

MacKinnon, James G.; Nielsen, Morten Ørregaard; Webb, Matthew D.

Working Paper

Leverage, influence, and the jackknife in clustered regression models: Reliable inference using summclust

Queen's Economics Department Working Paper, No. 1483

Provided in Cooperation with:

Queen's University, Department of Economics (QED)

Suggested Citation: MacKinnon, James G.; Nielsen, Morten Ørregaard; Webb, Matthew D. (2022) : Leverage, influence, and the jackknife in clustered regression models: Reliable inference using summclust, Queen's Economics Department Working Paper, No. 1483, Queen's University, Department of Economics, Kingston (Ontario)

This Version is available at:

<https://hdl.handle.net/10419/260488>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Queen's Economics Department Working Paper No. 1483

Leverage, Influence, and the Jackknife in Clustered Regression Models: Reliable Inference Using `summclust`

James G. MacKinnon
Queen's University

Morten Ørregaard Nielsen
Aarhus University

Matthew D. Webb
Carleton University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

3-2022

5-2022 (minor revisions)

Leverage, Influence, and the Jackknife in Clustered Regression Models: Reliable Inference Using `summclust`*

James G. MacKinnon[†] Morten Ørregaard Nielsen
Queen's University Aarhus University
`mackinno@queensu.ca` `mon@econ.au.dk`

Matthew D. Webb
Carleton University
`matt.webb@carleton.ca`

May 5, 2022

Abstract

Cluster-robust inference is widely used in modern empirical work in economics and many other disciplines. When data are clustered, the key unit of observation is the cluster. We propose measures of “high-leverage” clusters and “influential” clusters for linear regression models. The measures of leverage and partial leverage, and functions of them, can be used as diagnostic tools to identify datasets and regression designs in which cluster-robust inference is likely to be challenging. The measures of influence can provide valuable information about how the results depend on the data in the various clusters. We also show how to calculate two jackknife variance matrix estimators, CV_3 and CV_{3J} , as a byproduct of our other computations. All these quantities, including the jackknife variance estimators, are computed in a new Stata package called `summclust` that summarizes the cluster structure of a dataset.

Keywords: clustered data, cluster-robust variance estimator, grouped data, high-leverage clusters, influential clusters, jackknife, partial leverage, robust inference.

JEL Codes: C10, C12, C21, C23, C87.

*We are grateful to David Drukker for a very insightful suggestion. We are also grateful to seminar participants at York University. MacKinnon and Webb thank the Social Sciences and Humanities Research Council of Canada (SSHRC grants 435-2016-0871 and 435-2021-0396) for financial support. Nielsen thanks the Danish National Research Foundation for funding a DNRF Chair grant.

[†]Corresponding author. Address: Department of Economics, 94 University Avenue, Queen's University, Kingston, Ontario K7L 3N6, Canada. Email: `mackinno@queensu.ca`. Tel. 613-533-2293. Fax 613-533-6668.

1 Introduction

It is now standard in many fields of economics and other disciplines to employ cluster-robust inference for the parameters of linear regression models. In the most common case, each of the N observations is assigned to one of G disjoint clusters, which might correspond to, for example, families, schools, villages, hospitals, firms, industries, years, cities, counties, or states. The assignment of observations to clusters is assumed to be known, and observations in different clusters are assumed to be independent, but any pattern of heteroskedasticity and/or dependence is allowed within each cluster. Under these assumptions, it is easy to compute cluster-robust variance matrices that yield asymptotically valid t -tests, Wald tests, and confidence intervals. However, even when N is very large, the resulting inferences may be unreliable when G is not large or the clusters are not sufficiently homogeneous.

The literature on cluster-robust inference has grown rapidly in recent years. [Cameron and Miller \(2015\)](#) is a classic survey article. [Conley, Gonçalves and Hansen \(2018\)](#) surveys a broader class of methods for various types of dependent data. [MacKinnon, Nielsen and Webb \(2022a\)](#) explores the implications of key theoretical results for empirical practice. As that paper discusses, there are at least two situations in which cluster-robust t -tests and Wald tests are at risk of over-rejecting to an extreme extent, even when G is not small. The first is when one or a few clusters are much larger than the rest, and the second is when the only “treated” observations belong to just a few clusters; [Djogbenou, MacKinnon and Nielsen \(2019\)](#) discusses the first case, and [MacKinnon and Webb \(2017a, b, 2018\)](#) discuss the second. In both of these cases, one cluster (or a few of them) has high leverage, in the sense that omitting this cluster has the potential to change the OLS estimates substantially. When that actually happens, a cluster is said to be influential.

The concepts of leverage and influence are normally applied at the observation level ([Belsley, Kuh and Welsch 1980](#)), but they are equally applicable at the cluster level. Just as high-leverage observations can make heteroskedasticity-robust inference unreliable ([Chesher 1989](#)), so too can high-leverage clusters make cluster-robust inference unreliable. Just as highly influential observations may lead us to suspect that there is something wrong with the model or the data, so too may highly influential clusters. Although the two special cases discussed in the preceding paragraph are particularly important, any situation in which a few clusters have high leverage or high influence should be worrying.

Note that there are at least two different concepts of leverage. The usual one focuses on fitted values or, equivalently, residuals. A cluster is said to have high leverage if removing it has the potential to change the fitted values for that cluster by a lot. It can also be of interest to study partial leverage ([Cook and Weisberg 1980](#)). A cluster is said to have high

partial leverage for the j^{th} coefficient if removing that cluster has the potential to change the estimate of the j^{th} coefficient by a lot.

In this paper, we provide computationally efficient ways to identify high-leverage and influential clusters. These can be used as diagnostic tools to determine how reliable various methods for cluster-robust inference are likely to be. Of course, it is easy enough to check cluster sizes and the distribution of treatment status across clusters directly, and many careful empiricists already do these things. But whether a cluster has high leverage, is influential, or has both of these features can depend on the sample in rather complicated ways. In our view, it is best to investigate leverage, partial leverage, and influence at the cluster level directly. We provide a new Stata package called `summclust` that implements our procedures in order to summarize the cluster structure of a dataset.

A byproduct of our calculations is that they make it very fast to compute two jackknife variance matrix estimators called CV_3 and CV_{3J} . These often seem to perform much better in finite samples than other cluster-robust variance estimators (MacKinnon, Nielsen and Webb 2022b). Although these estimators have been known for some time (Bell and McCaffrey 2002), perhaps under different names, they have rarely been used because they were too expensive to compute. The `summclust` package can calculate both of them.

The remainder of the paper is organized as follows. In the next section, we review cluster-robust inference for the linear regression model. [Section 3](#) introduces our new measures of leverage, partial leverage, and influence at the cluster level. [Section 4](#) discusses several special cases in which some or all of these measures can be determined analytically. [Section 5](#) shows how our results can be used to compute the CV_3 and CV_{3J} jackknife variance matrix estimators. [Section 6](#) discusses what quantities should be reported, both for the case of one-way clustering on which the paper focuses and for the case of two or more clustering dimensions. [Section 7](#) describes some simulation experiments which suggest that it may be desirable to report many of these quantities. [Section 8](#) describes the `summclust` package. [Section 9](#) presents an empirical illustration in which our methods are highly informative, and [Section 10](#) concludes.

2 The Linear Regression Model with Clustering

Most of the paper deals with the linear regression model

$$\mathbf{y}_g = \mathbf{X}_g\boldsymbol{\beta} + \mathbf{u}_g, \quad g = 1, \dots, G, \tag{1}$$

where the data have been divided into G disjoint clusters. The g^{th} cluster has N_g observations, so that the sample size is $N = \sum_{g=1}^G N_g$. In (1), \mathbf{X}_g is an $N_g \times k$ matrix of regressors, $\boldsymbol{\beta}$ is a k -vector of coefficients, \mathbf{y}_g is an N_g -vector of observations on the regressand, and \mathbf{u}_g is an N_g -vector of disturbances (or error terms). The \mathbf{X}_g may of course be stacked into an $N \times k$ matrix \mathbf{X} , and likewise the \mathbf{y}_g and \mathbf{u}_g may be stacked into N -vectors \mathbf{y} and \mathbf{u} , so that (1) can be rewritten as $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}$.

Dividing the sample into clusters only becomes meaningful if we make assumptions about the disturbance vectors \mathbf{u}_g and, consequently, the score vectors $\mathbf{s}_g = \mathbf{X}_g^\top \mathbf{u}_g$. For a correctly specified model, $E(\mathbf{s}_g) = \mathbf{0}$ for all g . We further assume that

$$E(\mathbf{s}_g \mathbf{s}_g^\top) = \boldsymbol{\Sigma}_g \quad \text{and} \quad E(\mathbf{s}_g \mathbf{s}_{g'}^\top) = \mathbf{0}, \quad g, g' = 1, \dots, G, \quad g' \neq g, \quad (2)$$

where $\boldsymbol{\Sigma}_g$ is the symmetric, positive semidefinite variance matrix of the scores for the g^{th} cluster. The second assumption in (2) is crucial. It says that the scores for every cluster are uncorrelated with the scores for every other cluster. We take the number of clusters G and the allocation of observations to clusters as given. The important issue of how to choose the clustering structure, perhaps by testing for the correct level of clustering, is discussed in detail in [MacKinnon, Nielsen and Webb \(2020\)](#).

The OLS estimator of $\boldsymbol{\beta}$ is

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} = \boldsymbol{\beta}_0 + (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{u},$$

where the second equality depends on the assumption that the data are actually generated by (1) with true value $\boldsymbol{\beta}_0$. It follows that

$$\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 = (\mathbf{X}^\top \mathbf{X})^{-1} \sum_{g=1}^G \mathbf{X}_g^\top \mathbf{u}_g = \left(\sum_{g=1}^G \mathbf{X}_g^\top \mathbf{X}_g \right)^{-1} \sum_{g=1}^G \mathbf{s}_g. \quad (3)$$

From the rightmost expression in (3), we see that the distribution of $\hat{\boldsymbol{\beta}}$ depends on the disturbance subvectors \mathbf{u}_g only through the distribution of the score vectors \mathbf{s}_g . Asymptotic inference commonly uses the empirical score vectors $\hat{\mathbf{s}}_g = \mathbf{X}_g^\top \hat{\mathbf{u}}_g$, in which the \mathbf{u}_g are replaced by the residual subvectors $\hat{\mathbf{u}}_g$, to estimate the variance matrix of the \mathbf{s}_g . This should work well if the sum of the \mathbf{s}_g , suitably normalized, is well approximated by a multivariate normal distribution with mean zero, and if the \mathbf{s}_g are well approximated by the $\hat{\mathbf{s}}_g$. However, this type of asymptotic inference can be misleading when either of these approximations is poor.

It follows immediately from (3) that an estimator of the variance of $\hat{\boldsymbol{\beta}}$ should be based

on the usual sandwich formula,

$$(\mathbf{X}^\top \mathbf{X})^{-1} \left(\sum_{g=1}^G \boldsymbol{\Sigma}_g \right) (\mathbf{X}^\top \mathbf{X})^{-1}. \quad (4)$$

The natural way to estimate (4) is to replace the $\boldsymbol{\Sigma}_g$ matrices by their empirical counterparts, that is, the $\hat{\mathbf{s}}_g \hat{\mathbf{s}}_g^\top$. If, in addition, we multiply by a correction for degrees of freedom, we obtain the cluster-robust variance estimator, or CRVE,

$$\text{CV}_1: \quad \frac{G(N-1)}{(G-1)(N-k)} (\mathbf{X}^\top \mathbf{X})^{-1} \left(\sum_{g=1}^G \hat{\mathbf{s}}_g \hat{\mathbf{s}}_g^\top \right) (\mathbf{X}^\top \mathbf{X})^{-1}. \quad (5)$$

This is by far the most widely used CRVE in practice, and it is the default one implemented in **Stata**; alternatives to this estimator will be discussed in [Section 5](#). When $G = N$, the CV_1 estimator reduces to the familiar HC_1 estimator ([MacKinnon and White 1985](#)) that is robust only to heteroskedasticity of unknown form.

The fundamental unit of inference when the observations are clustered is the cluster, not the observation; this is evident from (3), (4), and (5). The asymptotic theory for cluster-robust inference has been analyzed recently by [Djogbenou et al. \(2019\)](#) and [Hansen and Lee \(2019\)](#) under the assumption that $G \rightarrow \infty$. In general, the quality of the asymptotic approximation is determined by the number of clusters G and the extent of heterogeneity of the score vectors ([MacKinnon et al. 2022a](#)). The more the distributions of the scores vary across clusters, the worse the asymptotic approximation will likely be. Heterogeneity can arise from variation in cluster sizes and/or from variation in the distributions of the disturbances, the regressors, or both. As we discuss in [Sections 3, 6 and 7](#), leverage, partial leverage, and summary statistics based on them provide useful measures of heterogeneity across clusters.

Inference about $\boldsymbol{\beta}$ is typically based on cluster-robust t -statistics and Wald statistics. If β_j denotes the j^{th} element of $\boldsymbol{\beta}$ and β_{0j} is its value under the null hypothesis, then the appropriate t -statistic is

$$t_j = \frac{\hat{\beta}_j - \beta_{0j}}{\text{s.e.}(\hat{\beta}_j)}, \quad (6)$$

where $\hat{\beta}_j$ is the OLS estimate, and $\text{s.e.}(\hat{\beta}_j)$ is the square root of the j^{th} diagonal element of (5). Under extremely strong assumptions ([Bester, Conley and Hansen 2011](#)), it can be shown that t_j asymptotically follows the $t(G-1)$ distribution. Conventional inference in **Stata** and other programs is based on this distribution.

As the articles cited in the second paragraph of [Section 1](#) discuss, inference based on t_j and the $t(G-1)$ distribution can be unreliable when G is small and/or the clusters are

severely heterogeneous. This is true to an even greater extent for Wald tests of two or more restrictions (Pustejovsky and Tipton 2018). The measures of leverage and partial leverage at the cluster level that we introduce in the next section may help to identify the sort of heterogeneity that is likely to make inference unreliable.

Instead of using the $t(G - 1)$ distribution, we can obtain both P values for t_j and confidence intervals for β_j by employing the wild cluster restricted (or WCR) bootstrap (Cameron, Gelbach and Miller 2008). It can sometimes provide much more reliable inferences than the conventional approach; see Section 7 and Djogbenou et al. (2019). Roodman, MacKinnon, Nielsen and Webb (2019) describes a computationally efficient implementation of this method in the Stata package `boottest`. When G is reasonably large and the clusters are not very heterogeneous, inferences based on the WCR bootstrap and inferences based on CV_1 t -statistics combined with the $t(G - 1)$ distribution will often be very similar. When they differ noticeably, neither should be relied upon without further investigation.

Section 5 discusses two other CRVEs, which we refer to as CV_3 and CV_{3J} , that are both based on the cluster jackknife. In practice, these estimators are often extremely similar. CV_3 and CV_{3J} tend to yield more reliable inferences in finite samples than does CV_1 , especially when the clusters are quite heterogeneous; see MacKinnon et al. (2022b) and Section 7. Based on this simulation evidence, we recommend computing either CV_3 or CV_{3J} essentially all the time. This is easy to do using `sumclust`.

3 Identifying High-Leverage and Influential Clusters

At the observation level, there are three classic measures of heterogeneity, namely, leverage, partial leverage, and influence (Belsley et al. 1980; Chatterjee and Hadi 1986). In this section, we propose analogous measures at the cluster level.

Measures of leverage at the observation level are based on how much the residual for observation i changes when that observation is omitted from the regression. If h_i denotes the i^{th} diagonal element of the “hat matrix” $\mathbf{H} = \mathbf{P}_X = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$, then omitting the i^{th} observation changes the i^{th} residual from \hat{u}_i to $\hat{u}_i/(1 - h_i)$. Because $0 < h_i < 1$, this delete-one residual is always larger in absolute value than \hat{u}_i . The factor by which the delete-one residual exceeds \hat{u}_i increases with h_i . Since the average of the h_i is k/N , observations with values of h_i substantially larger than k/N may reasonably be said to have high leverage.

Similarly, dropping the g^{th} cluster when we estimate β changes the g^{th} residual vector from $\hat{\mathbf{u}}_g$ to $(\mathbf{I} - \mathbf{H}_g)^{-1} \hat{\mathbf{u}}_g$, where

$$\mathbf{H}_g = \mathbf{X}_g(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}_g^\top \quad (7)$$

is the $N_g \times N_g$ diagonal block of \mathbf{H} that corresponds to cluster g . The matrix \mathbf{H}_g is the cluster analog of the scalar h_i . Of course, it is not feasible to report the \mathbf{H}_g . In fact, when any of the clusters is sufficiently large, even computing and storing these matrices may be challenging. As a measure of leverage, we therefore suggest using a matrix norm of the \mathbf{H}_g . Specifically, we suggest the scalar

$$L_g = \text{Tr}(\mathbf{H}_g) = \text{Tr}(\mathbf{X}_g^\top \mathbf{X}_g (\mathbf{X}^\top \mathbf{X})^{-1}). \quad (8)$$

When the g^{th} cluster contains just one observation, say the i^{th} , then $L_g = h_i$. Thus, in this special case, the leverage measure that we are proposing reduces to the usual measure of leverage at the observation level.

The trace in (8) is the nuclear norm of the matrix \mathbf{H}_g . In general, the nuclear norm of a matrix \mathbf{A} is the sum of the singular values of \mathbf{A} . When \mathbf{A} is symmetric and positive semidefinite, the singular values are equal to the eigenvalues, which are non-negative. Since the trace of any square matrix is equal to the sum of the eigenvalues, the trace of a symmetric and positive semidefinite matrix is also its nuclear norm. In principle, we could report any norm of the \mathbf{H}_g matrices, but the nuclear norm is particularly easy to compute, and it has the advantage of being linear. This implies that we can sum over g and take the sum inside the norm just as if the \mathbf{H}_g were scalars. Since $\sum_{g=1}^G \mathbf{X}_g^\top \mathbf{X}_g = \mathbf{X}^\top \mathbf{X}$, this result means that $G^{-1} \sum_{g=1}^G \text{Tr}(\mathbf{H}_g) = k/G$, which is analogous to the result that the average of the h_i over all observations is k/N .

High-leverage clusters can be identified by comparing the L_g to k/G , their average. If, for some cluster h , L_h is substantially larger than k/G , then cluster h may be said to have high leverage. Just how much larger L_h has to be is a matter of judgement. A cluster with $L_h = 2k/G$ probably does not qualify, but a cluster with $L_h = 5k/G$ probably does. Cluster h can have high leverage either because N_h is considerably larger than G/N or because the matrix \mathbf{X}_h is somehow extreme relative to the other \mathbf{X}_g matrices, or both. We can compare the leverage of any two clusters by forming ratios. For example, if $L_1 = 3$ and $L_2 = 1$, then we can say that the first cluster has three times the leverage of the second cluster.

Like the classic leverage measure h_i , the one we suggest in (8) shows the potential impact of a particular cluster on residuals and fitted values, but not on any particular regression coefficient. When interest focuses on just one such coefficient, say the j^{th} , it may be more interesting to calculate the partial leverage of each cluster. The concept of partial leverage was introduced, for individual observations, in [Cook and Weisberg \(1980\)](#). Let

$$\hat{\mathbf{x}}_j = (\mathbf{I} - \mathbf{X}_{[j]} (\mathbf{X}_{[j]}^\top \mathbf{X}_{[j]})^{-1} \mathbf{X}_{[j]}^\top) \mathbf{x}_j, \quad (9)$$

where \mathbf{x}_j is the vector of observations on the j^{th} regressor, and $\mathbf{X}_{[j]}$ is the matrix of observations on all the other regressors. Thus $\acute{\mathbf{x}}_j$ denotes \mathbf{x}_j after all the other regressors have been partialled out. The partial leverage of observation i is simply the i^{th} diagonal element of the matrix $\acute{\mathbf{x}}_j(\acute{\mathbf{x}}_j^\top \acute{\mathbf{x}}_j)^{-1} \acute{\mathbf{x}}_j^\top$, which is just $\acute{x}_{ji}^2/(\acute{\mathbf{x}}_j^\top \acute{\mathbf{x}}_j)$, where \acute{x}_{ji}^2 is the i^{th} element of $\acute{\mathbf{x}}_j$.

The analogous measure of partial leverage for cluster g is

$$L_{gj} = \frac{\acute{\mathbf{x}}_{gj}^\top \acute{\mathbf{x}}_{gj}}{\acute{\mathbf{x}}_j^\top \acute{\mathbf{x}}_j}, \quad (10)$$

where $\acute{\mathbf{x}}_{gj}$ is the subvector of $\acute{\mathbf{x}}_j$ corresponding to the g^{th} cluster. This is what (8) reduces to if we replace \mathbf{X} and \mathbf{X}_g by $\acute{\mathbf{x}}_j$ and $\acute{\mathbf{x}}_{gj}$, respectively. It is easy to calculate the partial leverage for every cluster for any coefficient of interest. The average of the L_{gj} is evidently $1/G$, so that if cluster h has $L_{hj} \gg 1/G$, it has high partial leverage for the j^{th} coefficient. Moreover, as we will see in Section 7, the empirical distribution of the L_{gj} across clusters seems to provide useful diagnostic information.

One possible consequence of heterogeneity is that the estimates may change a lot when certain clusters are deleted. It can therefore be illuminating to delete one cluster at a time, so as to see how influential each cluster is. To do this in a computationally efficient manner, we compute the cluster-level matrices and vectors

$$\mathbf{X}_g^\top \mathbf{X}_g \quad \text{and} \quad \mathbf{X}_g^\top \mathbf{y}_g, \quad g = 1, \dots, G. \quad (11)$$

This is almost costless if we calculate the quantities in (11) first and then use them to construct $\mathbf{X}^\top \mathbf{X}$ and $\mathbf{X}^\top \mathbf{y}$. The vector of least squares estimates when cluster g is deleted is then

$$\hat{\boldsymbol{\beta}}^{(g)} = (\mathbf{X}^\top \mathbf{X} - \mathbf{X}_g^\top \mathbf{X}_g)^{-1} (\mathbf{X}^\top \mathbf{y} - \mathbf{X}_g^\top \mathbf{y}_g). \quad (12)$$

Unless k is extremely large, it should generally be fairly inexpensive to compute $\hat{\boldsymbol{\beta}}^{(g)}$ for every cluster using (12). Especially when they vary a lot, the $\hat{\boldsymbol{\beta}}^{(g)}$ can reveal a great deal about the sample. In addition, as we shall see in Section 5, they may be used to calculate jackknife variance matrices.

When there is a parameter of particular interest, say β_j , it may be a good idea to report the $\hat{\beta}_j^{(g)}$ for $g = 1, \dots, G$ in either a histogram or a table. If $\hat{\beta}_j^{(h)}$ differs greatly from $\hat{\beta}_j$ for some cluster h , then cluster h is evidently influential. In a few extreme cases, there may be a cluster h for which it is impossible to compute $\hat{\beta}_j^{(h)}$. This will happen, for example, when the regressor corresponding to β_j is a treatment dummy and cluster h is the only treated one. This is an extreme example of the problem of few treated clusters, and inferences based on either the $t(G-1)$ distribution or the WCR bootstrap are likely to be completely unreliable

in this case (MacKinnon and Webb 2017b, 2018, 2020).

Identifying influential clusters by comparing the $\hat{\beta}^{(g)}$ from (12) with $\hat{\beta}$ is very similar to identifying influential observations using the classic methods discussed in Belsley et al. (1980) and Chatterjee and Hadi (1986); for an interesting recent extension, see Broderick, Giordano and Meager (2021). If influential clusters have been identified, the investigator will need to decide how to deal with them. Perhaps the influential clusters are atypical in ways that should make us mistrust the results, either because the model does not seem to apply to them or because there are measurement errors or observations that have been miscoded.

Regression models often include cluster fixed effects. When one of the regressors is a fixed-effect dummy for cluster g , the matrices $\mathbf{X}_g^\top \mathbf{X}_g$ and $\mathbf{X}^\top \mathbf{X} - \mathbf{X}_g^\top \mathbf{X}_g$ are singular. This will cause the calculation in (12) to fail unless a generalized inverse routine, such as the `invsym` routine in Mata, is used. The problem can be avoided, and some computer time saved, by partialing out the fixed-effect dummies prior to computing the cluster-level matrices and vectors in (11) using, for example, the `areg` procedure in Stata. Partialing these out means replacing \mathbf{X} and \mathbf{y} by $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{y}}$, the deviations from their cluster means. For example, the element of $\tilde{\mathbf{y}}$ corresponding to the j^{th} observation in the g^{th} cluster is $y_{g,j} - N_g^{-1} \sum_{i=1}^{N_g} y_{g,i}$. The g^{th} subvector of $\tilde{\mathbf{y}}$ is $\tilde{\mathbf{y}}_g$, and the g^{th} submatrix of $\tilde{\mathbf{X}}$ is $\tilde{\mathbf{X}}_g$. Since there is just one fixed effect per cluster, $\tilde{\mathbf{y}}_g$ depends solely on \mathbf{y}_g , and $\tilde{\mathbf{X}}_g$ depends solely on \mathbf{X}_g . The calculations in (8) and (12) are now based on $\tilde{\mathbf{X}}^\top \tilde{\mathbf{X}}$, $\tilde{\mathbf{X}}^\top \tilde{\mathbf{y}}$, the $\tilde{\mathbf{X}}_g^\top \tilde{\mathbf{X}}_g$, and the $\tilde{\mathbf{X}}_g^\top \tilde{\mathbf{y}}_g$.

We believe that investigators should routinely compute both the L_g and the L_{gj} for any coefficient(s) of particular interest. In some cases, the L_g and the L_{gj} will be roughly proportional to the N_g (the cluster sizes). That in itself would be informative. It may be even more interesting, however, if the relative size of L_h and/or L_{hj} for some cluster(s) h is much larger, or much smaller, than the relative size of N_h .

When G is small, it may be feasible to report the L_g , along with the L_{gj} and the $\hat{\beta}_j^{(g)}$ for one or a very few values of j , in a table. When G is not small, it may be interesting to report the empirical distributions of these quantities, either in tabular or graphical form. Measures of leverage and partial leverage will often help to identify cases in which inference, especially asymptotic inference, may be unreliable; see Section 7. Unlike the leverage measures, the $\beta_j^{(g)}$ may be either positive or negative, must depend on the \mathbf{y}_g , and necessarily vary across clusters. They may sometimes turn up features of the model or dataset that require further investigation, including data errors.

4 Simple Examples

In this section, we discuss a number of simple examples in which it is possible to calculate our measures of leverage and influence analytically. These examples are quite revealing.

Example 1 (Estimation of the mean). Finding the sample mean is equivalent to performing a least-squares regression in which the only regressor is $x_i = 1$ for all $i = 1, \dots, N$. In this case, it is easy to see that $\mathbf{X}_g^\top \mathbf{X}_g = N_g$ and $\mathbf{X}^\top \mathbf{X} = N$. Therefore,

$$L_g = \text{Tr}(\mathbf{H}_g) = \frac{N_g}{N} = \frac{N_g}{\sum_{h=1}^G N_h}. \quad (13)$$

In this simple case, cluster leverage is exactly proportional to cluster size. In other cases, we can interpret leverage as a generalization of cluster size that takes into account other types of heterogeneity as well.

Evidently, $\hat{\beta} = \bar{y} = N^{-1} \sum_{g=1}^G N_g \bar{y}_g$, where \bar{y} and \bar{y}_g denote the sample average for the full sample and for cluster g , respectively. This expression can be rewritten as

$$\hat{\beta} = \sum_{g=1}^G \frac{N_g}{N} \bar{y}_g = \sum_{g=1}^G L_g \hat{\beta}_g, \quad (14)$$

so that $\hat{\beta}$ is seen to be a weighted average of the G estimates $\hat{\beta}_g = \bar{y}_g$, with the weight for each cluster equal to its leverage. Similarly, we find that

$$\hat{\beta}^{(g)} = \frac{N}{N - N_g} \sum_{h \neq g} L_h \hat{\beta}_h, \quad (15)$$

where the first factor simply makes up for the fact that we are summing over $G - 1$ clusters instead of G as in (14). Subtracting (14) from (15), we conclude that

$$\hat{\beta}^{(g)} - \hat{\beta} = \frac{N_g}{N} (\hat{\beta}^{(g)} - \hat{\beta}_g) = L_g (\hat{\beta}^{(g)} - \hat{\beta}_g). \quad (16)$$

Therefore, cluster g will be influential whenever omitting it yields an estimate $\hat{\beta}^{(g)}$ that differs substantially from the estimate $\hat{\beta}_g$ for cluster g itself, especially when cluster g also has high leverage. \square

Example 2 (Single regressor plus constant). Consider a regression design with a single regressor, x_i , and a constant term. Then

$$\mathbf{X}_g^\top \mathbf{X}_g = \begin{bmatrix} N_g & \sum_{i=1}^{N_g} x_{g,i} \\ \sum_{i=1}^{N_g} x_{g,i} & \sum_{i=1}^{N_g} x_{g,i}^2 \end{bmatrix} \quad \text{and} \quad (\mathbf{X}^\top \mathbf{X})^{-1} = \frac{1}{N^2 \hat{\sigma}_x^2} \begin{bmatrix} \sum_{i=1}^N x_i^2 & -\sum_{i=1}^N x_i \\ -\sum_{i=1}^N x_i & N \end{bmatrix},$$

where $\hat{\sigma}_x^2$ denotes the sample variance of the x_i . After some algebra, we find that

$$L_g = \frac{N_g}{N\hat{\sigma}_x^2} \left(\hat{\sigma}_x^2 + \hat{\sigma}_{x,g}^2 + (\bar{x}_g - \bar{x})^2 \right), \quad (17)$$

where \bar{x}_g and $\hat{\sigma}_{x,g}^2$ denote the sample mean and sample variance of the x_i within cluster g . Expression (17) is a straightforward generalization of (13). The last two terms within the large parentheses are the sample variance of the $x_{g,i}$ within cluster g and the square of the difference between \bar{x}_g and \bar{x} . The sum of these terms is the sample variance of the $x_{g,i}$ around \bar{x} within cluster g . Thus cluster g will have high leverage when the variance of the $x_{g,i}$ around \bar{x} within that cluster is large relative to the variance $\hat{\sigma}_x^2$ for the full sample. If everything except cluster sizes were perfectly balanced, L_g would evidently reduce to $2N_g/N$.

The partial leverage for x is just

$$L_{g2} = \frac{N_g \left(\hat{\sigma}_{x,g}^2 + (\bar{x}_g - \bar{x})^2 \right)}{N\hat{\sigma}_x^2}, \quad (18)$$

the total variation around \bar{x} within cluster g divided by the total variation within the sample. If everything except cluster sizes were perfectly balanced, it would reduce to N_g/N . \square

Example 3 (Single regressor plus fixed effects). Suppose there is a single regressor, x_i , and there are cluster-level fixed effects, which have been partialled out. In this case, we can write all quantities as deviations from their cluster averages, and there is no distinction between leverage and partial leverage. Then $\tilde{\mathbf{X}}_g^\top \tilde{\mathbf{X}}_g = \sum_{i=1}^{N_g} (x_{g,i} - \bar{x}_g)^2 = N_g \hat{\sigma}_{x,g}^2$. Similarly, $\tilde{\mathbf{X}}^\top \tilde{\mathbf{X}} = \sum_{g=1}^G N_g \hat{\sigma}_{x,g}^2$ is the average variance of the x_i across all clusters. We find that

$$L_g = \frac{N_g \hat{\sigma}_{x,g}^2}{\sum_{h=1}^G N_h \hat{\sigma}_{x,h}^2}, \quad (19)$$

which is again a straightforward generalization of (13). The leverage of cluster g is proportional to N_g times the variance of the $x_{g,i}$ around \bar{x}_g . Thus, for example, doubling the variance of the $x_{g,i}$ has the same effect on leverage as doubling N_g .

In this case, using (19), it is easy to see that

$$\hat{\beta} = \frac{\sum_{g=1}^G N_g \hat{\sigma}_{xy,g}}{\sum_{g=1}^G N_g \hat{\sigma}_{x,g}^2} = \sum_{g=1}^G L_g \frac{\hat{\sigma}_{xy,g}}{\hat{\sigma}_{x,g}^2} = \sum_{g=1}^G L_g \hat{\beta}_g, \quad (20)$$

where $\hat{\sigma}_{xy,g} = (1/N_g) \sum_{i=1}^{N_g} (x_{g,i} - \bar{x}_g)(y_{g,i} - \bar{y}_g)$ is the sample covariance of x_i and y_i within cluster g . The rightmost expressions in (14) and (20) are identical. In both cases, $\hat{\beta}$ is seen to be a weighted average of the G cluster estimates, with the weight for each cluster equal to its leverage.

When cluster g is omitted, we obtain

$$\hat{\beta}^{(g)} = \frac{\sum_{h \neq g} N_h \hat{\sigma}_{xy,h}}{\sum_{h \neq g} N_h \hat{\sigma}_{x,h}^2} = \frac{\sum_{h \neq g} L_h \hat{\beta}_h}{\sum_{h \neq g} L_h}, \quad (21)$$

which would specialize to (15) if (13) were true. As before, $\hat{\beta}^{(g)}$ is a weighted average of the $\hat{\beta}_h$, with weights proportional to the L_g , which in this case are also the partial leverages. Subtracting (20) from (21), we find that

$$\hat{\beta}^{(g)} - \hat{\beta} = L_g (\hat{\beta}^{(g)} - \hat{\beta}_g), \quad (22)$$

which is formally identical to the rightmost expression in (16), although of course L_g is defined in (19) not (13). Cluster g will be influential whenever $\hat{\beta}^{(g)}$ differs substantially from the estimate $\hat{\beta}_g$ for cluster g itself, especially when cluster g also has high leverage. \square

Example 4 (Treatment model with a constant term). Now we specialize Example 2 to the case in which the single regressor is a treatment dummy denoted by d_i . Let \bar{d}_g and \bar{d} denote the proportion of treated observations in cluster g and in the sample, respectively. Then (17) becomes

$$L_g = \frac{N_g}{N} \left(\frac{\bar{d}_g}{\bar{d}} + \frac{1 - \bar{d}_g}{1 - \bar{d}} \right). \quad (23)$$

The first factor here is the relative size of the g^{th} cluster. The second factor depends on how much \bar{d}_g differs from \bar{d} . When $\bar{d}_g = \bar{d}$, we see that $L_g = 2N_g/N$. Otherwise, the first term inside the parentheses causes leverage to be high whenever \bar{d}_g is large relative to \bar{d} , and the second term causes leverage to be high whenever \bar{d}_g is small relative to \bar{d} . As \bar{d} increases for given \bar{d}_g , the first term becomes smaller relative to the second term. Thus the g^{th} cluster will tend to be influential either when it has a large proportion of treated observations and the overall proportion is small, or when it has a small proportion of treated observations and the overall proportion is large.

We can also obtain the partial leverage of the treatment dummy for this case. Expression (18) simply becomes

$$L_{g2} = \frac{N_g}{N} \left(\frac{\bar{d}_g}{\bar{d}} + \frac{\bar{d} - \bar{d}_g}{1 - \bar{d}} \right). \quad (24)$$

Once again, the first factor is the relative size of the g^{th} cluster. The second factor reduces to 1 when $\bar{d}_g = \bar{d}$, so that $L_{g2} = N_g/N$ in that special case.

We can further specialize (23) and (24) to models in which the treatment is applied at the cluster level. Suppose that all observations in clusters $g = 1, \dots, G_1$ are treated and no observations in the $G_0 = G - G_1$ control clusters from $G_1 + 1$ to G are treated. Then we find that $\bar{d}_g = 1$ for $g = 1, \dots, G_1$, and $\bar{d}_g = 0$ for $g = G_1 + 1, \dots, G$. Inserting these into

(23) shows that

$$L_g = \begin{cases} \frac{N_g}{N} \frac{1}{\bar{d}} & \text{for } g = 1, \dots, G_1, \\ \frac{N_g}{N} \frac{1}{(1-\bar{d})} & \text{for } g = G_1 + 1, \dots, G. \end{cases} \quad (25)$$

Inserting them into (24) shows that

$$L_{g^2} = \begin{cases} \frac{N_g}{N} \frac{\bar{d}+1}{\bar{d}} & \text{for } g = 1, \dots, G_1, \\ \frac{N_g}{N} \frac{\bar{d}}{(1-\bar{d})} & \text{for } g = G_1 + 1, \dots, G. \end{cases} \quad (26)$$

Thus any cluster tends to have high leverage if N_g/N is large. A treated cluster has high leverage and partial leverage if \bar{d} is small. Conversely, a control cluster has high leverage and partial leverage if \bar{d} is large. \square

Example 5 (Treatment with fixed effects). Finally, we consider the case of cluster-level fixed effects, where treatment is randomly applied at the individual level. This is a special case of [Example 3](#). We cannot consider cluster fixed effects with cluster-level treatment, because the treatment dummy would be invariant within clusters. We specialize (19) and find that

$$L_g = \frac{N_g \bar{d}_g (1 - \bar{d}_g)}{\sum_{h=1}^G N_h \bar{d}_h (1 - \bar{d}_h)}. \quad (27)$$

Thus, as before, the leverage of cluster g , relative to the average for the other clusters, is proportional to its size, N_g . It also depends on the proportion of treated observations in the cluster. The maximum (relative) leverage for cluster g occurs at $\bar{d}_g = 1/2$ and is symmetric around $1/2$. The result (22) continues to hold. It tells us that cluster g will be influential when its leverage (27) is large and $\hat{\beta}^{(g)}$ differs greatly from $\hat{\beta}_g$. \square

5 Two Jackknife Variance Matrix Estimators

Although the CV_1 variance estimator defined in (5) is very widely used, it does not have particularly good finite-sample properties. That is one reason why bootstrap methods are so commonly employed. Two alternative CRVEs, which are usually known as CV_2 and CV_3 , were proposed in [Bell and McCaffrey \(2002\)](#). They are the cluster analogs of the heteroskedasticity-consistent estimators HC_2 and HC_3 , which are appropriate when the u_i are independent. These names were coined in [MacKinnon and White \(1985\)](#), which proposed HC_3 as a jackknife variance estimator. In the remainder of this section, we focus on CV_3 , because CV_2 is not a jackknife estimator and is not amenable to the computational methods that we propose; for more on it, see [Imbens and Kolesár \(2016\)](#), [Pustejovsky and Tipton \(2018\)](#), and [Niccodemi, Alessie, Angelini, Mierau and Wansbeek \(2020\)](#).

The usual expression for CV_3 is

$$\text{CV}_3: \quad \frac{G-1}{G}(\mathbf{X}^\top \mathbf{X})^{-1} \left(\sum_{g=1}^G \ddot{\mathbf{s}}_g \ddot{\mathbf{s}}_g^\top \right) (\mathbf{X}^\top \mathbf{X})^{-1}, \quad (28)$$

where the modified score vectors $\ddot{\mathbf{s}}_g$ are defined as

$$\ddot{\mathbf{s}}_g = \mathbf{X}_g^\top \mathbf{M}_{gg}^{-1} \hat{\mathbf{u}}_g. \quad (29)$$

Here $\mathbf{M}_{gg} = \mathbf{I}_{N_g} - \mathbf{X}_g(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}_g^\top$ is the diagonal block corresponding to the g^{th} cluster of the projection matrix $\mathbf{M}_\mathbf{X}$, which satisfies $\hat{\mathbf{u}} = \mathbf{M}_\mathbf{X} \mathbf{u}$.

Although computing CV_3 using (28) works well when all the N_g are very small, it becomes expensive, or perhaps computationally infeasible, when even one of the N_g is sufficiently large. The problem is that an $N_g \times N_g$ matrix needs to be stored and inverted for every cluster. A method that is much faster for large clusters has recently been proposed in [Niccodemi et al. \(2020\)](#). Versions of this procedure apply to both CV_2 and CV_3 . However, recognizing that CV_3 is a jackknife estimator allows us to develop a much simpler and even faster method for it.

The most obvious cluster jackknife estimator of $\text{Var}(\hat{\boldsymbol{\beta}})$ is

$$\text{CV}_{3J}: \quad \frac{G-1}{G} \sum_{g=1}^G (\hat{\boldsymbol{\beta}}^{(g)} - \bar{\boldsymbol{\beta}})(\hat{\boldsymbol{\beta}}^{(g)} - \bar{\boldsymbol{\beta}})^\top, \quad (30)$$

where $\bar{\boldsymbol{\beta}}$ is the sample mean of the $\hat{\boldsymbol{\beta}}^{(g)}$, which were defined in (12). The expression in (30) is the cluster analog of the usual jackknife variance matrix estimator given in [MacKinnon and White \(1985, eqn. \(11\)\)](#). Each of the $\hat{\boldsymbol{\beta}}^{(g)}$ is obtained by deleting a cluster instead of an observation, and the summation is over clusters instead of observations.

If $\bar{\boldsymbol{\beta}}$ in (30) is replaced by $\hat{\boldsymbol{\beta}}$, we obtain the estimator CV_3 :

$$\text{CV}_3: \quad \frac{G-1}{G} \sum_{g=1}^G (\hat{\boldsymbol{\beta}}^{(g)} - \hat{\boldsymbol{\beta}})(\hat{\boldsymbol{\beta}}^{(g)} - \hat{\boldsymbol{\beta}})^\top. \quad (31)$$

This version of CV_3 is numerically identical to the one in (28); see [MacKinnon et al. \(2022b\)](#) for a proof. Unless all the clusters are very small, computing CV_3 using (31) is far faster than using (28); some timings are reported in [MacKinnon et al. \(2022b\)](#).

Many discussions of jackknife variance estimation follow [Efron \(1979\)](#) and use $\bar{\boldsymbol{\beta}}$ as in (30), but others, including [Bell and McCaffrey \(2002\)](#), use $\hat{\boldsymbol{\beta}}$ as in (31). Although these two jackknife variance estimators are asymptotically the same, they are rarely equal. For example, in the case of (20), they will coincide only when every cluster has the same (partial) leverage. Thus CV_3 and CV_{3J} will normally differ, with CV_3 always larger than CV_{3J} in

the sense that the difference between them is a positive semi-definite matrix. In practice, however, they tend to be very similar (MacKinnon et al. 2022b), and there seems to be no good reason to expect either CV_3 or CV_{3J} to perform better in general.

The factor of $(G - 1)/G$ in both (30) and (31) is designed to compensate for the tendency of the $\hat{\beta}^{(g)}$ to be too spread out. This factor is the analog of the usual factor of $(N - 1)/N$ for a jackknife variance matrix at the individual level. It implicitly assumes that all clusters are the same size and perfectly balanced, with disturbances that are independent and homoskedastic. In this special case, the estimators CV_3 and CV_{3J} would be identical and unbiased (Bell and McCaffrey 2002).

Interestingly, the original HC_3 estimator proposed in MacKinnon and White (1985) is actually the analog of CV_{3J} . The modern version of HC_3 , which is the analog of CV_3 , seems to be due to Davidson and MacKinnon (1993, Chapter 16). This version of HC_3 is normally computed by dividing each residual by the corresponding diagonal element of \mathbf{M}_X , and the factor of $(N - 1)/N$ is usually (but incorrectly) omitted.

Both jackknife estimators may readily be used to compute cluster-robust t -statistics. In view of the fact that there are G terms in the summation, it seems natural to compare these with the $t(G - 1)$ distribution, as usual. These procedures should almost always be more conservative than t -tests based on the widely-used CV_1 estimator. The results in Section 7 suggest that they can be much more conservative.

When a model includes fixed effects, some care needs to be taken when computing CV_3 and CV_{3J} . As noted in Section 3, it is computationally attractive to partial out fixed effects prior to calculating $\hat{\beta}$. However, if we were to partial out any arbitrary regressors prior to computing the delete-one-cluster estimates in (12), then the computed $\hat{\beta}^{(g)}$ would depend on the values of the partialled-out regressors for the full sample, including those in the g^{th} cluster, which would be incorrect. Consequently, the values of CV_3 and CV_{3J} will be incorrect if we partial out any regressor that affects more than one cluster (such as industry-level fixed effects with firm-level clustering). An important exception is when the regressors that are partialled out are cluster fixed effects or fixed effects at a finer level (such as firm-level fixed effects with industry-level clustering), because each of them affects only one cluster. See the discussion of the `absorb` and `fevar` options in Section 8.

It is possible that the vector β is identified for the full sample but not when one cluster is deleted. For example, consider the coefficient on a dummy variable that takes on non-zero values only for observations in the g^{th} cluster. This coefficient cannot be identified when cluster g is omitted. In such a case, the matrix $\mathbf{X}^\top \mathbf{X} - \mathbf{X}_g^\top \mathbf{X}_g$ in (12) is singular, and CV_3 and CV_{3J} should not be computable. However, because `sumclust` uses the `invSYM` function in `Stata`, and hence a generalized inverse, the offending element of $\hat{\beta}^{(g)}$ is simply replaced

by 0. The package therefore checks whether any of the $\beta^{(g)}$ coefficients of interest are equal to 0 and issues a warning when they are; see [Section 8](#).

[MacKinnon et al. \(2022b\)](#) studies how to modify the WCR bootstrap ([Section 2](#)) to use CV_3 instead of CV_1 and how to modify the bootstrap DGP in a similar way. This yields several new bootstrap procedures. Limited simulation results suggest that, because using CV_3 with the $t(G - 1)$ distribution often works so well, the gains from using the new procedures for inference based on CV_3 tend to be smaller than the gains from using the classic WCR bootstrap for inference based on CV_1 .

6 What Should Be Reported

When the number of clusters is small, it is easy enough to look at all the N_g , $\beta_j^{(g)}$, L_g , and L_{gj} to see whether any clusters are unusually large, unusually influential, or have unusually high leverage or partial leverage. Once G exceeds 10 or 12, however, it may be more informative to plot these quantities. This can be done in many different ways, including histograms and box plots. Another possibility is to report a number of summary statistics. The `summclust` package, to be discussed in [Section 8](#), always reports the minimum, first quartile, median, mean, third quartile, and maximum of the N_g and the L_g . It also reports these quantities for the L_{gj} and the $\beta_j^{(g)}$ for the specified regressor j .

In the remainder of this section, we focus on cluster sizes and leverage. We consider a generic quantity a_g , which might denote any of N_g , L_g , or L_{gj} for $g = 1, \dots, G$. It seems plausible that inference may be unreliable when any of the a_g vary substantially across clusters, and we provide some evidence to support this conjecture in [Section 7](#).

It may be attractive to report a single measure of how much the distribution of the a_g differs from what it would be in the perfectly balanced case. There are many such measures. One possibility is the scaled variance

$$V_s(a_\bullet) = \frac{1}{(G - 1)\bar{a}^2} \sum_{g=1}^G (a_g - \bar{a})^2, \quad (32)$$

where the argument a_\bullet is to be interpreted as the entire set of a_g for $g = 1, \dots, G$, and \bar{a} denotes the arithmetic mean, which is N/G for the N_g , k/G for the L_g , and $1/G$ for the L_{gj} . These are all positive numbers, so it is reasonable to scale by their squares. Larger values of V_s imply that the a_g are more variable across clusters, relative to their mean. We could report either V_s or its square root, which is often called the coefficient of variation. In the perfectly balanced case, $V_s = 0$. By default, `summclust` reports the coefficient of variation for the cluster sizes, the leverages, the partial leverages, and the $\hat{\beta}_j^{(g)}$.

Another possibility is to report one or more alternative sample means. The more these differ from the arithmetic mean, the more heterogeneous must be the clusters. Three common alternatives to the arithmetic mean are the harmonic, geometric, and quadratic means:

$$\bar{a}_{\text{harm}} = \left(\frac{1}{G} \sum_{g=1}^G 1/a_g \right)^{-1}, \quad \bar{a}_{\text{geo}} = \left(\prod_{g=1}^G a_g \right)^{1/G}, \quad \text{and} \quad \bar{a}_{\text{quad}} = \left(\frac{1}{G} \sum_{g=1}^G a_g^2 \right)^{1/2}. \quad (33)$$

Unless all the a_g are the same, the harmonic and geometric means will be less than the arithmetic mean \bar{a} , and the quadratic mean (which has the same form as the root mean squared error of an estimator) will be greater than \bar{a} . We could either report the three alternative means directly, or the ratio of each of them to \bar{a} . The three ratios provide scale-free measures of cluster heterogeneity; the closer they are to one, the more homogeneous are the clusters. `summcust` can, optionally, report all of these quantities.

Another way to quantify the heterogeneity of the cluster sizes and the regressors is to calculate G^* , the “effective number of clusters,” as proposed in [Carter, Schnepel and Steigerwald \(2017\)](#). The value of G^* depends on the coefficient j for which it is being computed and on a parameter ρ to be discussed below, so we denote it $G_j^*(\rho)$. It is defined as

$$G_j^*(\rho) = \frac{G}{1 + \Gamma_j(\rho)}, \quad \Gamma_j(\rho) = \frac{1}{G} \sum_{g=1}^G \left(\frac{\gamma_{gj}(\rho) - \bar{\gamma}_j(\rho)}{\bar{\gamma}_j(\rho)} \right)^2, \quad \bar{\gamma}_j(\rho) = \frac{1}{G} \sum_{g=1}^G \gamma_{gj}(\rho), \quad (34)$$

where $0 \leq \rho \leq 1$, and the $\gamma_{gj}(\rho)$ are given by

$$\gamma_{gj}(\rho) = \mathbf{e}_j^\top (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}_g^\top \boldsymbol{\Omega}_g(\rho) \mathbf{X}_g (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{e}_j, \quad g = 1, \dots, G. \quad (35)$$

Here \mathbf{e}_j is a k -vector with 1 in the j^{th} position and 0 everywhere else, so that $\mathbf{e}_j^\top (\mathbf{X}^\top \mathbf{X})^{-1}$ is the j^{th} row of $(\mathbf{X}^\top \mathbf{X})^{-1}$, and $\boldsymbol{\Omega}_g(\rho)$ is an $N_g \times N_g$ matrix with 1 on the principal diagonal and ρ everywhere else. It is easy to see that

$$\boldsymbol{\Omega}_g(\rho) = \rho \boldsymbol{\iota} \boldsymbol{\iota}^\top + (1 - \rho) \mathbf{I}, \quad (36)$$

where $\boldsymbol{\iota}$ is an N_g -vector of 1s, and \mathbf{I} is an $N_g \times N_g$ identity matrix. Notice that $\Gamma_j(\rho)$ is just the scaled variance of the $\gamma_{gj}(\rho)$; compare [\(32\)](#).

The parameter ρ may be interpreted as the intra-cluster correlation coefficient for a model with random effects at the cluster level. Since ρ is unknown, [Carter et al. \(2017\)](#) suggests calculating $G_j^*(1)$ as a sort of worst case. However, when there are fixed effects at the cluster level, or fixed effects at a finer level nested within clusters, those fixed effects will absorb all of the intra-cluster correlation. This can lead to numerical instabilities, as we discuss below. Thus it does not make sense to specify $\rho > 0$ in either of these cases. It does seem natural to

use $G_j^*(0)$, however, because the amount of intra-cluster correlation that remains in models with cluster fixed effects is often quite small.

From (35) and (36), we see that

$$\mathbf{X}_g^\top \boldsymbol{\Omega}_g(\rho) \mathbf{X}_g = \rho(\boldsymbol{\iota}^\top \mathbf{X}_g)^\top (\boldsymbol{\iota}^\top \mathbf{X}_g) + (1 - \rho) \mathbf{X}_g^\top \mathbf{X}_g. \quad (37)$$

This result makes it inexpensive to compute the $\gamma_{gj}(\rho)$ for any value of ρ by first computing them for $\rho = 0$ and $\rho = 1$. The needed equations are

$$\begin{aligned} \gamma_{gj}(0) &= \mathbf{w}_j^\top \mathbf{X}_g^\top \mathbf{X}_g \mathbf{w}_j, \\ \gamma_{gj}(1) &= (\boldsymbol{\iota}^\top \mathbf{X}_g \mathbf{w}_j)^\top (\boldsymbol{\iota}^\top \mathbf{X}_g \mathbf{w}_j), \text{ and} \\ \gamma_{gj}(\rho) &= \rho \gamma_{gj}(1) + (1 - \rho) \gamma_{gj}(0), \end{aligned} \quad (38)$$

where \mathbf{w}_j is the j^{th} column of $(\mathbf{X}^\top \mathbf{X})^{-1}$. After obtaining the $\gamma_{gj}(\rho)$ from (38), it is trivial to compute $G_j^*(\rho)$ using (34). Evidently, $G_j^*(\rho)$ is always less than G . When it is much smaller than G , it can provide a useful warning.

If we attempt to compute $G_j^*(\rho)$ for $\rho > 0$ when there are cluster fixed effects, numerical instabilities are likely to arise. Suppose that we have partialled out cluster fixed effects prior to computing $G_j^*(\rho)$. Then the first term on the right-hand side of (37) should in theory be a zero matrix, because every column of \mathbf{X}_g should add to zero. In practice, however, the limitations of floating-point arithmetic mean that this matrix will actually contain extremely small positive numbers. When the fixed effects are not partialled out, similar but more complicated numerical issues arise.

The Stata package `clusteff` discussed in Lee and Steigerwald (2018) is designed to calculate $G_j^*(\rho)$, with $\rho = 0.9999$ rather than $\rho = 1$ by default to avoid numerical instabilities. However, the only version of this package that we have used does so in a very computationally inefficient way. It apparently computes the $\gamma_g(\rho)$ by brute force instead of using (38). When any of the N_g is large, this can take a very long time, or even fail because Stata runs out of memory. For example, it failed with the samples used in MacKinnon et al. (2022a), where the largest values of N_g were either 35,995 or 144,914. To circumvent this problem, one could use a random sample of the data for each cluster to compute an approximation to $G_j^*(\rho)$, but with (38) that is no longer necessary.

Like $V_s(a_\bullet)$ and the alternative sample means for measures of leverage and partial leverage discussed above, $G_j^*(\rho)$ is sensitive not only to variation in cluster sizes but also to other features of the \mathbf{X}_g matrices. However, it is not sensitive to heteroskedasticity or to any other features of the disturbances. Our `sumclust` package computes $G_j^*(0)$, $G_j^*(1)$, and (optionally) $G_j^*(\rho)$ for a specified covariate. However, when there are cluster fixed effects,

or fixed effects nested within clusters, it will only compute $G_j^*(0)$. For example, it will not compute $G_j^*(\rho)$ for $\rho \neq 0$ whenever there are state-level fixed effects and clustering at the region level.

The quantity $G_j^*(0)$ is very closely related to $V_s(L_{\bullet j})$, where $L_{\bullet j}$ denotes the entire set of L_{gj} , for $g = 1, \dots, G$. It is not hard to see that the $\gamma_g(0)$ defined in (35) and (38) are equal to the L_{gj} defined in (10) divided by $\hat{\mathbf{x}}_j^\top \hat{\mathbf{x}}_j$. Since this makes the $\gamma_g(0)$ proportional to the L_{gj} , $V_s(L_{\bullet j})$ is numerically identical to $\Gamma(0)$; compare (32) and the middle equation in (34). Thus we see from the first equation in (34) that $G_j^*(0)$ is simply a monotonically decreasing function of the scaled variance of our measures of partial leverage at the cluster level. When $V_s(L_{\bullet j})$ is large, $G_j^*(0)$ is necessarily much smaller than G .

Up to this point, we have focused on the case of one-way clustering. However, it is also of interest to compute measures of leverage, partial leverage, and influence when there is multi-way clustering (Cameron, Gelbach and Miller 2011). When there is two-way clustering, which is the simplest and most commonly-encountered case, we recommend computing the usual one-way measures of leverage, partial leverage, and influence for each of the two clustering dimensions. This requires calling `summclust` twice. When the number of clusters in either dimension is small, or any of these measures suggests that the data are seriously unbalanced in either dimension, then conventional inference can be seriously unreliable. The WCR bootstrap, applied to the dimension with the fewest clusters or the most unbalanced clusters (MacKinnon, Nielsen and Webb 2021), is likely to provide less unreliable inferences, but they may still be problematic.

It would probably also be interesting to calculate the standard measures of leverage, partial leverage, and influence for the intersection of the two clustering dimensions. This means calling `summclust` a third time. Suppose there are two clustering dimensions, with G clusters in the first dimension and H clusters in the second. Then the number of intersection clusters is at most GH , but it can sometimes be smaller if some of the intersection clusters are empty. In order to use `summclust` for the intersections, it is necessary to create a new variable that uniquely identifies each of the intersection clusters.

It is important to remember that, when `summclust` is invoked three times for each of two clustering dimensions and their intersection, the CV_3 standard error that it reports for each of the three cases is based on a different pattern of one-way clustering. When two-way clustering is appropriate, none of these standard errors are valid.

7 Simulation Experiments

In order to investigate the predictive power of various summary measures of leverage and partial leverage, we perform a few simulation experiments. The objective is to see whether the rejection frequencies for cluster-robust t -tests, both asymptotic and WCR bootstrap, can be predicted from summaries of the features of the \mathbf{X} matrix. There are 3000 cases, each corresponding to a particular \mathbf{X} matrix. For each case, we generate 10,000 values of \mathbf{y} and use them to estimate rejection frequencies for tests at the .05 level.

In the experiments, there are either 20 clusters and 2000 observations or 30 clusters and 3000 observations. The cluster sizes N_g are determined by a parameter $\gamma \geq 0$, as follows:

$$N_g = \left\lfloor N \frac{\exp(\gamma g/G)}{\sum_{j=1}^G \exp(\gamma j/G)} \right\rfloor, \quad g = 1, \dots, G-1, \quad (39)$$

where $\lfloor \cdot \rfloor$ denotes the integer part of its argument, and $N_G = N - \sum_{j=1}^{G-1} N_g$. As γ increases, the cluster sizes become increasingly unbalanced. The value of γ is varied uniformly between 2 and 4, so that the cluster sizes tend to vary quite a lot. When $G = 20$, the smallest cluster has between 8 and 32 observations, and the largest has between 229 and 378. When $G = 30$, the smallest cluster has between 7 and 32 observations, and the largest has between 237 and 396.


There are five regressors, one of which is the test regressor, plus a constant term. The regressors are either 0 or 1. With probability $1 - p_c$, all the observations in a cluster are 0. With probability p_c , they randomly equal either 0 or 1, both with probability 0.5. Thus, when $p_c = 1$, all variation is at the individual level, and leverage tends to be proportional to cluster sizes. As p_c declines, the samples become more unbalanced. In the experiments, the values of p_c are chosen to be 0.25, 0.30, 0.35, 0.40, 0.50, and 0.60, each for one-sixth of the cases. Smaller values of p_c tend to be associated with larger discrepancies between actual rejection frequencies and .05, the nominal level of the tests.

For each experiment, we obtain 9000 estimated rejection frequencies. One-third of these are based on CV_1 and the $t(G-1)$ distribution, one-third are based on CV_3 and the $t(G-1)$ distribution, and one-third are based on the WCR bootstrap with 399 bootstrap samples. To predict these rejection frequencies, we use a generalized additive model based on smoothing splines; see [James et al. \(2021, Section 7.7\)](#). The base model can be written as

$$r_i = \beta_0 + f_1(V_{si}) + f_2(V_{si}^{1/2}) + \beta_1 G_{i0}^* + u_i, \quad (40)$$

where r_i is the rejection frequency for case i , V_{si} denotes $V_s(L_{\bullet j})$, the scaled variance of the partial leverages L_{gj} for the test regressor for case i , G_{i0}^* denotes $G_j^*(0)$ for the test regressor for case i (recall that it is a function of the L_{gj}), and $f_1(\cdot)$ and $f_2(\cdot)$ are smoothing splines

Figure 1: Predicted rejection frequencies for asymptotic and bootstrap tests at .05 level



Notes: Each of the curves shows fitted values from the generalized additive model (40) that predicts observed rejection frequencies, based on 10,000 replications, using nonlinear functions of the $V_s(L_{\bullet j})$; see the text for details. Bootstrap rejection frequencies are based on $B = 399$. For the asymptotic tests based on CV₁, the R^2 is 0.9487 for $G = 20$ and 0.9404 for $G = 30$. For the tests based on CV₃, the R^2 is 0.8709 for $G = 20$ and 0.8811 for $G = 30$. For the WCR bootstrap tests, the R^2 is 0.4593 for $G = 20$ and 0.3400 for $G = 30$.

with five degrees of freedom. Since everything on the right-hand side of (40) is a function of V_{si} , this model is simply using the V_{si} to predict the r_i in a potentially nonlinear way.

Figure 1 shows the fitted values from (40), which are predicted rejection frequencies, plotted against the scaled variance of the partial leverages L_{gj} for three methods of inference and two sample sizes. Panel (a) shows them for t -tests based on both CV₁ (solid lines) and CV₃ (dashed lines) for $G = 20$ and $G = 30$, and Panel (b) shows them for WCR bootstrap tests for the same two cases.

The t -tests based on CV₁ often over-reject to an extreme degree. For the very smallest values of $V_s(L_{\bullet j})$, the tests tend to over-reject modestly, with predicted rejection frequencies of 0.058 for $G = 20$ and 0.055 for $G = 30$. However, these then rise quite rapidly and almost linearly. In fact, for $G = 30$, (40) omits the f_2 term because it has very little explanatory power. Nevertheless, the analysis of variance for nonparametric effects strongly suggests that there actually is nonlinearity for both $G = 20$ and $G = 30$. For $G = 30$, there are four cases (out of 3000) for which $V_s(L_{\bullet j}) > 15$. These are not shown in the figure, but the approximately linear relationship continues to hold, and the fit for these extreme cases is reasonably good.

In contrast, the t -tests based on CV₃ tend to under-reject for small values of $V_s(L_{\bullet j})$. For the very smallest values, the predicted rejection frequencies are 0.033 for $G = 20$ and 0.039 for $G = 30$. Although it is not obvious from the figure, the CV₃ tests are predicted to

under-reject somewhat more than half the time, because, in our experiments, most values of $V_s(L_{\bullet j})$ are quite small. As $V_s(L_{\bullet j})$ increases, rejection frequencies increase, although for $G = 20$ they start to decline again once $V_s(L_{\bullet j})$ exceeds about 9.6. The predicted rejection frequencies never exceed 0.105 for $G = 20$ and 0.118 for $G = 30$. In a few cases (74 for $G = 20$ and 5 for $G = 30$), it was impossible to obtain meaningful CV_3 estimates because the matrix that is inverted in (12) was singular. This happened whenever one of the regressors took the same value for all observations in $G - 1$ of the clusters. These cases were dropped.

Panel (b) of Figure 1 shows the fitted values from (40) for WCR bootstrap t -tests plotted against the scaled variance of the L_{gj} . These curves are not even close to being linear. For both values of G , the predicted rejection frequencies are extremely close to .05 for the smallest values of $V_s(L_{\bullet j})$. They then rise for a while, achieving maxima of 0.061 for $G = 20$ and 0.057 for $G = 30$, and subsequently fall to very low levels. In fact, the predicted values for $G = 20$ are slightly negative for a few extreme values of $V_s(L_{\bullet j})$ that are not shown in the figure, where the actual values that are being predicted are mostly zero.

In one respect, Panel (b) of Figure 1 is a bit misleading. It suggests that the bootstrap tests under-reject for most values of $V_s(L_{\bullet j})$. They actually over-reject much more often than they under-reject, because most of the observed values of $V_s(L_{\bullet j})$ are relatively small. For $G = 20$, the crossover between over-rejection and under-rejection occurs at about 4.66, with only 380 values of $V_s(L_{\bullet j})$ (out of 3000) greater than 4.66. For $G = 30$, the crossover occurs at about 6.12, with only 139 values of $V_s(L_{\bullet j})$ greater than 6.12.

The bootstrap tests often under-reject when the number of “treated” clusters (that is, clusters that have any non-zero values of the test regressor) is small. When this happens, the asymptotic tests almost always over-reject, especially the ones based on CV_1 . Given the way we generate the regressors, this situation inevitably tends to occur more frequently, and to a greater extent, for $G = 20$ than for $G = 30$. Apparently $V_s(L_{\bullet j})$ tends to be quite large when this happens, so that (40) is able to predict both the over-rejection by the asymptotic tests and the under-rejection by the bootstrap tests.

It is natural to ask whether we can improve the fit of (40) by adding additional explanatory variables that are not simply functions of the $V_s(L_{\bullet j})$. The answer is that we can. When a large number of additional explanatory variables is added to this equation, the variables $\bar{a}_{\text{geo}}(L_{\bullet j})$ and $G_j^*(1)$ are always highly significant. In most cases, at least some functions of the L_g are also significant. However, functions of partial leverage generally seem to have much more explanatory power than functions of leverage. The spline $f_1(V_{si})$ is always highly significant, even when many other regressors are included. Thus, at least in these experiments, the scaled variance of the partial leverages, which is the square of their coefficient of variation, seems to be particularly revealing.

We make no effort to obtain the best possible predictive model, because the specification of any such model would be extremely dependent on the precise way in which the regressors in our experiments were generated. In particular, we simply add additional explanatory variables linearly, not within splines. Nevertheless, for all six experiments, we are able to reduce the residual standard errors substantially. In the best case, for $G = 30$ and CV_3 , we can reduce it by 21%. In the worst case, for $G = 20$ and CV_1 , we can reduce it by 9%. Thus it appears that no single summary statistic for either partial leverage or leverage can fully capture the effects of the \mathbf{X} matrix on the rejection frequencies of either asymptotic or bootstrap tests.

Based on these results, which are of course extremely dependent on the way in which the regressors are generated, it seems sensible for investigators to look at a number of different summary measures for both leverage and partial leverage. The `summclust` package therefore reports several of these. The most informative one seems to be the scaled variance, defined in (32), of the partial leverage measures L_{gj} , defined in (10), for the regressor of interest. However, this recommendation is necessarily very tentative.

8 The `summclust` Package

The `summclust` package calculates a large number of statistics to help assess cluster heterogeneity. It also provides CV_3 and CV_{3J} standard error estimates. The package does not rely on any other `Stata` packages, but it does require a version of `Stata` that provides Mata’s `pansum()` function (Version 13 or later). To illustrate `summclust`’s functionality and syntax, we consider a simple example using the online dataset `nlswork`, which contains a sample of women who were 14–26 years of age in 1968 from the National Longitudinal Survey of Young Working Women. Key output for this example is shown in this section. Additional output, along with a brief discussion, can be found in [Appendix A](#). All the output is displayed on screen, and most is also available in Mata and/or return memory using `r()`; see [Appendix A](#).

Suppose we are interested in estimating a simple Mincer regression using the `nlswork` dataset. We want to ask whether there is a marriage premium for wages. The variable `msp` is equal to 1 if the person is married and cohabits with their spouse, and equal to 0 otherwise. For the purposes of this example, we cluster by industry. The following code opens the dataset and estimates the regression using `Stata`’s `regress` command:

```
webuse nlswork, clear
reg ln_wage i.grade i.age i.birth_yr union race msp, cluster(ind)
```

The `Stata` output provides CV_1 standard errors. Alternatively, we can estimate the coefficient

on `msp` and two standard errors for it using `summclust`. The command is:

```
summclust msp, yvar(ln_wage) xvar(union race) \\\
fevar(grade age birth_yr) cluster(ind)
```

This code results in the default output from `summclust`, which is mostly contained in two tables. The first one includes the coefficient on the variable specified (in this case `msp`), the CV_1 and CV_3 standard errors for this coefficient, and the associated t -statistics, P values, and confidence intervals.

Cluster summary statistics for `msp` when clustered by `ind_code`.
There are 19130 observations within 12 `ind_code` clusters.

Regression Output

s.e.	Coeff	Sd. Err.	t-stat	P value	CI-lower	CI-upper
CV1	-0.027515	0.009293	-2.9608	0.0130	-0.047969	-0.007061
CV3	-0.027515	0.014064	-1.9564	0.0763	-0.058470	0.003440

Unsurprisingly, the first row matches the output from the `regress` command. In this case, the CV_1 and CV_3 standard errors are noticeably different, with the latter being considerably larger. The second table from `summclust` may provide some insight as to why this is so. It reports summary statistics for N_g , L_g , L_{gj} , and $\beta_j^{(g)}$.

Cluster Variability

Statistic	Ng	Leverage	Partial L.	beta no g
min	38.00	0.093321	0.001622	-0.033200
q1	159.00	0.672924	0.008649	-0.029275
median	995.50	3.515491	0.056682	-0.027765
mean	1594.17	5.416667	0.083333	-0.026920
q3	2335.50	7.731883	0.120933	-0.025975
max	6335.00	20.289183	0.312995	-0.015835
coefvar	1.19	1.152965	1.141326	0.162898

We see that the clusters are extremely heterogeneous. The largest cluster contains almost one-third of the sample and is 167 times the size of the smallest. There are also extreme

differences in both leverage and partial leverage across clusters. The ratio of the largest to the smallest value is 217 for leverage and 193 for partial leverage. The sum of the leverages is $12 \times 5.416667 = 65$, which is the number of estimated coefficients. Although the $\beta_j^{(g)}$, denoted “beta no g” in the table, vary quite a bit, dropping one cluster never changes the sign of the coefficient.

The syntax of `summclust` requires that the first argument be the variable corresponding to the coefficient of interest, in this case `msp`. The arguments `yvar`, `xvar`, and `cluster`, which follow the comma after `msp`, are required. The dependent variable is specified in `yvar`. The clustering variable is specified in the `cluster` argument.

The treatment of additional independent variables requires a bit more discussion. The argument `xvar` is used for listing ‘ordinary’ regressors, such as continuous and binary variables. In the example above, these are the binary indicators for union and race. Time-series operators and factor variables are not permitted within `xvar`. The optional argument `fevar` should be used for the latter. This argument includes all the listed variables as factor variables to be treated as additional independent variables. In the above example, these are `grade`, `age`, and `birth_yr`. Each argument in `fevar` will result in the creation of a set of temporary dummy variables for each value of the argument. These dummy variables are included in the regression, and there is no constant term if they are present.

The sample code above does not illustrate several additional options. The most important of these is the `absorb` option, which operates like `fevar`. It treats its argument, a single variable, as an additional factor variable to include in the set of regressors. `absorb(varname)` can be used when including `i.varname` in a regression would result in many fixed effects. Speed is increased, perhaps substantially, by partialing out the absorbed fixed effects from the dependent and all the independent variables. It is advisable to use `absorb` rather than `fevar` whenever their argument corresponds to a set of cluster fixed effects, since the elements of $\hat{\beta}^{(g)}$ that correspond to the fixed effects cannot be identified in that case; see [Section 3](#).

The `absorb` option should be used with care. Partialing out fixed effects is valid for the measures of leverage and influence and for the jackknife variance matrices only when the absorbed variable yields fixed effects that can be partialled out on a cluster-by-cluster basis. That is, `absorb` should only be used for straight cluster fixed effects or for fixed effects at a finer level, such as state \times year fixed effects for a panel with clustering at the state level. It is not valid to partial out fixed effects that are not limited to a single cluster. In that case, the $\beta^{(g)}$ and quantities based on them would be different for the original data and the data after partialing out, because the partialled-out observations for the g^{th} cluster would depend on other clusters as well. Accordingly, `summclust` checks to ensure that the clustering variable is invariant within each value of the absorb variable. When it is not, a warning is displayed,

and the estimates for L_{gj} , $\beta_j^{(g)}$, CV_3 , and CV_{3J} are not available.

To see the difference between `fevar` and `absorb`, we can estimate an expanded regression that includes industry fixed effects. The following two commands produce identical estimates for N_g , L_{gj} , $\beta_j^{(g)}$, and CV_3 . However, the estimates of the L_g are different.

```
summclust msp, yvar(ln_wage) xvar(union race) \\
fevar(grade age birth_yr ind) cluster(ind)

summclust msp, yvar(ln_wage) xvar(union race) \\
fevar(grade age birth_yr) absorb(ind) cluster(ind)
```

Including the cluster fixed effects changes the coefficient and the two standard errors. Both commands produce the following regression output:

Regression Output

s.e.	Coeff	Sd. Err.	t-stat	P value	CI-lower	CI-upper
-----+-----						
CV1	-0.020895	0.007086	-2.9487	0.0132	-0.036492	-0.005299
CV3	-0.020895	0.007931	-2.6345	0.0232	-0.038352	-0.003438

However, they produce somewhat different output for cluster variability, as measured by leverage; see [Appendix A](#).

The option `sample` allows for sample restrictions. The argument(s) for this option are whatever would follow the “if” in a standard `regress` command. For instance, in order to restrict the analysis to individuals 25 years of age or older in the `nlswork` example, `sample(age>=25)` should be added to the list of options.

The remaining options instruct `summclust` to print additional output. For the `nlswork` example, this output is shown in [Appendix A](#).

`jackknife` calculates the CV_{3J} standard error and the associated t -statistic, P value, and confidence interval. These results are displayed in an additional row in the Regression Output table.

`svars` prints the alternative means of the N_g , L_g , L_{gj} , and $\beta_j^{(g)}$, as described in [Section 6](#).

For the N_g , L_g , and L_{gj} , it reports the harmonic, geometric, and quadratic means, as well as the ratio of each of them to the arithmetic mean. For the $\beta_j^{(g)}$, which can be negative, only the quadratic mean and its ratio are reported, because the harmonic and geometric means are not defined for negative numbers.

gstar calculates the effective number of clusters $G_j^*(0)$. When there are no cluster fixed effects, or fixed effects nested within clusters, it also calculates $G_j^*(1)$.

rho calculates the effective number of clusters, $G_j^*(\rho)$, but only if it is valid to do so. It either prints $G_j^*(\rho)$ in addition to $G_j^*(0)$ and $G_j^*(1)$, or it reports that $G_j^*(\rho)$ cannot be computed and prints only $G_j^*(0)$. This option can be specified with or without also specifying **gstar**, and the output will be the same in either case. The value of ρ is specified as **rho**(ρ), where ρ must be between 0 and 1. The program stops and displays an error message when an invalid value of ρ is entered.

table prints a table of the G cluster-specific estimates of N_g , L_g , L_{gj} , and $\beta_j^{(g)}$. If $G > 52$, then the unformatted matrix is displayed instead of a table.

9 Empirical Example

We consider an empirical example from [Busso and Galiani \(2019\)](#), which studies an experiment where retail firms were randomly assigned to enter different geographic markets within the Dominican Republic. Firm entries were assigned across 72 districts, with clustering by district. After randomization, 21 districts had no entrants and so were in the control group, 18 districts had one entrant, another 18 had two, and the remaining 15 had three. The primary analysis only distinguishes between the 51 treated districts and the 21 control districts. The number of observations (stores) per district varies from 20 to 55.

This example is interesting because conventional wisdom (e.g., [MacKinnon et al. 2022a](#)) suggests that, with 72 clusters that do not vary much in size, and with neither few treated nor few control clusters, inference based on CV_1 standard errors and the $t(71)$ distribution should work well. However, our leverage measures suggest otherwise, and alternative inference methods yield noticeably different results.

The model we estimate is

$$Y_{sd} = \alpha + \gamma Z_d + \mathbf{X}_{sd}\boldsymbol{\beta} + \epsilon_{sd}. \quad (41)$$

Here s indexes stores, and d indexes districts. The treatment variable Z_d equals 1 if district d is treated (there was entry) and 0 if it was a control (there was no entry). The coefficient of interest is γ , which measures the causal effect of increased competition on an outcome Y . The original paper deals with many outcomes, but we focus on just one of them, the log of demeaned prices after treatment. The results from this regression are found in Table 5, Panel B, column 4, row 1 of [Busso and Galiani \(2019\)](#). The table states that there are 72 clusters and 2,025 observations; however, the replication dataset that we use contains just

Table 1: Estimates of the Treatment Effect

Method	$\hat{\gamma}$	Standard Error	P value	Confidence Interval
CV ₁	−0.01469	0.00724	0.0461	[−0.02913, −0.00025]
CV ₃	−0.01469	0.00909	0.1105	[−0.03281, 0.00343]
WCR bootstrap	−0.01469		0.0898	[−0.03125, 0.00245]

Notes: There are $N = 1,926$ observations and $G = 72$ clusters. The WCR bootstrap uses $B = 999,999$.

Table 2: Leverage and Partial Leverage for $\hat{\gamma}$

Statistic	N_g	Leverage	Partial Leverage	$\hat{\gamma}^{(g)}$
Minimum	20	0.130842	0.000099	−0.017550
First quartile	24	0.204104	0.003166	−0.015089
Median	26	0.235813	0.009001	−0.014791
Mean	26.75	0.263889	0.013889	−0.014663
Third quartile	27	0.292042	0.020926	−0.014070
Maximum	55	0.737797	0.064242	−0.010723
Coef. of variation	0.21	0.388686	1.059813	0.074061

Notes: There are $N = 1,926$ observations and $G = 72$ clusters. The effective numbers of clusters are $G_\gamma^*(0) = 34.16$ and $G_\gamma^*(1) = 33.33$.

1,926 observations.

Regression (41) includes 17 control variables in the row vector \mathbf{X}_{sd} . These are the first lag of the outcome variable, the number of retailers in each neighborhood pre-treatment, a lagged quality index, eight province fixed effects, total neighborhood beneficiaries of a conditional cash transfer program, percent beneficiaries of that program, average income in the district, two district education measures, and a binary indicator for the urban status of the district. Thus the total number of regressors is 19.

The OLS estimate of γ , its CV₁ standard error, the P value for a test that $\gamma = 0$, and a .95 confidence interval are shown in the first row of Table 1. Allowing for different numbers of reported digits, these estimates accord with the ones in Busso and Galiani (2019). The estimate of −0.01469 has the expected sign (average prices declined). However, the P value is just slightly less than 0.05, and the confidence interval barely excludes 0.

We next use the `summclust` package to calculate the cluster-level characteristics of the model and dataset. Some key ones are reported in Table 2. It is evident that cluster sizes are well balanced, varying from 20 to 55, with the first and third quartiles equal to 24 and 27. However, both the leverages L_g and the partial leverages L_{g1} vary considerably. The former range from 0.1308 to 0.7378, and the latter from 0.0001 to 0.0642. The coefficients of variation are 0.3887 and 1.0598, respectively. The latter is moderately large, although not

enormous. The two values of G^* are slightly smaller than $G/2$, which also suggests that the sample is not well balanced.

Most of the $\hat{\gamma}^{(g)}$ do not vary much, and thus their coefficient of variation is small. However, the most extreme values are notable. The estimate of γ , which is -0.01469 , could be as small as -0.01755 or as large as -0.01072 if just one out of 72 clusters were dropped.

These results suggest that CV_1 , the default CRVE, may not be particularly reliable in this case. We therefore consider two alternative procedures. The second row of [Table 1](#) reports the CV_3 standard error as well as the P value and confidence interval associated with it, computed using `summcust`. At 0.1105, the CV_3 P value does not even allow us to reject the null at the .10 level. The third row of [Table 1](#) reports a WCR bootstrap P value and a .95 confidence interval, both computed using `boottest` ([Roodman et al. 2019](#)). At 0.0898, the bootstrap P value also does not allow us to reject the null hypothesis at the .05 level. Of course, both these confidence intervals include 0.


In view of the reasonably large number of clusters and the fact that cluster sizes do not vary much, the large discrepancy between the results for CV_1 and the other two procedures may seem surprising. However, it is not all that surprising when we note how much the leverages and, especially, the partial leverages vary.

To explore what is driving the differences in partial leverage, we create a scatter plot. [Figure 2](#) plots partial leverage against the number of observations per cluster, with different colors and symbols depending on whether or not a given district (cluster) was treated. The figure has two interesting features. The first is that the three rather large clusters have fairly small partial leverage. The second is that the 12 clusters with the highest partial leverage are all control districts. The first result is quite surprising, since large clusters often tend to have high leverage. But [Figure 2](#) makes it clear that there is, in general, no simple relationship between cluster sizes and partial leverage. The second result is not so surprising, because only 21 out of the 72 clusters are controls. Many of the control clusters presumably have high partial leverage because control clusters are relatively rare.

10 Concluding Remarks

In modern empirical work, it is very often assumed that the observations belong to a number of disjoint clusters. No dependence across clusters is allowed, but, as we discuss in [Section 2](#), there may be arbitrary patterns of heteroskedasticity and correlation within each cluster. In such cases, the key unit of observation is the cluster. It therefore makes sense to examine measures of influence, leverage, and partial leverage at the cluster level. As we discuss in [Section 3](#), these measures are easy to compute and are conceptually very similar to classic

Figure 2: Partial Leverage vs Cluster Size



Notes: The figure plots partial leverage against cluster size for 72 clusters. A green X marks a treated cluster, and an orange circle marks a control cluster.

measures of influence, leverage, and partial leverage at the observation level. In simple cases, our measures can be obtained theoretically, and studying such cases, as we do in [Section 4](#), provides some useful insights.

Our measure of influence, which is closely related to the jackknife, can provide valuable information about how empirical results depend on the data in the various clusters. Investigators should be wary if dropping one or two clusters changes the results dramatically.

In [Section 5](#), we show how to calculate two jackknife variance matrix estimators, CV_3 and CV_{3J} , very rapidly as a byproduct of our other computations. Although these estimators are not new ([Bell and McCaffrey 2002](#)), they have until recently been difficult or impossible to compute for large samples. In related work ([MacKinnon et al. 2022b](#)), we show that CV_3 and CV_{3J} generally yield extremely similar results, and that they are often much more reliable than the widely-used CV_1 variance matrix estimator. [Figure 1](#) provides additional evidence on the latter point.

We recommend that our measures of leverage and partial leverage be used as diagnostic tools to identify datasets and regression designs in which cluster-robust inference is likely to be challenging. In [Section 6](#), we discuss several ways of summarizing these measures. In

Section 7, we show that many of these summary statistics, especially the scaled variance of the partial leverages, can help to explain how well both asymptotic and bootstrap cluster-robust inference performs when the number of clusters is modest (20 or 30 in our experiments). We argue that investigators should examine a number of these summary statistics and exercise caution whenever there seems to be a lot of variation across clusters.

In Section 9, we present an empirical example in which conventional asymptotic inference based on CV_1 seems to work poorly, even though there are 72 clusters that do not vary much in size, and neither the number of treated clusters nor the number of control clusters is particularly small. However, since our measures of partial leverage vary quite a lot, with a coefficient of variation that exceeds one, there is reason to worry about the accuracy of conventional inferences based on CV_1 and the $t(71)$ distribution. In fact, both CV_3 and the WCR bootstrap yield inferences that differ substantially from the conventional ones.

The procedures that we discuss in this paper are implemented in a new Stata package called `summclust`, which may be obtained from SSC or <https://github.com/mattdwebb/summclust>. A detailed description of the `summclust` command and its options is given in Section 8.

A Appendix

Appendix A.1 discusses `summclust`'s options and their output in more detail than Section 8 did. Appendix A.2 describes where to find the stored results in memory.

A.1 Additional Output

The `jackknife` option prints out an additional line in the Regression Output table based on the CV_{3J} standard error.

Regression Output

s.e.	Coeff	Sd. Err.	t-stat	P value	CI-lower	CI-upper
-----+-----						
CV1	-0.027515	0.009293	-2.9608	0.0130	-0.047969	-0.007061
CV3	-0.027515	0.014064	-1.9564	0.0763	-0.058470	0.003440
CV3J	-0.027515	0.013925	-1.9760	0.0738	-0.058164	0.003133

The `table` option prints out the values of the four quantities N_g , L_g , L_{gj} , and $\beta_j^{(g)}$ for all clusters. For the example in Section 8, the code is:

```
summclost msp, yvar(ln_wage) xvar(union race) \\
fevar(grade age birth_yr) cluster(ind) table
```

In addition to the two standard tables, this command prints the following table:

Cluster by Cluster Statistics

ind_code	Ng	Leverage	Partial L.	beta no g
-----+				
1	130	0.609953	0.005946	-0.027754
2	38	0.093321	0.001622	-0.027775
3	185	0.735895	0.009619	-0.027846
4	3747	13.687178	0.200637	-0.022575
5	1069	3.436310	0.060675	-0.026229
6	2912	10.372989	0.150633	-0.030672
7	1759	5.090777	0.091233	-0.025721
8	572	2.801372	0.027535	-0.029964
9	922	3.863913	0.052689	-0.033200
10	133	0.424438	0.007678	-0.028587
11	6335	20.289182	0.312995	-0.015835
12	1328	3.594673	0.078738	-0.026881

This table makes it easy to see whether the high leverage clusters are also the largest clusters. That is clearly the case here. After running the program, this table is stored as the Mata matrix `scall`.

To obtain summary statistics on the four measures of cluster variability, we can use the `svars` option:

```
summclost msp, yvar(ln_wage) xvar(union race) \\
fevar(grade age birth_yr) cluster(ind) svars
```

This command produces the following table:

Alternative Sample Means and Ratios to Arithmetic Mean

	Ng	Leverage	Partial L.	beta no g
-----+				
Harmonic Mean	227.315	0.679104	0.010629	.
Harmonic Ratio	0.143	0.125373	0.127543	.

Geometric Mean	687.061	2.390709	0.035683	.
Geometric Ratio	0.431	0.441362	0.428198	.
Quadratic Mean	2413.502	8.068016	0.123437	0.027245
Quadratic Ratio	1.514	1.489480	1.481240	-1.012089

Once again, we see that there is extreme variability across the clusters. This is particularly noticeable for the ratio of the harmonic mean to the arithmetic mean, which is between 0.125 and 0.143 for the cluster size, leverage, and partial leverage measures. Recall that these ratios would be close to one if the clusters were relatively homogeneous. This table is stored in Mata's memory as `bonus`.

To obtain estimates of the effective number of clusters, we can use either the `gstar` option or the `rho()` option. The former prints $G_j^*(0)$ and $G_j^*(1)$. The latter requires a specified value of ρ and prints $G_j^*(0)$ and $G_j^*(1)$ along with $G_j^*(\rho)$. For the `nlswork` example, the first option could be called as:

```
summclust msp, yvar(ln_wage) xvar(union race) \\\
fevar(grade age birth_yr) cluster(ind) gstar
```

This yields:

Effective Number of Clusters

```
-----
G*(0)  =  5.469
G*(1)  =  1.366
-----
```

The second option would be called (using $\rho = 0.5$ to illustrate) as:

```
summclust msp, yvar(ln_wage) xvar(union race) \\\
fevar(grade age birth_yr) cluster(ind) rho(0.5)
```

This yields:

Effective Number of Clusters

```
-----
G*(0)  =  5.469
G*(.5) =  1.421
G*(1)  =  1.366
-----
```

In this example, it is clear that the effective number of clusters is substantially less than the actual number of clusters. This provides more evidence that inference using the CV_1 standard error together with the $t(G - 1)$ distribution may be unreliable. These three quantities can be accessed in Mata's memory as `gstarzero`, `gstarrho`, and `gstarone`, respectively.

Finally, the following two examples highlight the differences between the `fevar` and `absorb` options. Up to this point, we have only used the `fevar` option. But suppose that we wish to control for industry fixed effects. Since the industry fixed effects can be partialled out on a cluster-by-cluster basis, we can control for them using the `absorb` option. We could alternatively use the `fevar` option, but it will produce different values for leverage. The first command below uses `absorb`, and the second one uses `fevar`.

```
summcclust msp, yvar(ln_wage) xvar(union race) \\\
    fevar(grade age birth_yr) absorb(ind) cluster(ind)

summcclust msp, yvar(ln_wage) xvar(union race) \\\
    fevar(grade age birth_yr ind) cluster(ind)
```

Using the `absorb` option yields:

Cluster Variability

Statistic	Ng	Leverage	Partial L.	beta no g
min	38.00	0.087112	0.001561	-0.023382
q1	159.00	0.656606	0.008621	-0.022428
median	995.50	3.442673	0.056073	-0.021258
mean	1594.17	5.333333	0.083333	-0.020770
q3	2335.50	7.605927	0.121546	-0.020189
max	6335.00	20.011074	0.312377	-0.015001
coefvar	1.19	1.155829	1.141658	0.120094

In contrast, using the `fevar` option yields:

Cluster Variability

Statistic	Ng	Leverage	Partial L.	beta no g
min	38.00	1.087112	0.001561	-0.023382
q1	159.00	1.656606	0.008621	-0.022428

median	995.50	4.442673	0.056073	-0.021258
mean	1594.17	6.333333	0.083333	-0.020770
q3	2335.50	8.605927	0.121546	-0.020189
max	6335.00	21.011074	0.312377	-0.015001
-----+-----				
coefvar	1.19	0.973329	1.141658	0.120094

Notice that the estimates for partial leverage and the $\beta_j^{(g)}$ are identical for both options. However, the leverage estimates are smaller when we use the **absorb** option. Recall that, for the original model with no industry fixed effects, the leverages summed to 65. In the first case just above, where the industry fixed effects are partialled out using **absorb**, the regression has 64 coefficients, and the leverages therefore sum to $12 \times 5.333333 = 64$. In the second case, where the industry fixed effects are included as regressors in **fevar**, the regression has 76 coefficients, and the leverages therefore sum to $12 \times 6.333333 = 76$. In fact, for the second case, each of the leverages is larger than the corresponding one for the first case by precisely 1.

A.2 List of Stored Results

All the results that are displayed as output can also be found in Mata's memory. To access one of these after running **summclust**, simply add the following line:

```
mata: object_name
```

The `object_name` can take one of the following values.

cvstuff: This matrix stores the table with the title "Regression Output". It is 2×6 when the **jackknife** option is not used (the default), and 3×6 when **jackknife** is used.

scall: This matrix stores the $G \times 4$ table created by the **table** option with the title "Cluster by Cluster Statistics".

bonus: This 6×4 matrix contains the alternative sample means and their ratios to the arithmetic mean created by the **svars** option.

gstarzero: This scalar contains $G^*(0)$ created by the **gstar** or **rho** options.

gstarone: This scalar contains $G^*(1)$ created by the **gstar** or **rho** options.

gstarrho: This scalar contains $G^*(\rho)$ created by the **rho** option.

Scalars within matrices can be referenced on a cell-by-cell basis. For example, the CV_3 standard error is stored in the second row and second column of **cvstuff**, and to display it one can enter the following command:

`mata: cvstuff[2,2]`

Additionally, several results are available as scalars or matrices in return memory using `r()`. The available scalars are:

beta: The estimate $\hat{\beta}$ for the coefficient of interest.

cv1se: The CV_1 standard error for the coefficient of interest.

cv1t: The CV_1 t -statistic for the coefficient of interest.

cv1p: The P value for the null hypothesis that $\beta = 0$ for the coefficient of interest using the CV_1 standard error.

cv1lci: The lower bound of the 95% confidence interval for β using the CV_1 standard error.

cv1uci: The upper bound of the 95% confidence interval for β using the CV_1 standard error.

gstarzero: The effective number of clusters for the coefficient of interest using $\rho = 0$.

gstarone: The effective number of clusters for the coefficient of interest using $\rho = 1$.

gstarrho: The effective number of clusters for the coefficient of interest using the value of ρ specified in `rho(ρ)`.

The standard error, t -statistic, P value, and confidence interval bounds are also available for the CV_3 and CV_{3J} standard errors. To access these, replace “1” in the above with either “3” or “3J”; for example, the P value using CV_{3J} is available in `cv3Jp`.

The available matrices are:

ng: This $G \times 1$ matrix contains the number of observations, N_g , for each cluster.

leverage: This $G \times 1$ matrix contains the leverage, L_g , for each cluster.

partlev: This $G \times 1$ matrix contains the partial leverage, L_{gj} , for each cluster.

betanog: This $G \times 1$ matrix contains the $\beta_j^{(g)}$ for each cluster.

References

- Bell, R.M., McCaffrey, D.F., 2002. Bias reduction in standard errors for linear regression with multi-stage samples. *Survey Methodology* 28, 169–181.
- Belsley, D.A., Kuh, E., Welsch, R.E., 1980. *Regression Diagnostics*. Wiley, New York.
- Bester, C.A., Conley, T.G., Hansen, C.B., 2011. Inference with dependent data using cluster covariance estimators. *Journal of Econometrics* 165, 137–151.
- Broderick, T., Giordano, R., Meager, R., 2021. An automatic finite-sample robustness metric: Can dropping a little data change conclusions? *ArXiv e-prints* 2011.14999. London School of Economics.

- Busso, M., Galiani, S., 2019. The causal effect of competition on prices and quality: Evidence from a field experiment. *American Economic Journal: Applied Economics* 11, 33–56.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics* 90, 414–427.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29, 238–249.
- Cameron, A.C., Miller, D.L., 2015. A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 50, 317–372.
- Carter, A.V., Schnepel, K.T., Steigerwald, D.G., 2017. Asymptotic behavior of a t test robust to cluster heterogeneity. *Review of Economics and Statistics* 99, 698–709.
- Chatterjee, S., Hadi, A.S., 1986. Influential observations, high-leverage points, and outliers in linear regression. *Statistical Science* 1, 379–416.
- Chesher, A., 1989. Hájek inequalities, measures of leverage and the size of heteroskedasticity robust tests. *Econometrica* 57, 971–977.
- Conley, T.G., Gonçalves, S., Hansen, C.B., 2018. Inference with dependent data in accounting and finance applications. *Journal of Accounting Research* 56, 1139–1203.
- Cook, R.D., Weisberg, S., 1980. Characterizations of an empirical influence function for detecting influential cases in regression. *Technometrics* 22, 495–508.
- Davidson, R., MacKinnon, J.G., 1993. *Estimation and Inference in Econometrics*. Oxford University Press, New York.
- Djogbenou, A.A., MacKinnon, J.G., Nielsen, M.Ø., 2019. Asymptotic theory and wild bootstrap inference with clustered errors. *Journal of Econometrics* 212, 393–412.
- Efron, B., 1979. Bootstrapping methods: Another look at the jackknife. *Annals of Statistics* 7, 1–26.
- Hansen, B.E., Lee, S., 2019. Asymptotic theory for clustered samples. *Journal of Econometrics* 210, 268–290.
- Imbens, G.W., Kolesár, M., 2016. Robust standard errors in small samples: Some practical advice. *Review of Economics and Statistics* 98, 701–712.
- James, G.M., Witten, D.M., Hastie, T.J., Tibshirani, R.J., 2021. *An Introduction to Statistical Learning*. Second ed., Springer, New York.
- Lee, C.H., Steigerwald, D.G., 2018. Inference for clustered data. *Stata Journal* 18, 447–460.
- MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2020. Testing for the appropriate level of clustering in linear regression models. QED Working Paper 1428. Queen’s University.
- MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2021. Wild bootstrap and asymptotic inference with multiway clustering. *Journal of Business & Economic Statistics* 39, 509–519.

- MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2022a. Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics* xx, to appear.
- MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2022b. Fast jackknife and bootstrap methods for cluster-robust inference. QED Working Paper 1485. Queen’s University.
- MacKinnon, J.G., Webb, M.D., 2017a. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics* 32, 233–254.
- MacKinnon, J.G., Webb, M.D., 2017b. Pitfalls when estimating treatment effects using clustered data. *The Political Methodologist* 24, 20–31.
- MacKinnon, J.G., Webb, M.D., 2018. The wild bootstrap for few (treated) clusters. *Econometrics Journal* 21, 114–135.
- MacKinnon, J.G., Webb, M.D., 2020. Randomization inference for difference-in-differences with few treated clusters. *Journal of Econometrics* 218, 435–450.
- MacKinnon, J.G., White, H., 1985. Some heteroskedasticity consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics* 29, 305–325.
- Niccodemi, G., Alessie, R., Angelini, V., Mierau, J., Wansbeek, T., 2020. Refining clustered standard errors with few clusters. Working Paper 2020002-EEF. University of Groningen.
- Pustejovsky, J.E., Tipton, E., 2018. Small sample methods for cluster-robust variance estimation and hypothesis testing in fixed effects models. *Journal of Business & Economic Statistics* 36, 672–683.
- Roodman, D., MacKinnon, J.G., Nielsen, M.Ø., Webb, M.D., 2019. Fast and wild: Bootstrap inference in Stata using `boottest`. *Stata Journal* 19, 4–60.