathbb{V}[\hat{\theta}]).$$

Finally,

$$\max_g (\tilde{x}'_g \beta)^2 = O \left(\sum_{g=1}^N B_{gg}^2 \right) = O \left(\max_g B_{gg} \text{trace}(\tilde{\underline{A}}) \right)$$

where

$$\begin{aligned} \max_g B_{gg} &= \max_g \Delta \dot{f}'_g \frac{J+1}{n^2} (\underline{\mathcal{L}}^\dagger)^2 \Delta \dot{f}_g = 2 \frac{J+1}{n^2} + \frac{1-4p_b^2}{(np_b)^2} = O \left(\text{trace}(\tilde{\underline{A}}^2) \right) \\ \text{trace}(\tilde{\underline{A}}) &= \frac{J-1}{n} + \frac{1}{2p_b n} = o(1) \end{aligned}$$

so $\max_g B_{gg} \text{trace}(\tilde{\underline{A}}) = O(\text{trace}(\tilde{\underline{A}}^2))o(1)$.

Finally, we use the previous lemma to transfer the above results to their relevant sample analogues.

$$\begin{aligned} \max_g |P_{gg} - \underline{P}_{gg}| &= \max_g |\Delta \dot{f}'_g (S_{\Delta \dot{f} \Delta \dot{f}}^\dagger - \frac{J+1}{n} \underline{\mathcal{L}}^\dagger) \Delta \dot{f}_g| \\ &= \frac{J+1}{n} \max_g \left| \Delta \dot{f}'_g (\underline{\mathcal{L}}^\dagger)^{1/2} \left(\underline{\mathcal{L}}^{1/2} \frac{n}{J+1} S_{\Delta \dot{f} \Delta \dot{f}}^\dagger \underline{\mathcal{L}}^{1/2} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right) (\underline{\mathcal{L}}^\dagger)^{1/2} \Delta \dot{f}_g \right| \\ &= O \left(\left\| \underline{\mathcal{L}}^\dagger \frac{J+1}{n} S_{\Delta \dot{f} \Delta \dot{f}} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| \right) \max_g P_{gg} = o \left(\max_g P_{gg} \right) \\ \left| \text{trace}(\tilde{\underline{A}}^2 - \underline{\tilde{A}}^2) \right| &= \left| \sum_{\ell=1}^J \frac{1}{n^2 \lambda_\ell^2} - \frac{1}{n^2 \dot{\lambda}_\ell^2} \right| = \text{trace}(\tilde{\underline{A}}^2) O \left(\max_\ell \left| \frac{\dot{\lambda}_\ell - \lambda_\ell}{\dot{\lambda}_\ell} \right| \right) = o_p \left(\text{trace}(\tilde{\underline{A}}^2) \right) \\ \left| \frac{\lambda_1^2}{\sum_{\ell=1}^J \lambda_\ell^2} - \frac{\dot{\lambda}_1^2}{\sum_{\ell=1}^J \dot{\lambda}_\ell^2} \right| &= \frac{\lambda_1^2}{\sum_{\ell=1}^J \lambda_\ell^2} O \left(\frac{|\dot{\lambda}_J - \lambda_J|}{\dot{\lambda}_J} + \frac{|\text{trace}(\tilde{\underline{A}}^2 - \underline{\tilde{A}}^2)|}{\text{trace}(\tilde{\underline{A}}^2)} \right) = o_p(1) \end{aligned}$$

with a similar argument applying to $\frac{\lambda_2^2}{\sum_{\ell=1}^J \lambda_\ell^2} - \frac{\lambda_2^2}{\sum_{\ell=1}^J \underline{\lambda}_\ell^2}$. Furthermore,

$$\max_g w_{g1}^2 = \max_g \left(\Delta \dot{f}_g \left(\frac{J+1}{n} \underline{\mathcal{L}}^\dagger \right)^{1/2} \left(\underline{\mathcal{L}} \frac{n}{J+1} S_{\Delta f \Delta f}^\dagger \right)^{1/2} q_1 \right)^2 \leq \| (\underline{\mathcal{L}} \frac{n}{J+1} S_{\Delta f \Delta f}^\dagger)^{1/2} \| \max_g \underline{P}_{gg} = o_p(1)$$

and $\max_g |(\tilde{x}'_g \beta)^2 - (\tilde{x}'_g \beta)^2| = o_p(\text{trace}(\tilde{A}^2))$ since

$$\begin{aligned} \max_g (\tilde{x}'_g \beta - \tilde{x}'_g \beta)^2 &= \frac{J+1}{n^2} \max_g \left(\Delta \dot{f}'_g \underline{\mathcal{L}}^\dagger \left(\underline{\mathcal{L}} S_{\Delta f \Delta f} \dot{D} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right) \frac{\psi}{\sqrt{J+1}} \right)^2 \\ &\leq \left\| \underline{\mathcal{L}} S_{\Delta f \Delta f} \dot{D} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| \max_g \underline{B}_{gg} \frac{\|\psi\|^2}{J+1} \\ &= o_p(\text{trace}(\tilde{A}^2)) \end{aligned}$$

and this also handles $\max_i |(\tilde{x}'_{g1} \beta)^2 - (\tilde{x}'_{g1} \beta)^2| = o_p(1)$ as the previous result does not depend on the behavior of $\sqrt{J} p_b$. Finally,

$$\begin{aligned} \max_g |B_{gg} - \underline{B}_{gg}| &= \frac{J+1}{n^2} \max_g \left| \Delta \dot{f}'_g \underline{\mathcal{L}}^\dagger \left(\frac{n}{J+1} \underline{\mathcal{L}} S_{\Delta f \Delta f}^\dagger \dot{D} S_{\Delta f \Delta f}^\dagger \underline{\mathcal{L}} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right) \underline{\mathcal{L}}^\dagger \Delta \dot{f}_g \right| \\ &\leq \left\| \frac{n}{J+1} \underline{\mathcal{L}} S_{\Delta f \Delta f}^\dagger \frac{J+1}{n} \dot{D} \frac{n}{J+1} S_{\Delta f \Delta f}^\dagger \underline{\mathcal{L}} - I_{J+1} + \frac{\mathbf{1}_{J+1} \mathbf{1}'_{J+1}}{J+1} \right\| \max_g \underline{B}_{gg} \\ &= o_p(\max_g \underline{B}_{gg}) \\ \left| \text{trace}(\tilde{A} - \tilde{A}) \right| &= \left| \sum_{\ell=1}^J \frac{1}{n \dot{\lambda}_\ell} - \frac{1}{n \underline{\dot{\lambda}}_\ell} \right| = \text{trace}(\tilde{A}) O \left(\max_\ell \left| \frac{\dot{\lambda}_\ell - \underline{\dot{\lambda}}_\ell}{\dot{\lambda}_\ell} \right| \right) = o_p \left(\text{trace}(\tilde{A}) \right) \end{aligned}$$