

Correcting Small Sample Bias in Linear Models with Many Covariates

Miren Azkarate-Askasua

Miguel Zerecero*

November 10, 2022

Abstract

Estimations of quadratic forms in the parameters of linear models exhibit small-sample bias. The direct computation for a bias correction is not feasible when the number of covariates is large. We propose a bootstrap method for correcting this bias that accommodates different assumptions on the structure of the error term including general heteroscedasticity and serial correlation. Our approach is suited to correct variance decompositions and the bias of multiple quadratic forms of the same linear model without increasing the computational cost. We show with Monte Carlo simulations that our bootstrap procedure is effective in correcting the bias and find that it is faster than other methods in the literature. Using administrative data for France, we apply our method by carrying out a variance decomposition of a linear model of log wages with person and firm fixed effects. We find that the person and firm effects are less important in explaining the variance of log wages after correcting for the bias and depending on the specification their correlation becomes positive after the correction.

JEL Codes: C13, C23, C55, J30, J31

Keywords: Variance components, many regressors, fixed effects, bias correction, limited mobility bias.

*Azkarate-Askasua: University of Mannheim (azkarate-askasua@uni-mannheim.de); Zerecero: University of California, Irvine (mzerecer@uci.edu). First version: May 2018. We thank our advisor Christian Hellwig for his guidance throughout this project. We thank Fabrice Collard, Thomas Crossley, Patrick Fève, Simen Gaure, Silvia Goncalves, Cristina Gualdani, Koen Jochmans, Elia Lapenta, Tim Lee, Thierry Magnac, Nour Meddahi, Jean-Marc Robin and Uta Schönberg for helpful comments. We gratefully acknowledge financial support from the Toulouse School of Economics to get access to the administrative data provided by CASD. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the 'Investissements d'avenir' program (reference: ANR-10-EQPX-17 – Centre d'accès sécurisé aux données – CASD). Zerecero acknowledges the financial support from CONACYT (reference: 329103/383874). Azkarate-Askasua acknowledges the support from the German Research Foundation (through the CRC-TR-224 project B06). All errors are our own.

1 Introduction

With the increased availability of large panel data sets, researchers have been interested in understanding to what extent unobserved heterogeneity can explain the variation of an outcome of interest. Usually, econometricians include fixed effects in a standard linear model to control for this unobserved heterogeneity and then perform a variance decomposition. These methods have been used in the context of education to study the importance of classroom effects (e.g. [Chetty, Friedman, Hilger, Saez, Schanzenbach, and Yagan \(2011\)](#)) and extensively in the labor market context where log-additive models of wages are used to study the determinants of labor income and sources of wage inequality (e.g. [Abowd, Kramarz, and Margolis \(1999\)](#); [Card, Heining, and Kline \(2013\)](#); [Iranzo, Schivardi, and Tosetti \(2008\)](#); [Lopes de Melo \(2018\)](#)).

The elements of a variance decomposition of a linear model are quadratic objects in the parameters. As long as the parameters are estimated with noise, these quadratic objects are subject to small-sample bias. This bias can be substantial in empirical applications and can even change the sign of estimated covariances. Moreover, in most applications this bias does not fade away by increasing the sample size. This is the case when using panel data, as the number of parameters to estimate, i.e. the number of fixed effects, grows with the sample size.

Focusing on the context of labor economics, researchers have used employer-employee matched datasets to study the sorting patterns of workers into firms. Various papers have estimated a linear model of log wages with person and firm fixed effects, following the seminal work of [Abowd, Kramarz, and Margolis \(1999\)](#) (AKM henceforth). These studies compute the correlation between the person and firm fixed effects to determine the degree of sorting in the labor market. Most studies have found zero or negative correlations, casting doubt on whether there is sorting in the labor market. However, as noted by [Abowd, Kramarz, Lengeremann, and Pérez-Duarte \(2004\)](#) this correlation is likely to suffer from small-sample bias, dubbed *limited mobility* bias in their paper. [Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler \(2023\)](#) show that the limited mobility bias is substantial when performing a variance decomposition of log wages for several countries.

[Andrews, Gill, Schank, and Upward \(2008\)](#) derive formulas for correcting the bias when the errors are homoscedastic. [Gaure \(2014\)](#) provides formulas for more general variance structures. Unfortunately, the direct implementation of these corrections in high dimensional models is infeasible. The reason is that the corrections entail computing the inverse of an impractically large matrix, which has prevented the direct application of the correction formulas.¹²

In this paper we propose a bootstrap method to correct for small-sample bias in quadratic forms in the estimated parameters of linear models with a large number of covariates. The main application of the method is the correction of variance decompositions of multi-way fixed effects models. The contribution of the paper is to provide an easy-to-implement and fast alternative to the bias correction methods for quadratic forms present in the literature. We see our paper as an application of the bootstrap in variance decompositions of multi-way fixed effects.

Our method consists on re-estimating the same quadratic forms of a linear model on boot-

¹By large matrix we mean a matrix with dimension in the order of hundreds of thousands or millions.

²Some recent examples of papers doing a variance decomposition of log wages into worker and firm fixed effects without correcting the limited mobility bias are: [Song, Price, Guvenen, Bloom, and Von Wachter \(2019\)](#), [Sorkin \(2018\)](#), [Card, Cardoso, Heining, and Kline \(2018\)](#).

strapped data. The sample means of these quadratic forms are our bias correction terms. Using Monte Carlo simulations we show that our method successfully corrects the bias of quadratic forms for multiple assumptions on the variance structure of the error term, such as heteroscedasticity, serial correlation or clustering. In practice, under the assumption of a diagonal covariance matrix, we use a wild bootstrap. When the covariance matrix is assumed non-diagonal, we use a wild block bootstrap (Cameron, Gelbach, and Miller, 2008) that is valid for unrestricted dependence of the error terms within group and heteroscedasticity. The wild block bootstrap is flexible in the definition of the group and therefore allows, for example, the clustering of the errors depending on geographical area or serial correlation within the worker-firm match.³

Our approach is similar to the ones proposed by Gaure (2014) and by Kline, Saggio, and Sølvssten (2020). The bias appears as the trace of a matrix, but when the number of covariates of the linear model is large, the explicit computation of this trace is not feasible. Like ours, both of their methods rely on iterative procedures to compute an estimate of the trace term. Gaure exploits the fact that the trace can be represented as the expectation of a more manageable quadratic form in a random vector. This expectation can in turn be approximated by estimating a sample mean after simulating different random vectors.⁴

Kline, Saggio, and Sølvssten (2020) (KSS henceforth) follow a similar approach to Gaure (2014). In their large-scale computation procedure, they estimate the trace term leading to the bias and implement a bias correction assuming either heteroscedasticity or serial correlation of the errors. An important point of their paper is that their leave-one-out covariance-matrix estimate is unbiased. Our approach differs in the way we estimate the trace term which allows us to be faster and more flexible. One drawback is that we can not use their leave-one-out covariance-matrix estimate as its not suitable for bootstraps.⁵ However, our bootstrap procedure can accommodate easily more complicated variance structures. In terms of speed, Monte Carlo simulations show that our correction takes between 56% to 70% of the computing time of KSS, and has similar accuracy.

The computational cost in Gaure and KSS comes from estimating a bias correction for each interested quadratic form, as it requires solving a large system of linear equations in each iteration that are particular to each quadratic form. In contrast, we re-estimate the model with bootstrapped data and show that a sample mean of the *bootstrapped* moment estimates is an unbiased and consistent estimator of the direct bias correction term. In our method, the computational cost comes from estimating the linear model in each bootstrap but does not increase depending on the number of moments to correct. We need to solve one system of linear equations per bootstrap regardless of the number of moments to correct, while with the Gaure and KSS methods, one needs to solve as many systems of equations per iteration as needed corrections.⁶ They implement corrections of the

³Other examples include errors correlated within firms, workers or occupations.

⁴In particular, the way Gaure estimates the trace is known as the Hutchinson method. Denote a random vector $x \in \mathbb{R}^n$, where each individual entry is independently distributed Rademacher (entries can take values of 1 or -1 with probability 1/2). Then, for a square matrix $A \in \mathbb{R}^{n \times n}$ we have that $tr(A) = E(x'Ax)$. The Hutchinson estimator of the trace of matrix A is $\frac{1}{M} \sum_{i=1}^M x_i'Ax_i$, where x_i is the i -th draw of the random vector x ; see Hutchinson (1989) and Avron and Toledo (2011). Gaure's R package *lfe* implements the correction when the error terms are assumed homoscedastic. The function applying the correction is *bccorr*; see Gaure (2013). Gaure (2014) sketches the procedure to correct for the bias when the error terms are heteroscedastic, but to the best of our knowledge he does not implement it in his R package.

⁵In the heteroscedastic case, when using their leave-one-out procedure, some diagonal elements of the estimated covariance-matrix could be negative. In practice, we can not implement a bootstrap procedure with such covariance matrix estimate as we need to take the squared root of each diagonal element.

⁶For example, consider the linear model $y_t = X_{1t}\beta_1 + X_{2t}\beta_2 + \varepsilon_t$ where one is interested in doing a variance decomposition for each period t . This would yield three quadratic objects to correct ($Var(X_1\hat{\beta}_1)$, $Var(X_2\hat{\beta}_2)$, $Cov(X_1\hat{\beta}_1, X_2\hat{\beta}_2)$) per period.

second order moments of the two leading fixed effects while we can directly perform a full variance decomposition, which is therefore suited for corrections on multi-way fixed effect regressions.

Our method is easy to implement as it only requires estimating linear models. While of course this requires solving a system of linear equations, like in the case of Gaure and KSS, these systems are more ubiquitous. Therefore, there is a wide range of algorithms that estimate linear models for different softwares. We provide codes in Matlab, but the user can easily implement the correction method by taking profit of other algorithms in alternative softwares.

We apply our method to French administrative data and perform a variance decomposition of an estimated AKM type model. Consistent with the Andrews et al. (2008) formulation, we find that sample variances of person and firm effects are reduced and their covariance increased after the correction. The estimated correlation at the connected set passes from -0.10 to almost zero under the assumption of serial correlation of the error terms within the match.⁷

Labor economists have been aware of the small-sample bias problem with quadratic forms in the parameters and the difficulty in estimating a correction at least since Andrews et al. (2008). There have been several attempts to correct this bias when performing variance decompositions of estimated linear models. Some methods are based on variations of the jackknife, such as the split-panel jackknife estimator by Dhaene and Jochmans (2015) or the leave-one-out estimator by KSS mentioned above. Bonhomme, Lamadon, and Manresa (2019) relax the exogenous mobility assumption from the AKM model and mitigate the small-sample bias by reducing the dimensionality of the estimated parameters. Borovičková and Shimer (2017) propose an alternative method to estimate the correlation between worker and firm types.

2 The Bias

For clarity of exposition we layout the source of the bias. Consider the following linear model

$$Y = X\beta + u, \quad (1)$$

where Y is a $n \times 1$ vector representing the endogenous variable, X is a matrix of covariates of size $n \times k$, and β is a vector of parameters.⁸ The error term u satisfies mean independence $\mathbb{E}(u|X) = 0$.

The OLS estimate of β is,

$$\hat{\beta} = \beta + Qu,$$

where $Q = (X'X)^{-1} X'$.

We are interested in estimating the following quadratic form $\varphi = \beta' A \beta$ for some matrix A of dimensions $k \times k$, where $\mathbb{E}(A|X) = A$. From the expression for $\hat{\beta}$ we can decompose the plug-in estimator $\hat{\varphi}_{PI} = \hat{\beta}' A \hat{\beta}$ as,

$$\hat{\varphi}_{PI} = \beta' A \beta + u' Q' A Q u + 2u' Q' A \beta. \quad (2)$$

Using the general formula for the expectation of quadratic forms, the exclusion restriction $\mathbb{E}(u|X) =$

⁷Abowd et al. (2004), also using French data but a different sample, found a correlation of -0.28.

⁸We follow loosely the notation in Kline et al. (2020) for the interested reader to compare the papers.

0, and $\mathbb{E}(A|X) = A$ we obtain,⁹

$$\mathbb{E}(\widehat{\varphi}_{PI}|X) = \beta' A \beta + \text{trace}(Q' A Q \mathbb{V}(u|X)) = \varphi + \delta, \quad (3)$$

where the bias $\delta \equiv \text{trace}(Q' A Q \mathbb{V}(u|X))$ comes from the fact that $\widehat{\beta}$ is estimated with noise.

To get a bias correction one needs an estimate of the trace term δ . One option is to just plug-in the estimate for the conditional covariance matrix $\widehat{\mathbb{V}}(u|X)$. We define $\widehat{\delta}$ as the direct bias correction term, which is equal to

$$\widehat{\delta} \equiv \text{trace}(Q' A Q \widehat{\mathbb{V}}(u|X)). \quad (4)$$

Computing $\widehat{\delta}$ is difficult when the number of covariates is large because it requires to calculate first the matrix Q , which is itself a function of the inverse of a very large matrix.¹⁰ In the next section we propose a methodology to apply a computationally feasible correction.

We define the following bias-corrected estimate of the quadratic form φ as

$$\widehat{\varphi} = \widehat{\beta}' A \widehat{\beta} - \widehat{\delta}.$$

As long as $\mathbb{E}(\widehat{\delta}|X) = \delta$, then it follows that $\mathbb{E}(\widehat{\varphi}|X) = \varphi$.

Proposition 1. *The direct bias correction $\widehat{\delta}$ is an unbiased estimate of the bias term δ if and only if $\widehat{\mathbb{V}}(u|X)$ is an unbiased estimator of $\mathbb{V}(u|X)$.*

Thus, it is necessary to have an unbiased estimate of the covariance matrix $\mathbb{V}(u|X)$ to have an unbiased estimate of the quadratic form φ .

2.1 Components of a variance decomposition as quadratic objects

When performing a variance decomposition of a linear model, one can think of each element as a particular form of $\widehat{\beta}' A \widehat{\beta}$ with the appropriate choice of A . To see this, we can rewrite (1) as

$$Y = X_1 \beta_1 + X_2 \beta_2 + u, \quad (5)$$

where X_1 and X_2 are matrices of covariates of size $n \times k_1$ and $n \times k_2$, $k = k_1 + k_2$ with $X = [X_1 \ X_2]$ and $\beta' = [\beta_1' \ \beta_2']$.

We are interested in the sample variances ($\widehat{\text{var}}(X_1 \beta_1)$, $\widehat{\text{var}}(X_2 \beta_2)$) and covariance ($\widehat{\text{cov}}(X_1 \beta_1, X_2 \beta_2)$), denoted, respectively, as σ_1^2 , σ_2^2 and σ_{12} .¹¹ Define $\mathbf{1}$ as a vector of ones with appropriate length. Then, denote the demeaning operator as $M_1 = \mathbf{I} - \frac{1}{n} \mathbf{1} \mathbf{1}'$. We can then write the sample variances and covariances in matrix notation as

$$\sigma_j^2 = \beta' A_j \beta, \quad \text{for } j = \{1, 2\} \text{ and}$$

$$\sigma_{12} = \beta' A_{12} \beta,$$

⁹Given a random vector x and a symmetric matrix B we have that $\mathbb{E}(x' B x) = \mathbb{E}(x') B \mathbb{E}(x) + \text{trace}(B \mathbb{V}(x))$.

¹⁰The dimension of this matrix is related to the number of covariates that are estimated in the linear model. In a typical AKM type model the data will comprise of hundreds of thousands or millions of workers and tens of thousands of firms, each representing a covariate in the model.

¹¹The sample variance for a vector $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$ is $\widehat{\text{var}}(\mathbf{x}) = \frac{1}{n-1} \sum_{i=1}^N (x_i - \bar{x})^2$, where \bar{x} is the sample mean. Similarly, the sample covariance for vectors \mathbf{x} and \mathbf{y} is $\widehat{\text{cov}}(\mathbf{x}, \mathbf{y}) = \frac{1}{n-1} \sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})$.

where the symmetric matrices A_1 , A_2 and A_{12} are equal to

$$A_1 = \frac{1}{n-1} \begin{pmatrix} X_1' M_1 X_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad A_2 = \frac{1}{n-1} \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & X_2' M_1 X_2 \end{pmatrix}, \quad A_{12} = \frac{1}{2(n-1)} \begin{pmatrix} \mathbf{0} & X_1' M_1 X_2 \\ X_2' M_1 X_1 & \mathbf{0} \end{pmatrix}.$$

The plug-in estimators of σ_1^2 , σ_2^2 and σ_{12} , obtained by substituting β with the OLS estimate $\hat{\beta}$, are just particular examples of $\hat{\varphi}_{PI}$. Therefore, these estimates are biased.

3 Bootstrap Correction

The bootstrap correction estimates the direct bias correction (4) by replicating the bias structure of the plug-in estimates (2). In this section we present the bootstrap correction and discuss different implementations depending on the choice of the covariance matrix estimate.

Suppose that we have the residuals of our original regression $\hat{u} = Y - X\hat{\beta}$. Using these residuals we can construct an estimate of the covariance matrix, $\hat{V}(u|X)$. We generate a new dependent variable for the bootstrap Y^* as:

$$Y^* = v^*,$$

where v^* is a vector containing the bootstrapped residuals. This is equivalent to performing a linear regression on bootstrapped data, while setting $\hat{\beta} = \mathbf{0}$. The construction of v^* will depend ultimately in the assumption that we are making about the error term. In particular, we need that the variance of the bootstrapped errors $\mathbb{V}(v^*|X, u)$ to be equal to $\hat{V}(u|X)$. The following proposition states the main result of the paper and all the proofs are left to the Appendix.

Proposition 2. *Suppose the regression model (1) is correctly specified. Let p denote the number of bootstraps. Define β_j^* as the OLS estimate of regressing v_j^* over X for the j -th bootstrap iteration. If the conditional variance-covariance matrix of the bootstrapped residuals $\mathbb{V}(v_j^*|X, u)$ is equal to $\hat{V}(u|X)$, and $\mathbb{E}(v_j^*|X, u) = 0$, then*

$$\delta^* \equiv \frac{1}{p} \sum_{j=1}^p \left(\beta_j^{*'} A \beta_j^* \right)$$

is an unbiased and consistent estimator of the direct bias correction $\hat{\delta}$.

The proposition tells us that instead of computing directly the direct bias correction term $\hat{\delta}$, we can estimate it using a sample average of estimated quadratic forms.

The intuition behind our bias estimator is that in every bootstrap iteration we are replicating the source of the bias, which is the noise embedded in the estimated parameters. The computational burden of our method comes from estimating β_j^* for each bootstrap.¹² The advantage of our method is twofold. First, we can correct several moments simultaneously, without increasing the computational time. For example, assume we are interested in doing a variance decomposition exercise for each year in the sample of an estimated linear model. Then, we would need to do a correction for the variances of each group of covariates and the covariance term for *every* year but do the bootstrap only one time. Second, to estimate β_j^* in every iteration one just needs to solve for

¹²Current softwares avoid the inversion of the $X'X$ matrix to estimate linear models and are therefore able to estimate linear models even when the number of covariates is large.

a least squares regression. There are extremely efficient procedures to compute these regressions, especially in cases where the high dimensionality of the covariates is the result of a large number of fixed effects.

The key for the bootstrap correction to work is that $\mathbb{V}(v^*|X, u)$ is equal to the sample variance-covariance matrix $\widehat{\mathbb{V}}(u|X)$, so the bootstrap correction δ^* is an unbiased and consistent estimator of the direct bias correction term $\widehat{\delta}$. Therefore, the bootstrap procedure has to be compatible with the underlying assumption on the structure of the error term.

The small sample properties of the bootstrap estimate δ^* would depend ultimately on the choice of estimate for the covariance matrix $\mathbb{V}(u|X)$. In particular, we have the following corollary for the bias which is just a consequence of Propositions 1 and 2.

Corollary 1. *Conditioning on X , if $\widehat{\mathbb{V}}(u|X)$ is an unbiased estimator of $\mathbb{V}(u|X)$, then the bootstrap correction δ^* is an unbiased estimator of the bias δ .*

In what follows we provide examples for some popular choices for estimators of the covariance matrix and how to implement the bootstrap correction.

Example 1: Homoscedasticity. When the errors are homoscedastic, we can use the well-known unbiased estimate of the covariance matrix $\widehat{\sigma}^2 \mathbf{I}$, where $\widehat{\sigma}^2 = n/(n-k) \sum_{i=1}^n \widehat{u}_i^2$.¹³ A suitable bootstrap could be a residual bootstrap with a degrees of freedom correction. This would mean re-sampling with replacement the estimated residuals and multiplying them by $\sqrt{n/(n-k)}$. Thus the variance of the bootstrapped errors would be equal to the estimated covariance-matrix $\widehat{\sigma}^2 \mathbf{I}$. Another possibility could be to simulate errors from a normal distribution with zero mean and variance $\widehat{\sigma}^2$. In the case of homoscedastic errors, the proposed bootstraps can replicate the variance of an unbiased estimate of the covariance matrix. Thus, the bootstrap bias correction δ^* is an unbiased estimate of the bias term δ ; see Corollary 1.

Example 2. Heteroscedasticity. Another popular assumption is when the covariance matrix is diagonal, with non-zero i th diagonal element equal to ψ_i . Let $\widehat{\psi}_i$ be the estimate of the variance for the i th observation error term. [MacKinnon and White \(1985\)](#) explore different consistent variance estimates $\widehat{\psi}_i$. These include

$$HC_0 = \widehat{u}_i^2, \quad HC_1 = \frac{n}{n-k} \widehat{u}_i^2 \quad \text{and} \quad HC_2 = \frac{\widehat{u}_i^2}{1 - h_{ii}},$$

where h_{ii} is the i th diagonal element of the projection matrix $H = X(X'X)^{-1}X'$. The term h_{ii} is sometimes known as the *leverage* of observation i , because, as explained by [Angrist and Pischke \(2008\)](#), it tell us how much *pull* a particular observation exerts over the regression line.

A suitable bootstrap for the different covariance matrix estimators is the Wild bootstrap. In our exercises below, we implement this bootstrap by first generating i.i.d. Rademacher random variables, meaning they take values of 1 or -1 with probability $1/2$. Then we multiply $\sqrt{\widehat{\psi}_i}$ to the i th Rademacher entry r_i . This would constitute the i th bootstrapped residual. For example, if we use HC_0 , the bootstrap residual would be $v_i^* = \widehat{u}_i r_i$. The Online Appendix contains the specific algorithms with the steps for this procedure.

¹³The origin of the bias is again a trace term that under homoscedasticity is equal to $n-k$. For a textbook explanation see Proposition 1.2. in [Hayashi \(2000\)](#).

When using the HC_2 estimates, we need first to calculate the leverage h_{ii} . When the number of covariates is large, a direct computation of the leverage is unfeasible. In the Online Appendix we show how to estimate this leverages by means of averaging the squared fitted values of linear regressions. We also provide a diagnostic and correction method to ensure that the estimated leverages are bounded above by 1.

In general, the three alternatives of covariance matrix estimates (HC_0, HC_1 and HC_2) are biased.¹⁴ For example, for HC_0 we have

$$\mathbb{E}(\hat{u}_i^2|X) = \psi_i - 2\psi_i h_{ii} + h_i' \mathbb{V}(u|X) h_i,$$

where h_i is the i th column of the projection matrix H .¹⁵ Thus, while δ^* is an unbiased estimate of $\hat{\delta}$ (Proposition 2), it would be biased with respect to δ (Proposition 1).

Recently, [Kline et al. \(2020\)](#) and [Jochmans \(2018\)](#) have proposed the following unbiased estimator of the i th conditional variance:¹⁶

$$HC_U = \frac{Y_i \hat{u}_i}{1 - h_{ii}}.$$

In practice, when estimating $\hat{\psi}_i$ with HC_U , some of the estimates are negative. This would prevent us from taking the square root of $\hat{\psi}_i$, which is needed in the bootstrap algorithm.¹⁷ However, even though HC_U is unbiased and HC_2 is not, it is not clear that minimizes the mean squared error compared to other variance estimates. For example, let $\hat{Y}_i = h_i' Y$ be the fitted value for observation i . Then,

$$HC_U = \frac{Y_i \hat{u}_i}{1 - h_{ii}} = \frac{(\hat{Y}_i + \hat{u}_i) \hat{u}_i}{1 - h_{ii}} = \frac{\hat{Y}_i \hat{u}_i}{1 - h_{ii}} + HC_2.$$

While the expectation of HC_U is equal to ψ_i , it can be the case that its variance is larger than the one of HC_2 . Thus, it is not clear that using the KSS correction would yield a more efficient bias corrected estimate of the quadratic forms compared to our bootstrap method. In fact, we show in the simulation exercises below that our method is in general more efficient in terms of mean squared errors.

Example 3: Clustered errors and serial correlation. When the error terms are clustered or present serial correlation within group, the covariance matrix is no longer diagonal. We restrict our attention to dependence of the errors only within a given group. Thus, we restrict to the case where the variance covariance matrix is block diagonal, as there are non zero elements around the diagonal corresponding to the dependence of the errors within the group g , but not across groups.¹⁸ One particular example is when the group is a worker-firm match and errors are autocorrelated within

¹⁴A particular case where the estimate is unbiased is when using HC_1 and the error terms are homoscedastic.

¹⁵A textbook exposition of these issues can be found in Chapter 8 of [Angrist and Pischke \(2008\)](#).

¹⁶See page 1862 of [Kline et al. \(2020\)](#).

¹⁷Negative estimates of individual variances are also prevalent in KSS. However, the way they approximate the bias does not require to take the square root of $\hat{\psi}_i$.

¹⁸Assume that the errors have a first order autocorrelation within group g and the true innovations are i.i.d. and therefore homoscedastic. We consider that the error term u of worker i at group g at time t in (1) is:

$$u_{i,g,t} = \rho u_{i,g,t-1} + \varepsilon_{i,g,t}, \quad \varepsilon_{i,g,t} \text{ i.i.d.}$$

We denote the variance of the innovation ε as σ_ε^2 . Ordering the data by group, suppose the first group has three observations and the

match. Following [Roodman, Nielsen, MacKinnon, and Webb \(2019\)](#) we estimate the variance of observation i , $\widehat{\psi}_i$, with a variant of HC_1 from Example 2 that takes into account the number of groups G : $\widehat{\psi}_i = \frac{G}{G-1} \frac{n}{n-k} \widehat{u}_i^2$.

When the errors present dependence within the group we use a wild block bootstrap as proposed by [Cameron et al. \(2008\)](#). This consists of a wild bootstrap that takes into account the group or cluster dependence of the data. It has the benefit of accommodating any structure of the dependence within group and also heteroscedasticity of the true innovations. The Online Appendix describes the algorithm to implement our bias correction that keeps the dependence structure among groups through a wild block bootstrap.

While the method proposed by KSS can also be adapted to "*settings where the data are organized into mutually and independent 'clusters'*" (page 1863 of [Kline et al. \(2020\)](#)), our method can also accomodate more general settings as the example below explains.

Example 4: Non-block-diagonal covariance matrices. Sometimes, the assumption on the error dependence do not yield a diagonal or block diagonal covariance matrix. This can happen when there are two (or more) dimensions of dependency. For example, when there are temporal and spatial dependency, as in [Driscoll and Kraay \(1998\)](#). In the AKM context, for example, there would be a non-block-diagonal covariance matrix if there is temporal dependence at the person *and* firm dimensions. With workers changing firms, the resulting dependence across observations would break any block-diagonal structure in the covariance matrix.

In short, our method can be applied whenever one can estimate a covariance matrix with positive diagonal entries. Then, the dependent variable used in the bootstrap can be generated by simulating an error vector with zero mean and a covariance matrix equal to the estimated one.

Efficiency gains. Using the bootstrap to correct for biases is ubiquitous in the literature. [MacKinnon and Smith Jr \(1998\)](#) (MS, henceforth) propose a similar bootstrap to correct for flat biases like the one considered here.¹⁹ MS propose building the bootstrapped dependent variable by using the original estimate of β , $Y^* = X\widehat{\beta} + v^*$. In the context of our application, that would mean to compute the quadratic objects $\beta_{p,MS}^{*'} A \beta_{p,MS}^*$ for each bootstrap p and use them to create a bias correction equal to

$$\delta_{MS}^* = \frac{1}{p} \sum_{j=1}^p \beta_{j,MS}^{*'} A \beta_{j,MS}^* - \widehat{\beta}' A \widehat{\beta}.$$

second group two observations. Then, $\mathbb{V}(u|X)$ is:

$$\mathbb{V}(u|X) = \frac{\sigma_\varepsilon^2}{1-\rho^2} \begin{pmatrix} 1 & \rho & \rho^2 & 0 & \dots & 0 \\ \rho & 1 & \rho & \vdots & \ddots & \vdots \\ \rho^2 & \rho & 1 & 0 & \dots & 0 \\ 0 & \dots & 0 & 1 & \rho & 0 & \dots & 0 \\ & & & \rho & 1 & 0 & \ddots & \\ \vdots & \ddots & \vdots & & & \ddots & & 0 \\ 0 & & 0 & & & & & 1 \end{pmatrix}.$$

The covariance matrix under clustering of the errors is similar but with all non-zero elements out of the diagonal equal to ρ .

¹⁹A flat bias is one that does not depend on the levels of the original estimates. The bias from the quadratic forms is flat because the trace term in (3) is independent of $\widehat{\beta}$.

MS already note that one can estimate a flat bias correction by using any $\hat{\beta}$ to generate Y^* . For example, one can use $\hat{\beta} = \mathbf{0}$, as we do in Proposition 2.

Analogously to equation (2) we have that in bootstrap j , $\beta_{j,MS}^* A \beta_{j,MS}^* = \hat{\beta}' A \hat{\beta} + (v_j^*)' Q' A Q v_j^* + 2v_j^{*'} Q' A \hat{\beta}$. When the errors are independent and the third moment is zero, it can be shown that the covariance of the last two terms conditional on X and u is equal to zero.²⁰ Thus we have that the variance of their bias correction conditional on X and u is

$$\mathbb{V}(\delta_{MS}^* | X, u) = \frac{1}{p} \mathbb{V}((v^*)' Q' A Q v^* | X, u) + \frac{4}{p} \mathbb{V}(v^{*'} Q' A \hat{\beta} | X, u).$$

The expression above can be rewritten as

$$\mathbb{V}(\delta_{MS}^* | X, u) = \mathbb{V}(\delta^* | X, u) + \frac{4}{p} \mathbb{V}(v^{*'} Q' A \hat{\beta} | X, u), \quad (6)$$

which is larger than the variance of our estimator, $\mathbb{V}(\delta^* | X, u)$, attributable to the presence of the last term, similarly to equation (2). While both methods yield an unbiased and consistent estimate of the direct bias correction $\hat{\delta}$, δ^* is more efficient.

Why is our method flexible and easy to implement? First, our method is flexible because it allows for a wide range of assumptions on the error's covariance matrix. It is only limited by the capacity of the bootstrap to replicate the assumed covariance matrix. Meaning, as long as there is a bootstrap that allows for a re-sampling where the conditional variance of the bootstrap errors is equal to $\hat{V}(u|X)$, then one could implement our correction.

Second, our method is easy to implement as it relies only on estimating linear regressions. Thus, our method can profit from the development of any fast estimation procedure handling linear regressions with many covariates.

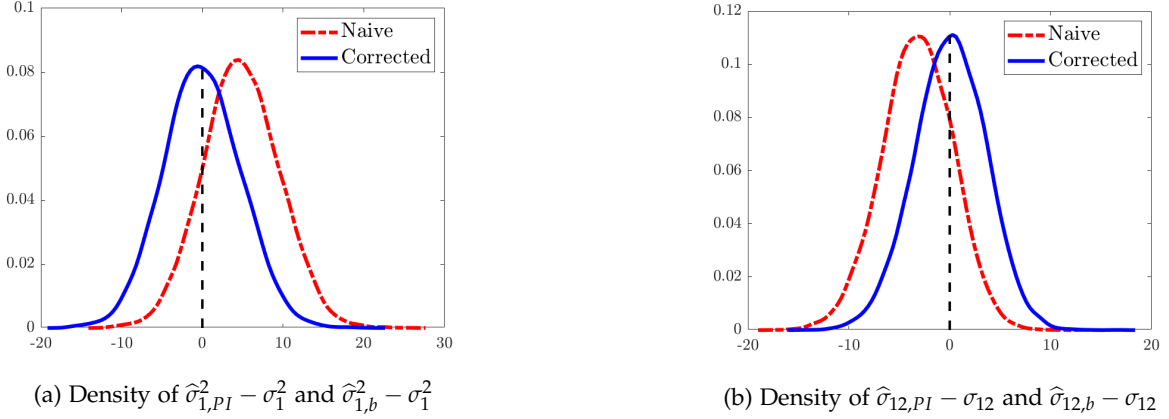
In principle, one could think that both estimating a linear regression and our bias correction deal with the problem of inverting $X'X$. The estimated coefficients of a linear regression $\hat{\beta}$ are defined as the unknowns that solve for the normal equations, $X'X\hat{\beta} = X'Y$. One could find this solution by directly inverting the matrix $X'X$, but this is not efficient. Instead, standard algorithms avoid the direct computation of such inverse.²¹

When the number of covariates is large, even standard methods have trouble in doing a linear estimation. But, with the proliferation of large datasets, there has been significant progress in the development of efficient algorithms to estimate linear models with a large number of covariates. Especially when that large number stems from fixed effects. For example, when we compare our method with existing alternatives in Section 4, we use the preconditioned conjugate gradient method in Matlab to solve for the normal equations. However, the choice of algorithm for solving the normal equations is up to the user.

²⁰These conditions would be satisfied, for example, if the error terms would be distributed normal, or, as in our applications, when we use the Rademacher errors for the bootstrap. A formal proposition of the statement and its proof can be found in the Online Appendix.

²¹In practice, standard algorithms for estimating linear regressions do a QR decomposition of X , which transforms the normal equations into an upper triangular system. This allows to use a sequential method to solve for $\hat{\beta}$. While this method directly avoids the computation of the inverse, it is however impractical when the number of covariates is large, as the number of sequential steps is equal to the number of covariates. In addition, the QR decomposition itself is computationally costlier the larger the matrix.

Figure 1: Differences between the true moments with plug-in and bootstrap estimators



Notes: The left figure presents the distributions of the differences between the true variance σ_1^2 and both, the naive plug-in estimated variance $\hat{\sigma}_{1,PI}^2$ and the bias corrected estimated variance $\hat{\sigma}_{1,b}^2$. The distribution of the difference between the true moment and the bias corrected estimated covariance is centered at zero. The right figure does the same but for the covariance σ_{12} .

3.1 Simple illustration

We illustrate the effectiveness of our bias correction method with some simple Monte Carlo simulations. The model design is the same as in equation (5) with homoscedastic errors and sample size $n = 500$. The number of covariates is $k_1 = k_2 = 200$. We keep this number relatively low to be able to compute what we dubbed previously the direct bias correction $\hat{\delta}$. We do 10,000 simulations in total. In each simulation, keeping X fixed, we draw new error terms to form the dependent variable. We estimate $\hat{\beta}$ and compute the direct bias correction terms. After the estimation, we perform $p = 100$ bootstraps and use them to compute the estimation of the bootstrap correction. We do a Wild bootstrap consistent with using HC_1 as the covariance estimator.²²

Figures 1a and 1b show the effectiveness of our method. Figure 1a plots the distribution of the difference between the plug-in estimate of the variance ($\hat{\sigma}_{1,PI}^2$) and the true variance ($\sigma_{1,PI}^2$). The figure also plots the difference between the bootstrap corrected variance ($\hat{\sigma}_{1,b}^2$) and the true variance. Figure 1b shows analogous distributions of differences between estimates and true moments but for the covariance ($\sigma_{12,PI}$). The figures show that the distribution of the differences between the plug-in estimates the true moment are not centered at zero, reflecting the bias. On the other hand the distribution of difference between the bootstrap corrected moments and the true ones are centered at zero, suggesting our method is effective in reducing the bias.

In terms of efficiency, our methods is very close to the direct correction—which is the best one can do—but outperforms more traditional bootstrap methods. Table 1 presents the mean and variance of the differences of our bootstrap method δ^* and the bootstrap following MacKinnon and Smith Jr (1998) δ_{MS}^* with respect to the direct correction $\hat{\delta}$.²³ The mean differences of our method are very small as well as the variances, meaning that the estimated bootstrap correction is performing almost as well as the direct correction in almost all simulations. The alternative bootstrap correction

²²In other words, for each observation and bootstrap iteration, we sample a Rademacher random variable and multiply it to each observation's residual times $\sqrt{N/(N-K)}$.

²³As previously stated, MacKinnon and Smith Jr (1998) propose to generate the bootstrap dependent variable as $Y^* = X\hat{\beta} + v^*$. Their correction is: $\delta_{MS}^* = \frac{1}{p} \sum_{j=1}^p (\beta_{j,MS}^* A \beta_{j,MS}^*) - \hat{\beta}' A \hat{\beta}$, where the last term is the plug-in estimate.

Table 1: Comparison Bootstrap and Direct Estimations.

	$\widehat{\delta} - \delta^*$		$\widehat{\delta} - \delta_{MS}^*$		Mean Squared Error			
	Mean	Variance	Mean	Variance	Plug-In	Direct	Boot	Boot MS
$\widehat{\text{var}}(X_1\beta_1)$	-0.00015	0.0015	-0.0037	0.08	47.78	24.34	24.34	24.44
$\widehat{\text{var}}(X_2\beta_2)$	-1.2×10^{-5}	0.0015	0.0054	0.19	79.00	55.55	55.55	55.74
$\widehat{\text{cov}}(X_1\beta_1, X_2\beta_2)$	9.3×10^{-5}	0.0014	-0.0014	0.05	25.83	15.18	15.18	15.22

The first two columns represent, respectively, the mean and the variance of the difference between the direct correction $\widehat{\delta}$ and the bootstrap correction δ^* . Columns 3 and 4 are analogous but using the bootstrap correction proposed by [MacKinnon and Smith Jr \(1998\)](#), δ_{MS}^* . Columns 5 to 8 compute the Mean Squared Error between the different estimated moments and the true ones. *Plug-In* refers to the non-corrected estimated moments using the estimates of the linear regression. *Direct* uses the estimated moments with the direct bias correction. *Boot* and *Boot MS* refer, respectively, to the moments with our bootstrap correction and with the bootstrap correction proposed by [MacKinnon and Smith Jr \(1998\)](#).

δ_{MS}^* in Columns 3 and 4 performs worse in terms of bias and variance.

Table 1 also shows the Mean Squared Error (MSE) of the different estimated moments. The MSE of naive plug-in estimators is larger than the one obtained with the directly corrected and bootstrap corrected moments. As our estimator is a noisy estimate of the direct correction, it is expected that the MSE of the corrected moments using our estimator to be larger than the directly corrected moments, although very close. In fact, to the level of rounding presented in the table, the two are indistinguishable. Also, as expected, our bootstrap has lower MSE than the alternative bootstrap corrected moments which follows [MacKinnon and Smith Jr \(1998\)](#).

4 Comparison of Methods

In this section we first compare our method to [Gaure \(2014\)](#) and [Kline et al. \(2020\)](#). Both methods aim to estimate the trace term in equation (3). In the Online Appendix we also compare our method with [Borovičková and Shimer \(2017\)](#) who propose an alternative method to estimate directly some quadratic forms without first estimating a linear model.

The differences between Gaure, KSS and our method are on the scope of error structures allowed, the covariance matrix estimation and how easy they are to apply. All three methods are in principle suited to perform corrections with homoscedastic and heteroscedastic errors. Nevertheless, Gaure implemented his bias correction method on the R package *lfe* only under the assumption of homoscedastic errors. In contrast, KSS and ourselves provide corrections under heteroscedasticity and serial correlation or clustering of the errors. Additionally, our method is the only one capable of doing multiple corrections at a time without increasing the computational cost. KSS and Gaure, on the contrary, need to solve new sets of equations in order to approximate each trace term that corresponds to any additional correction. Finally, our method can accommodate more complicated error structures with two or more dimensions of serial dependence, which would yield a non-block-diagonal covariance matrix.

4.1 Labor market simulations

An important application of two-way fixed effect models are the AKM type log wage regressions with worker and firm fixed effects. We closely follow [Card et al. \(2013\)](#) to implement the estimation

Table 2: Monte Carlo simulations. Homoscedastic errors.

	Time	Mean Squared Error (MSE $\times 10^2$)			Average
		$\hat{\sigma}_\theta^2$	$\hat{\sigma}_\psi^2$	$\hat{\sigma}_{\theta,\psi}$	
Plug-in		6.637	0.341	0.114	2.364
Gaure	17.3	0.050	0.109	0.015	0.058
Boot	0.9	0.050	0.105	0.014	0.057
KSS	1.3	0.050	0.106	0.014	0.057

Notes: *Plug-in* is the naive plug-in estimator, *Gaure* refers to the method [Gaure \(2014\)](#) implemented through the R package *lfe*, *Boot* is our method with HC_2 covariance matrix estimator, and *KSS* is the [Kline et al. \(2020\)](#) method. *Time* is the computing time in seconds. True moments are computed at the final sample for each method, i.e. largest connected set for *Gaure* and the largest leave-one-out connected set for *Boot* and *KSS*. $\hat{\sigma}_\theta^2$, $\hat{\sigma}_\psi^2$ and $\hat{\sigma}_{\theta,\psi}$ present respectively the mean squared errors (MSE) of the corrected estimates of the variance of the worker fixed effects, variance of the firm fixed effects and the covariance between worker and firm effects. All the MSE are multiplied by 100. *Average* is the average MSE (also scaled).

of the following regression model for the log of the wage of worker i at time t :

$$w_{it} = \theta_i + \psi_{J(i,t)} + q_{it}\gamma + \varepsilon_{it}, \quad (7)$$

where the function $J(i, t)$ gives the identity of the unique firm that employs worker i at time t , θ_i is a worker fixed effect, $\psi_{J(i,t)}$ is the firm $J(i, t)$ fixed effect, q_{it} are time varying observables (age and education interacted with year effects), and ε_{it} is the error term.

Equation (7) can be estimated by OLS where the person and firm fixed effects estimators have the same structure as the ones in Section 2. Thus, the second order moments exhibit a similar bias and the implementation of the correction is analogous.

We compare the correction methods by simulating many labor markets under different assumptions on the error terms and evaluate them in terms of computation time and mean squared errors. We also explore differences between the covariance estimation methods described in Section 3.

We first compare all the methods under conditional homoscedasticity of the errors. Results are in Table 2. All the methods reduce the initial bias of the plug-in estimate. *Gaure*, *KSS* and our method are very similar in terms of MSE, and even look identical after rounding the numbers up, with *Gaure* doing slightly worse.²⁴ *Gaure* is the slowest method and the bootstrap correction is also faster than *KSS*.

Table 3 presents the comparison of our method to *KSS* under conditional heteroscedasticity for different degrees of worker mobility.²⁵ Both methods reduce by more than 97% the MSE compared to the plug-in estimates in the low mobility case.²⁶ Our method is slightly more accurate for both mobility cases, and it also outperforms *KSS* in terms of time.²⁷ In the Online Appendix, we show that HC_2 estimate for the covariance matrix outperforms both HC_0 and HC_1 in terms of MSE when doing the bias correction with heteroscedastic errors.

Table 4 presents the results from a simulation with a non diagonal covariance matrix. In par-

²⁴We implement *Gaure* and *KSS* corrections as follows. We use the *bccor* command of *Gaure's lfe* R package with 300 maximum samples and tolerance of 1e-6. We run Version 3 of *KSS* Matlab code eliminating observations (instead of matches) for the leave-one-out estimation with 300 simulations to estimate the leverage and corrections at once. We run our corrections in Matlab with tolerance of 1e-6 and 300 simulations.

²⁵When workers are more mobile, the firm fixed effect estimates are less noisy. As this noise is the source of the bias of the quadratic objects, more precise estimates will yield a smaller bias as one can see for the *Plug-in* estimates in Table 3.

²⁶Table 1 in [Kline et al. \(2020\)](#) shows that their connected set is similar to our low mobility scenario with 2.7 movers per firm and average firm size of 12.

²⁷In the Online Appendix, we compare the densities of the bias for the different methods. The densities show that both corrections (*KSS* and our bootstrap) are similar but the bootstrap method has smaller variance for the reasons suggested in Section 3. We also show in the Online Appendix that the results are similar even when using a more realistic sample size of roughly 5 million observations.

Table 3: Monte Carlo simulations. Heteroscedastic errors.

Mov/firm	Model	Time	Mean Squared Error (MSE $\times 10^2$)			
			$\hat{\sigma}_\theta^2$	$\hat{\sigma}_\psi^2$	$\hat{\sigma}_{\theta,\psi}$	Average
<i>Low Mobility</i>						
3	Plug-in		22.885	7.702	6.451	12.346
3	Boot	0.7	0.224	0.665	0.192	0.360
3	KSS	1.0	0.268	0.708	0.233	0.403
<i>Mid Mobility</i>						
5	Plug-in		10.518	1.670	1.070	4.419
5	Boot	0.8	0.085	0.255	0.048	0.129
5	KSS	1.1	0.086	0.257	0.048	0.131

Notes: *Plug-in* is the naive plug-in estimator, *Boot* refers to our method with HC_2 covariance matrix estimator, and *KSS* is the Kline et al. (2020) method. True moments are computed at the leave-one-out connected set. *Mov/firm* is the number of movers per firm and the average firm has 12 employees. *Time* is the computing time in seconds. $\hat{\sigma}_\theta^2$, $\hat{\sigma}_\psi^2$ and $\hat{\sigma}_{\theta,\psi}$ present respectively the mean squared errors of the corrected estimates of the variance of the worker fixed effects, variance of the firm fixed effects and the covariance between worker and firm effects. All the MSE are multiplied by 100. *Average* is the average MSE (also scaled).

ticular, we assume that there is serial correlation of the wages within a given match and the true innovation is homoscedastic. The table compares the plug-in estimate to our bootstrap correction and to KSS. *Boot* is the best performing correction method in terms of accuracy and time. We show in the Online Appendix that the differences in performance are amplified when we have heteroscedastic innovations at the match level.

Why is our method faster? In the above simulations, our method takes between half and two thirds as much time as the method proposed by KSS.²⁸ The underlying reason is that our method needs to do *at most* two iterative procedures regardless of the number of corrections: one for estimating the leverage—for example, if one uses HC_2 for the covariance matrix estimator—and one for the bootstrap. On the other hand, KSS method needs to do, in general, the same number of iterative procedures as number of corrections plus the iteration for the leverage estimation. In some particular cases, like in the AKM model, they can reduce the minimum number of iterative procedures to three: one for the leverage and two extra for the variance of the worker fixed effects, the variance of firm fixed effects, and their covariance.²⁹

5 Application

In the application we use a panel data from the French statistical agency (INSEE) from 2002 to 2014.³⁰ Our dependent variable is (log) gross daily wage of full time employees with ages between 20 and 60 working at private firms.

The goal is to use our bootstrap method to do a bias corrected variance decomposition of log

²⁸This range is computed taking into account the simulations and the application in cases where the bootstrap correction is used with the HC_2 covariance matrix estimator and therefore both methods need to compute the leave-one-out connected set and estimate the leverage of observations. We think this is the fairest comparison. In other cases where we would not need to estimate the leverages our method would be even faster.

²⁹See section 2.3.2 in the computational appendix of KSS.

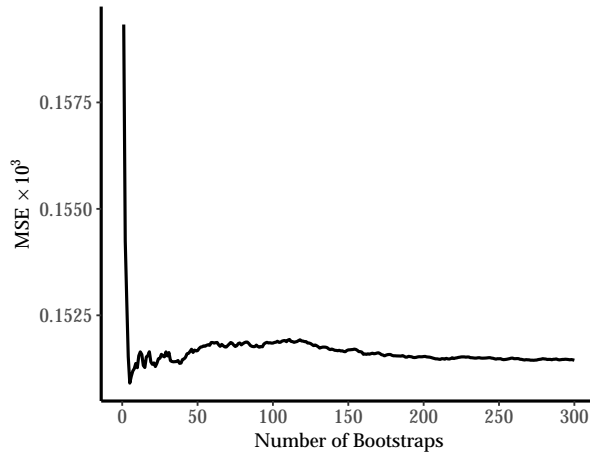
³⁰In particular we use *Panel tous salariés-EDP* that consists of a random subsample of workers with firm identifiers and socio-demographic variables. The sample consists of workers born in October on certain days. The sample size was increased in 2002 so we took this as the starting year.

Table 4: Monte Carlo simulations. Serial correlation with homoscedasticity.

	Time	Mean Squared Error ($\text{MSE} \times 10^2$)			Average
		$\hat{\sigma}_\theta^2$	$\hat{\sigma}_\psi^2$	$\hat{\sigma}_{\theta,\psi}$	
Plug-in		94.352	1.670	0.603	32.208
Boot	0.6	9.674	0.264	0.053	3.330
KSS	1.3	21.571	0.254	0.052	7.292

Notes: *Plug-in* is the naive plug-in estimator, *Boot* refers to our method with a wild block bootstrap where each match defines a block and we skip the pruning of the data. *KSS* is the Kline et al. (2020) method leaving a match out. The average firm has 10 movers and 12 employees. *Time* is the computing time in seconds. True moments are computed at the largest connected set for *Boot* and at the largest leave-one-out connected set for *KSS*. $\hat{\sigma}_\theta^2$, $\hat{\sigma}_\psi^2$ and $\hat{\sigma}_{\theta,\psi}$ present respectively the mean squared errors (MSE) multiplied by 100 of the corrected estimates of the variance of the worker fixed effects, variance of the firm fixed effects and the covariance between worker and firm effects. *Average* is the average MSE (also scaled).

Figure 2: MSE of corrected $\hat{\sigma}(\theta, \psi)$ by number of bootstraps.



Notes: This figure presents the mean squared error (MSE) of the covariance between worker-firm fixed effects $\hat{\sigma}(\theta, \psi)$ across 1000 homoscedastic error simulations. The bootstrap correction assumes a diagonal covariance matrix and we use the HC_1 covariance matrix estimator.

wages. To guide our choice of number of bootstraps, we perform some simulations with a fixed set of covariates with low mobility and simulate a thousand samples by simulating the error. With each dataset we perform corrections from one to 300 bootstraps as in the Monte Carlo simulations of Section 4. Figure 2 shows the mean squared error between the true covariance of person and firm fixed effects and the corrected one for different number of bootstraps.³¹ The figure shows that with the first 25 bootstraps the MSE reduces significantly and around 150 it flattens. This suggests that few bootstraps are enough to gain accuracy.³² In the Online Appendix we show a more formal way of choosing the number of bootstraps based on the Chebyshev's inequality.³³

Table 5 shows the variance decomposition of log wages as well as the correlation between firm and worker fixed effects using the plug-in moments and the corrected ones under the assumption of serial correlation within a match. The variance of the person and firm effects are both reduced and they explain a lower share of the total variance after the correction. The correlation becomes closer to zero and it approaches values that have been found in other countries with a larger number

³¹For all the samples we take the corrections obtained with different bootstraps and take the mean squared error against the true moment.

³²Throughout the application corrections we run corrections with 300 simulation to estimate the leverage and 1000 bootstraps to estimate the corrections of second order moments.

³³This criterion yields a quite conservative number of bootstraps. It reflects the generality of the result as this criterion would work regardless of the distribution of the error terms.

Table 5: Application. Plug-in vs corrected decomposition.

	Plug-in	Boot Serial		Plug-in	Boot Serial
$Var(y)$	0.2162	0.2162	$2Cov(\hat{\theta}_i, \hat{\psi}_j)$	-0.0325	-0.0062
$Var(\hat{\theta}_i)$	0.1688	0.1409	$2Cov(\hat{\theta}_i, \mathbf{q}\hat{\gamma})$	-0.0135	-0.0134
$Var(\hat{\psi}_j)$	0.0493	0.0319	$2Cov(\hat{\psi}_j, \mathbf{q}\hat{\gamma})$	-0.0006	-0.0006
$Var(\mathbf{q}\hat{\gamma})$	0.0146	0.0146	$Corr(\hat{\theta}_i, \hat{\psi}_j)$	-0.0964	-0.0004
$Var(\hat{\varepsilon})$	0.0301	0.0490	Obs.	5108399	5108399

Notes: *Plug-in* refers to the uncorrected estimates of each of the variance components at the largest connected set and *Boot Serial* refers to the estimates after our bootstrapped correction using a wild block bootstrap. $Var(y)$ is the variance of log wages, $Var(\hat{\theta}_i)$ the variance of worker fixed effects (naive $\hat{\sigma}_\theta^2$ or corrected $\tilde{\sigma}_\theta^2$), $Var(\hat{\psi}_j)$ is the variance of firm fixed effects, $Var(\mathbf{q}\hat{\gamma})$ is the variance of other covariates and $Var(\hat{\varepsilon})$ is the variance of the error term. The other terms of the decomposition are twice the covariances between the fixed effects and the covariates ($2Cov(\hat{\theta}_i, \hat{\psi}_j)$, $2Cov(\hat{\theta}_i, \mathbf{q}\hat{\gamma})$ and $2Cov(\hat{\psi}_j, \mathbf{q}\hat{\gamma})$). Finally, $Corr(\hat{\theta}_i, \hat{\psi}_j)$ is the estimated correlation between worker and firm fixed effects and *Obs.* is the number of observations.

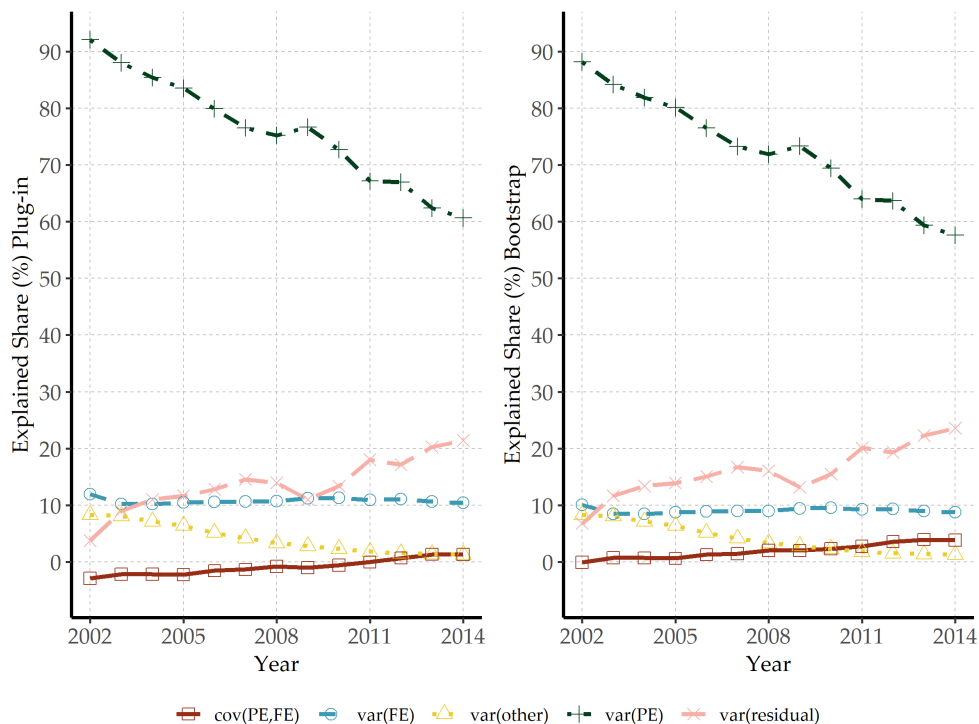
of movers per firm, which should attenuate the bias, as reported by Table 1 of [Lopes de Melo \(2018\)](#). Naturally, the variance and covariance of the person and firm effects are the moments that change the most after the correction. The reason is that the underlying estimates of the person and fixed effects are very noisy. In contrast, when the underlying estimates of a particular moment are estimated with precision, as it is in the case of the parameters $\hat{\gamma}$ associated with the common covariates \mathbf{q} , the change between the plug-in and corrected moments is negligible.

To fully exploit the benefit of our bootstrap correction method we also perform a yearly variance decomposition. In Figure 3 we compare the year-to-year evolution of the different explained shares using the naive estimated moments and the corrected ones. The main takeaway from this figure is that the correction changes the levels but not the slopes of explained shares. This leads to a change in the relative importance of each component. In particular, the corrected variance of the residuals is relatively more important than the corrected variance of the firm effect in almost every year, while both are similar when considering uncorrected variances. A very interesting trend is the decline in explanatory power of the individual fixed effects for recent years. It might be just a feature of the French data and explanations for this phenomenon are outside the scope of this paper.

The fixed effects we are estimating are constant across the entire sample. The relative importance of each set of fixed effects changes across periods because the composition in the sample changes via the exit and entry of firms and workers. Another potential reason is the variance of the residual changes across periods. The reason for changes in the relative importance of the fixed effects is different than the exercise done by [Lachowska, Mas, Saggio, and Woodbury \(2022\)](#), which allows for the estimated firm effects to change across periods. In their setting, the changes can arise because of sample composition changes and/or changes in the estimated fixed effects for the same individuals across periods. We want to emphasize that the reasons why our method is faster than KSS will also apply in such settings.

Finally, in the Online Appendix we compare our method against KSS using the French data. To make the comparison as fair as possible, we use the HC_2 covariance matrix estimator, which requires estimating the leverage of each observation as well. We also use the leave-one-out connected set sample, which is smaller than the connected set sample used in the baseline. In the application, our bootstrap method takes a little more than a quarter of the time of the KSS correction. Both methods yield slightly different estimates of second order moments that also lead to differences in

Figure 3: Application. Evolution of the explained shares.



Notes: This figure presents the year-to-year evolution of the explained shares of the total log wage variance of the plug-in and corrected estimates of the person and firm fixed effects, their covariance and the variances of other covariates and the residual.

the estimate of the correlation between firm and worker fixed effects. When assuming a diagonal covariance matrix with heteroscedastic errors (*Boot HC₂* and *KSS*), both methods yield a positive estimate for the correlation between firm and worker fixed effects. The estimate of this correlation with the bootstrap correction is 0.14 while the one of *KSS* is 0.11. Also, when assuming that the error terms are correlated at the match level, the estimated correlation between worker and firm fixed effects is positive with both methods when using the leave-one-out connected set.³⁴ The estimated correlation with the bootstrap is 0.15 and 0.33 with *KSS*.³⁵ We can conclude that, even after correcting for the limited-mobility bias, the estimates of the correlation between workers and firms fixed effects are sensitive to sample selection. The leave-one-out connected set is comprised of more mobile workers who could have a different sorting pattern than the rest of the labor force. Thus, it could be that the suggested small, yet positive correlation, is driven by those workers who change jobs more frequently.

6 Conclusion

In this paper, we propose a computationally feasible bootstrap method to correct for the small-sample bias found in all quadratic forms in the parameters of linear models with a very large number of covariates. We show using Monte Carlo simulations that the method is effective at reducing the bias. The application to French labor market data shows that the correction increases the correlation between firm and worker fixed effects. Depending on the sample and on the speci-

³⁴The results from *Boot Serial (Connected)* are analogous to the ones in Table 5 but computed with residualized log wages.

³⁵The column *Boot Serial* in the Online Appendix is computed at the leave-one-out connected set for comparison of the estimates.

fication, our bias correction method changes the sign of that correlation and in all cases it changes the relative importance of the different components in explaining the variance of log wages.

The only requirements to implement our correction is to have a bootstrap procedure that is consistent with the assumption on the variance-covariance matrix of the error term and to estimate the model several times. The correction can thus be applied easily to any study running an AKM type regression or multi-way fixed effects regressions. Our method is faster than [Kline et al. \(2020\)](#) and as accurate in the simulations. Besides the speed, another advantage of our approach is that it allows to increase the number of moments to correct without increasing the computational costs. Plus, it allows to consider more elaborated variance structures. In the case of heteroscedastic errors, although KSS estimator would yield to unbiased estimates of the bias correction, we find that using HC_2 for the covariance matrix estimator leads, in general, to more efficient bias correction estimators.

APPENDIX FOR PUBLICATION

A Proofs

Proposition 1.

Proof. By the linearity of the trace and expectation operators we have that

$$\mathbb{E}(\widehat{\delta}|X) = \mathbb{E} \left(\text{trace} \left(Q' A Q \widehat{\mathbb{V}}(u|X) \right) | X \right) = \text{trace} \left(Q' A Q \mathbb{E} \left(\widehat{\mathbb{V}}(u|X) | X \right) \right) = \text{trace} (Q' A Q \mathbb{V}(u|X)) = \delta$$

□

Proposition 2.

Proof. First, note that for any bootstrap estimate j of the quadratic form $\beta_j^{*'} A \beta_j^*$ we have that

$$\beta_j^{*'} A \beta_j^* = (v_j^*)' Q' A Q v_j^*.$$

Under the bootstrap, i.e. conditional on X and u , the only source of randomness is v_j^* . Taking expectations under the bootstrap of $\beta_j^{*'} A \beta_j^*$, conditionally on X and u , and using the assumption $\mathbb{E}(v_j^* | X, u) = 0$, we get

$$\mathbb{E}_{v^*} \left(\beta_j^{*'} A \beta_j^* \mid X, u \right) = \text{trace} \left(Q' A Q \mathbb{V}(v_j^* | X, u) \right).$$

By assumption $\mathbb{V}(v_j^* | X, u) = \widehat{\mathbb{V}}(u|x)$, then $\mathbb{E}_{v^*} \left(\beta_j^{*'} A \beta_j^* \mid X, u \right) = \widehat{\delta}$.

Unbiased. Taking expectations over $\delta^* \equiv \frac{1}{p} \sum_{j=1}^p \left(\beta_j^{*'} A \beta_j^* \right)$, conditionally on X and u we obtain

$$\mathbb{E}_{v^*}(\delta^* | X, u) = \frac{1}{p} \sum_{j=1}^p \mathbb{E}_{v^*} \left(\beta_j^{*'} A \beta_j^* \mid X, u \right) = \frac{1}{p} \sum_{j=1}^p \widehat{\delta} = \widehat{\delta}.$$

Consistent. From the definition of $\delta^* \equiv \frac{1}{p} \sum_{j=1}^p \left(\beta_j^{*'} A \beta_j^* \right)$, we have that

$$\frac{1}{p} \sum_{j=1}^p \left(\beta_j^{*'} A \beta_j^* \right) \xrightarrow{p} \mathbb{E}_{v^*} \left(\beta_i^{*'} A \beta_i^* \mid X, u \right) = \widehat{\delta}.$$

□

Corollary 1

Proof. Using the Law of Iterated Expectations we get

$$\mathbb{E}(\delta^* | X) = \mathbb{E}_u \left(\mathbb{E}_{v^*}(\delta^* | X, u) \mid X \right) = \mathbb{E}_u(\widehat{\delta} | X) = \delta.$$

□

References

- ABOWD, J., F. KRAMARZ, P. LENGERMANN, AND S. PÉREZ-DUARTE (2004): "Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States," *Unpublished Manuscript*.
- ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.
- ANDREWS, M. J., L. GILL, T. SCHANK, AND R. UPWARD (2008): "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171, 673–697.
- ANGRIST, J. D. AND J.-S. PISCHKE (2008): *Mostly harmless econometrics: An empiricist's companion*, Princeton university press.
- AVRON, H. AND S. TOLEDO (2011): "Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix," *Journal of the ACM (JACM)*, 58, 8.
- BONHOMME, S., K. HOLZHEU, T. LAMADON, E. MANRESA, M. MOGSTAD, AND B. SETZLER (2023): "How Much Should we Trust Estimates of Firm Effects and Worker Sorting?" *Journal of Labor Economics*, 41.
- BONHOMME, S., T. LAMADON, AND E. MANRESA (2019): "A distributional framework for matched employer employee data," *Econometrica*, 87, 699–739.
- BOROVÍČKOVÁ, K. AND R. SHIMER (2017): "High wage workers work for high wage firms," Tech. rep., National Bureau of Economic Research.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2008): "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 90, 414–427.
- 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, S13–S70.
- CARD, D., J. HEINING, AND P. KLINE (2013): "Workplace heterogeneity and the rise of West German wage inequality," *The Quarterly journal of economics*, 128, 967–1015.
- 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, 1593–1660.
- DHAENE, G. AND K. JOCHMANS (2015): "Split-panel jackknife estimation of fixed-effect models," *The Review of Economic Studies*, 82, 991–1030.
- DRISCOLL, J. C. AND A. C. KRAAY (1998): "Consistent covariance matrix estimation with spatially dependent panel data," *Review of economics and statistics*, 80, 549–560.
- GAURE, S. (2013): "lfe: Linear group fixed effects," *The R Journal*, 5, 104–117.
- (2014): "Correlation bias correction in two-way fixed-effects linear regression," *Stat*, 3, 379–390.
- HAYASHI, F. (2000): "Econometrics. 2000," *Princeton University Press. Section, 1*, 60–69.
- HUTCHINSON, M. (1989): "A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines," *Communications in Statistics-Simulation and Computation*, 18, 1059–1076.
- IRANZO, S., F. SCHIVARDI, AND E. TOSETTI (2008): "Skill dispersion and firm productivity: An analysis with employer-employee matched data," *Journal of Labor Economics*, 26, 247–285.
- JOCHMANS, K. (2018): "Heteroskedasticity-robust inference in linear regression models," *arXiv preprint arXiv:1809.06136*.
- KLINE, P., R. SAGGIO, AND M. SØLVSTEN (2020): "Leave-out estimation of variance components," *Econometrica*, 88, 1859–1898.
- LACHOWSKA, M., A. MAS, R. SAGGIO, AND S. A. WOODBURY (2022): "Do firm effects drift? Evidence from Washington administrative data," *Journal of Econometrics*.

- LOPES DE MELO, R. (2018): "Firm wage differentials and labor market sorting: Reconciling theory and evidence," *Journal of Political Economy*, 126, 000–000.
- MACKINNON, J. G. AND A. A. SMITH JR (1998): "Approximate bias correction in econometrics," *Journal of Econometrics*, 85, 205–230.
- MACKINNON, J. G. AND H. WHITE (1985): "Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties," *Journal of econometrics*, 29, 305–325.
- ROODMAN, D., M. Ø. NIELSEN, J. G. MACKINNON, AND M. D. WEBB (2019): "Fast and wild: Bootstrap inference in Stata using boottest," *The Stata Journal*, 19, 4–60.
- SONG, J., D. J. PRICE, F. GUVENEN, N. BLOOM, AND T. VON WACHTER (2019): "Firming up inequality," *The Quarterly journal of economics*, 134, 1–50.
- SORKIN, I. (2018): "Ranking firms using revealed preference," *The quarterly journal of economics*, 133, 1331–1393.