

Jackknife Standard Errors for Clustered Regression

Bruce E. Hansen*

University of Wisconsin, USA

November 2025

This paper presents a theoretical case for replacement of conventional heteroskedasticity-consistent and cluster-robust variance estimators with jackknife variance estimators, in the context of linear regression with heteroskedastic and/or cluster-dependent observations. We examine the bias of variance estimation and the coverage probabilities of confidence intervals. Concerning bias, we show that conventional variance estimators have full downward worst-case bias, while our jackknife variance estimator is never downward biased. Concerning confidence intervals, we show that intervals based on conventional standard errors have worst-case coverage equalling zero, while the jackknife-based confidence interval has coverage probability bounded by the Cauchy distribution, under the auxiliary assumption of normal errors. We also extend the Bell-McCaffrey (2002) student t approximation to our jackknife t -ratio, resulting in confidence intervals with improved coverage probabilities. Our theory holds under broad assumptions, allowing arbitrary cluster sizes, regressor leverage, within-cluster correlation, heteroskedasticity, regression with a single treated cluster, fixed effects, and delete-cluster invertibility failures. Our theoretical findings are consistent with the extensive simulation literature investigating heteroskedasticity-consistent and cluster-robust variance estimation.

*Research support from the NSF and the Phipps Chair are gratefully acknowledged. My thanks to the editor Xavier D'Haultfoeuille and four referees for careful and detailed comments. Over the course of this research, I have received helpful comments and suggestions from many individuals, including Andrew Chesher, Harold Chiang, Grant Hillier, Rustam Ibragimov, James MacKinnon, Ulrich Müller, Morten Nielsen, Marc Paoletta, Peter Phillips, Matthew Webb, and Thilo Welz. Thanks to J.C. Lazzaro for helpful research assistance.

1 Introduction

Heteroskedasticity-consistent (HC) and cluster-robust (CRVE) variance estimators for regression coefficient estimators are foundational for applied economic analysis. They are used for measuring estimation precision, confidence intervals, and tests of hypotheses. Unfortunately, under standard conditions, conventional HC and CRVE variance estimators can be fully downward biased, conventional HC and CRVE t -tests can exhibit unbounded size distortions, and conventional HC and CRVE confidence intervals can have coverage rates equal to zero, even when the regression errors are normally distributed. This situation – full downward bias, unbounded size, and zero coverage probability – can be corrected by replacing conventional variance estimators by an appropriate jackknife estimator, and this option is now available in base Stata. This simple change – the use of jackknife instead of conventional standard errors – has the result that variance estimation is never downward biased and size distortion is bounded.

Our case for jackknife standard errors is based on two new theoretical insights concerning the linear regression model. First, our jackknife variance estimator is conservative (its expectation is larger than the exact variance), while the standard CRVE variance estimator has arbitrarily large downward worst-case estimation bias. The conservative property of the jackknife estimator holds under broad assumptions, allowing arbitrary cluster sizes, regressor leverage, within-cluster correlation, heteroskedasticity, and delete-cluster invertibility failures. Concerning the latter – delete-cluster invertibility failures – this allows for the context where a regressor is non-zero only for a single cluster; for example, regressions with included cluster-level fixed effects, regressions with cluster-level treatment indicators when only one cluster is treated, and regressions with potentially sparse¹ dummy variables. In contrast, conventional variance estimators (including conventional jackknife variance estimators) can fail miserably in such contexts. A new discovery we highlight is that for jackknife variance estimation to be robust to invertibility failures, it is critical to carefully modify the jackknife formula so not to discard clusters. Existing methods (which make *ad hoc* modifications) fail to be robust.

Our second theoretical finding is that under the assumptions of cluster-dependent normality and a stronger rank condition, the worst-case coverage probability of the CRVE-based confidence interval equals zero, while the worst-case coverage probability of our jackknife-based confidence interval is uniformly bounded and controlled by the Cauchy distribution, where uniformity is across all regressor and covariance matrix configurations. Both of these theoretical results assume that the linear regression model is correctly specified, so that the error has a zero conditional mean.

¹Binary variables which are non-zero for only a few observations. This often occurs when categorical variables are interacted.

Furthermore, to provide improved finite sample confidence intervals and p-values, we propose a “Satterthwaite” adjusted student t approximation to the finite sample distribution of the jackknife t -ratio. This adjustment is similar to those of Bell and McCaffrey (2002), Imbens and Kolesár (2016), and Young (2016), but is the first specifically designed for jackknife t -ratios. We present a detailed set of simulation results which demonstrate that our proposed confidence interval has excellent coverage rates in finite samples across a broad range of regression designs. A limitation of this adjustment (similar to the others in this literature) is that it is confined to inference on real-valued parameters and thus excludes joint tests. We recommend that conventional statistical software calculate default confidence intervals and p-values using this Satterthwaite adjustment, instead of the current practice of the student t distribution. The Satterthwaite adjustment is computationally convenient, more conservative than student t inference, and more reliable in finite samples.

The family of heteroskedasticity-consistent variance estimators are often written with the labels HC_0 , HC_1 , HC_2 , and HC_3 . The HC_0 version was introduced by Eicker (1963), Huber (1967), and White (1980). The degrees of freedom correction known as HC_1 was suggested by Hinkley (1977), and became ubiquitous in applied econometric practice by its designation as the default “r” robust option in Stata. Together, HC_0 and HC_1 are known as Eicker-Huber-White (EHW) variance estimators. HC_2 and HC_3 were introduced by MacKinnon and White (1985) as unbiased estimators under homoskedasticity and the jackknife principle, respectively.

The cluster-robust variance estimator (CRVE) was introduced by Liang and Zeger (1986) and Arellano (1987), is available in Stata through its ubiquitous `cluster` standard error option, and currently dominates applied econometric practice. Bell and McCaffrey (2002) introduced two generalizations for clustered regression similar to HC_2 and HC_3 . The review by MacKinnon, Nielsen, and Webb (2023a) use the labels CV_1 , CV_2 , and CV_3 to denote these three estimators. The recognition that finite-sample inference based on EHW confidence intervals can be severely distorted is a recurrent theme in econometrics. Some investigations include Chesher and Jewitt (1987), Chesher (1989), Chesher and Austin (1991), Long and Ervin (2000), and Young (2019). There has been a substantial recent literature proposing improved standard errors over the EHW class under independent sampling. This includes Bera, Suprayitno, and Premaratne (2002), Cattaneo, Jansson, and Newey (2018), and Kline, Saggio, and Sølvssten (2020). These new variance estimators have the advantage that they are (approximately) unbiased, but have as disadvantages that the variance estimators are computationally burdensome in large samples and are not necessarily positive semi-definite. These methods, while promising, have not been generalized to the clustered sampling setting and are not investigated in this paper.

Jackknife standard errors can be paired with conventional critical values (student t or normal) or with alternative methods. The latter include the distributional adjustments of Bell and

McCaffrey (2002), bootstrap percentile- t methods (Cameron, Gelbach, and Miller (2008)), and conditional critical values (Pötscher and Preinerstorfer (2025)). Our recommendation is that default inference should be based on the Bell-McCaffrey-style distributional adjustment, as it is simple to calculate and has excellent performance. Another excellent option for inference is the cluster wild bootstrap, if paired with jackknife standard errors. The recommendation to use jackknife/HC₃ variance estimators is not new. Authors making this recommendation include Efron and Stein (1981), MacKinnon and White (1985), Andrews (1991), Chesher and Austin (1991), Long and Ervin (2000), and MacKinnon, Nielsen and Webb (2023abc).

Our analysis applies to all regression contexts. One where the inadequacy of conventional approximations has received particular attention is regression with a small number of clusters and/or a small number of treated clusters. This literature includes Conley and Taber (2011), Ibragimov and Müller (2016), Ferman and Pinto (2019), Hagemann (2019, 2025), MacKinnon and Webb (2020), Canay, Santos, and Shaikh (2021), and Niccodemi and Wansbeek (2022). These applications may also benefit from jackknife standard errors and our adjusted degrees of freedom approximation, as their application will reduce size distortions.

Our recommended standard errors with adjusted inference is simple to implement in Stata (version 19.5). The syntax for linear regression of y on x clustered by $cvar$ is `regress y x, vce(hc3 cvar, dfadjust hansen)`.

The organization of the paper is as follows. Section 2 introduces cluster-robust variance estimation. Section 3 introduces jackknife variance estimation. Section 4 presents results on variance estimation bias. Section 5 presents the Cauchy distribution bound. Section 6 proposes the Satterthwaite adjusted t distributional approximation. Section 7 discusses computation. Section 8 presents simulation evidence. Section 9 presents an empirical application, extending Meng, Qian and Yared (2015) and Canay, Santos, and Shaikh (2021). A conclusion is presented in Section 10. Some technical proofs are presented in the Appendix. An Online Appendix presents additional technical proofs, additional asymptotic results, and computation details. A companion paper, Hansen (2025), presents several additional applications, to further illustrate the empirical relevance of the proposed methods.

A word on notation. We say that the symmetric square matrix \mathbf{A} is positive-definite if $x' \mathbf{A} x > 0$ for all conformable vectors $x \neq 0$. Similarly, \mathbf{A} is positive semi-definite if $x' \mathbf{A} x \geq 0$ for all $x \neq 0$. We write $\mathbf{A} > \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is positive-definite, and write $\mathbf{A} \geq \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is positive semi-definite.

2 Clustered Regression and the CRVE

The model is a standard clustered sampling regression. The observations are separated into G unbalanced mutually independent clusters. We write the observations on the i th individ-

ual in the g th cluster as (Y_{ig}, X_{ig}) , for $i = 1, \dots, n_g$ and $g = 1, \dots, G$, where Y_{ig} is scalar and X_{ig} is $k \times 1$. We stack the observations by cluster, so that $\mathbf{Y}_g = (Y_{1g}, \dots, Y_{n_gg})'$ is $n_g \times 1$ and $\mathbf{X}_g = (X_{1g}, \dots, X_{n_gg})'$ is $n_g \times k$. The number of observations in the g th cluster is n_g and the total number of observations is $n = \sum_{g=1}^G n_g$. Stacking the observations conventionally, we obtain the full sample (\mathbf{Y}, \mathbf{X}) .

The observations are assumed to satisfy the linear regression model $Y_{ig} = X'_{ig}\beta + e_{ig}$ where β is a $k \times 1$ coefficient vector and e_{ig} is an error. Written at the level of the cluster, the model is

$$\mathbf{Y}_g = \mathbf{X}_g\beta + \mathbf{e}_g \quad (1)$$

$$\mathbb{E}[\mathbf{e}_g] = 0 \quad (2)$$

where $\mathbf{e}_g = (e_{1g}, \dots, e_{n_gg})'$. It is also sometimes convenient to use the full-sample notation $\mathbf{Y} = \mathbf{X}\beta + \mathbf{e}$. We will treat the regressors as fixed, but all results go through in the random regressor setting by conditioning. Notice that model (1)-(2) specifies the conditional expectation of \mathbf{Y}_g as linear in \mathbf{X}_g , and hence excludes the linear projection model. Define the cluster-level covariance matrices

$$\mathbb{E}[\mathbf{e}_g\mathbf{e}'_g] = \boldsymbol{\Sigma}_g. \quad (3)$$

The specification (3) allows the covariance matrices $\boldsymbol{\Sigma}_g$ to be a function of the regressors (and hence conditionally heteroskedastic), and/or to be a function of the cluster g (and hence unconditionally heteroskedastic). We follow the clustering literature and impose no structure on $\boldsymbol{\Sigma}_g$. We also define the full-sample covariance matrix $\boldsymbol{\Sigma} = \text{diag}\{\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_G\}$.

The model (1)-(3) includes heteroskedastic regression as the special case where $n_g = 1$ for all g . We call this the “no clustering” or “absence of clustering” case.

Assumption 1 *Model (1)-(3) holds, \mathbf{X} is full rank, and $\mathbb{E}[e^2_{ig}] < \infty$ for all i and g .*

We focus on the least squares estimator $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{Y})$. It is well known that under Assumption 1, $\hat{\beta}$ is unbiased for β with exact covariance matrix

$$\mathbf{V} = \text{var}[\hat{\beta}] = (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \boldsymbol{\Sigma}_g \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1}.$$

The nearly ubiquitous cluster-robust variance estimator (CRVE₁) for \mathbf{V} is

$$\hat{\mathbf{V}}_1 = \frac{G(n-1)}{(G-1)(n-k)} (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \hat{\mathbf{e}}_g \hat{\mathbf{e}}'_g \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1}, \quad (4)$$

where $\hat{\mathbf{e}}_g = \mathbf{Y}_g - \mathbf{X}_g \hat{\beta}$ denotes the least squares residual vector for the g th cluster. The estimator

(4) with no degrees of freedom adjustment was introduced by Liang and Zeger (1986) and Arellano (1987). The degrees of freedom correction appearing in (4) is added by the Stata “cluster” option. In the absence of clustering, (4) specializes to the HC₁ “heteroskedasticity-robust” or Eicker-White (EHW) estimator of Eicker (1963), Huber (1967), and White (1980), multiplied by the $n/(n-k)$ degrees of freedom correction suggested by Hinkley (1977).

As an alternative to CRVE₁, Bell and McCaffrey (2002), Imbens and Kolesár (2016), and Kolesár (2023) recommended the CRVE₂ variance estimator

$$\hat{\mathbf{V}}_2 = (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \mathbf{M}_g^{+1/2} \hat{\mathbf{e}}_g \hat{\mathbf{e}}'_g \mathbf{M}_g^{+1/2} \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1} \quad (5)$$

where $\mathbf{M}_g^{+1/2}$ is the Moore-Penrose inverse of the square root of the partial projection matrix

$$\mathbf{M}_g = \mathbf{I}_{n_g} - \mathbf{X}_g (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'_g. \quad (6)$$

In the absence of clustering and when all \mathbf{M}_g are invertible, CRVE₂ specializes to the HC₂ estimator of MacKinnon and White (1985). The original definitions of HC₂ and CRVE₂ required that all \mathbf{M}_g are invertible. As discussed by Kolesár (2023), the generalized inverse in (5) allows CRVE₂ to be defined even when \mathbf{M}_g is non-invertible, and this is the implementation in Stata 18 through its `vce(hc2 clustvar)` option.

3 Jackknife Variance Estimation

The jackknife estimator of variance of Tukey (1958) extended to clustered dependence is

$$\hat{\mathbf{V}}_3 = \frac{G-1}{G} \sum_{g=1}^G (\hat{\boldsymbol{\beta}}_{-g} - \bar{\boldsymbol{\beta}}) (\hat{\boldsymbol{\beta}}_{-g} - \bar{\boldsymbol{\beta}})' \quad (7)$$

where

$$\hat{\boldsymbol{\beta}}_{-g} = \left(\sum_{j \neq g} \mathbf{X}'_j \mathbf{X}_j \right)^{-1} \left(\sum_{j \neq g} \mathbf{X}'_j \mathbf{Y}_j \right) = \left(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g \mathbf{X}_g \right)^{-1} \left(\mathbf{X}'\mathbf{Y} - \mathbf{X}'_g \mathbf{Y}_g \right) \quad (8)$$

and $\bar{\boldsymbol{\beta}} = \frac{1}{G} \sum_{g=1}^G \hat{\boldsymbol{\beta}}_{-g}$. The delete-one-cluster estimator $\hat{\boldsymbol{\beta}}_{-g}$ in (8) is obtained by applying least squares to the sample after deleting the observations in cluster g . A variant of $\hat{\mathbf{V}}_3$ is

$$\hat{\mathbf{V}}_4 = \frac{G-1}{G} \sum_{g=1}^G (\hat{\boldsymbol{\beta}}_{-g} - \hat{\boldsymbol{\beta}}) (\hat{\boldsymbol{\beta}}_{-g} - \hat{\boldsymbol{\beta}})', \quad (9)$$

which centers at the full-sample estimator $\hat{\boldsymbol{\beta}}$ rather than at $\bar{\boldsymbol{\beta}}$.

The estimators (7) and (9) as written are not uniquely defined if there is a cluster g for which $\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g$ is noninvertible². The implementation in Stata 18 circumvents this difficulty by excluding from the sums in (7)-(9) any cluster where (8) is undefined³. We follow this interpretation; henceforth, we assume that (7)-(9) are implemented with this modification. We describe (7)-(9) as the “conventional” jackknife variance estimators.

As we show in the next section, the conventional estimators $\widehat{\mathbf{V}}_3$ and $\widehat{\mathbf{V}}_4$ can exhibit downward bias. To eliminate this possibility, we propose the following jackknife estimator:

$$\widehat{\mathbf{V}}_5 = \sum_{g=1}^G (\widehat{\beta}_{-g} - \widehat{\beta}) (\widehat{\beta}_{-g} - \widehat{\beta})' \quad (10)$$

where

$$\widehat{\beta}_{-g} = \left(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g \right)^- \left(\mathbf{X}'\mathbf{Y} - \mathbf{X}'_g\mathbf{Y}_g \right) \quad (11)$$

and \mathbf{A}^- denotes a generalized inverse of \mathbf{A} . $\widehat{\beta}_{-g}$ is a generalized delete-one-cluster estimator. It is defined for all g and therefore (10) includes all clusters. Most of our results will hold for any choice of generalized inverse; in practice, we recommend the Moore-Penrose inverse. With this choice, $\widehat{\beta}_{-g}$ is the unique minimum-length minimizer of the delete-one-cluster least squares criterion. The Moore-Penrose inverse is computationally stable and is available in all statistical packages. Our estimator $\widehat{\mathbf{V}}_5$ differs from $\widehat{\mathbf{V}}_3$ and $\widehat{\mathbf{V}}_4$ in three respects, all of which contribute to the inequality $\widehat{\mathbf{V}}_5 > \widehat{\mathbf{V}}_4 > \widehat{\mathbf{V}}_3$. First, $\widehat{\mathbf{V}}_5$ does not drop noninvertible clusters. Second, $\widehat{\mathbf{V}}_5$ is centered at the full-sample estimator $\widehat{\beta}$ rather than at $\bar{\beta}$. Third, $\widehat{\mathbf{V}}_5$ does not have the degrees of freedom correction $(G-1)/G$. In some applications these differences will be negligible, but in others, as we show later, they can be substantial. In Stata (version 19.5), the estimator $\widehat{\mathbf{V}}_5$ can be obtained by the `vce(hc3 cvar)` option, where `cvar` is a variable which identifies clusters.

Jackknife estimation is ideally suited for the context where the delete-one-cluster estimators $\widehat{\beta}_{-g}$ are well-defined for all clusters, which requires that the matrices $\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g$ are invertible for all g . We call this context *clusterwise invertibility* and its failure *clusterwise noninvertibility*. Noninvertibility occurs when deletion of a single cluster renders the regressor design matrix singular. Most typically, this occurs when a regressor (or some linear combination of regressors) only takes non-zero values for a single cluster. Examples include regressions with included cluster-level fixed effects, regressions with cluster-level treatment indicators when only one cluster is treated, and saturated regressions with sparse cell proportions. Such regressions are commonplace in applications, so it is desirable for variance estimation to be sufficiently flexible to handle their occurrence, which is presumably the motivation for the “drop nonin-

²This is identical to the context where \mathbf{M}_g in (6) is noninvertible.

³This follows the recommendation in Shao and Tu (1995) for noninvertible bootstrap replications.

vertible clusters” modification described above. However, the properties of modifications need to be investigated to preclude undesirable outcomes. When the sample satisfies clusterwise invertibility then the estimators \widehat{V}_4 and \widehat{V}_5 only differ by the degrees of freedom correction $(G - 1)/G$, which is inconsequential when G is large. Under clusterwise noninvertibility, however, they can differ substantially.

The clustered jackknife estimators \widehat{V}_3 and \widehat{V}_4 were developed by Cochran (1977), Rust and Rao (1996), and Bell and McCaffrey (2002). See MacKinnon, Nielsen, and Webb (2023abc). Stata has codified the modification to delete noninvertible clusters. MacKinnon, Nielsen, and Webb (2023c) propose a generalized delete-one-cluster estimator similar⁴ to (11), though they do not investigate its statistical properties.

For alternative algebraic representations see MacKinnon, Nielsen, and Webb (2023b). One we use is based on the delete-one-cluster prediction errors

$$\widehat{\mathbf{e}}_{-g} = \mathbf{Y}_g - \mathbf{X}_g \widehat{\boldsymbol{\beta}}_{-g}. \quad (12)$$

In the proof of Theorem 1 we show that

$$\widehat{\boldsymbol{\beta}}_{-g} - \widehat{\boldsymbol{\beta}} = -(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'_g \widehat{\mathbf{e}}_{-g}. \quad (13)$$

This equality holds even under clusterwise noninvertibility. Given (13), we can write (10) as

$$\widehat{V}_5 = (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \widehat{\mathbf{e}}_{-g} \widehat{\mathbf{e}}'_{-g} \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1}. \quad (14)$$

The jackknife estimators (7) and (10) simplify in the absence of clustering, though existing proposals and analysis assume that \mathbf{M}_g are all invertible⁵. Under these conditions the estimator (7) corresponds to that proposed by MacKinnon and White (1985) and the estimator (10) corresponds to that proposed by Andrews (1991) and Davidson and MacKinnon (1993) and is known as HC₃. When clusterwise invertibility fails, the estimators HC₂ and HC₃ are undefined⁶. For more details, see the monographs of Efron (1982) and Shao and Tu (1995).

While this paper is primarily focused on finite sample theory, it is instructive to observe that under the conditions for asymptotic normality of tests constructed with the CRVE₁ variance estimator (as in Theorem 9 of Hansen and Lee (2019)), tests constructed with the jackknife variance estimator (14) are also asymptotically normal. We present this theory in the Online Appendix.

⁴The differences are that they multiply (10) by $(G - 1)/G$, and implement (11) via the QR algorithm, mimicking the Stata implementation.

⁵In the absence of clustering, this means that the leverage values $h_{ii} = \mathbf{X}'_i (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}_i$ all satisfy $h_{ii} < 1$.

⁶A word of caution: Stata 18 reports HC₂ and HC₃ standard errors even when the regressor matrix is clusterwise noninvertible by giving these observations a weight of zero in the variance calculation. See Kranz (2024).

4 Biased vs Conservative Variance Estimation

In the classical regression model, the classical variance estimator is unbiased for the exact finite sample variance. Not surprisingly, outside the classical model, robust variances estimators are not unbiased. We now examine the worst-case downward bias of the five variance estimators described in the previous sections. We focus on downward bias, as this is the issue which causes undercoverage of confidence intervals and oversized tests. The main contribution of this section is the following result.

Theorem 1 *Under Assumption 1, $\mathbb{E}[\widehat{\mathbf{V}}_5] \geq \mathbf{V}$.*

Theorem 1 shows that our recommended jackknife estimator $\widehat{\mathbf{V}}_5$ is *never downward biased*, or *conservative*. This means that in any regression context, and any sample size, we can be confident that the jackknife estimator is not downward biased. We will find that this bias property is important as it is directly connected to the coverage probabilities of confidence intervals.

Theorem 1 holds quite broadly, holding for all sample sizes, regressor matrices, variance matrices, and violations of clusterwise noninvertibility. In particular, the robustness to clusterwise noninvertibility is new and surprising. Theorem 1 augments Theorem 2 of Bell and McCaffrey (2002), which established that $\widehat{\mathbf{V}}_4$ is never downward biased when the regressors satisfy clusterwise invertibility and the errors e_{ig} are i.i.d. (that is, when $\boldsymbol{\Sigma} = \mathbf{I}_n \sigma^2$).

Theorem 1 is also related to the seminal work of Efron and Stein (1981). Their results are typically described as stating that the same result holds for $\widehat{\mathbf{V}}_3$ but this is incorrect. Instead, Efron and Stein's Theorem 2 states that $\widehat{\mathbf{V}}_3$ is never-downward-biased as an estimator of $\text{var}[\widehat{\beta}]$, not as an estimator of $\text{var}[\widehat{\beta}]$. Furthermore, Theorem 2 below shows $\widehat{\mathbf{V}}_3$ does not satisfy the never-downward-biased property.

We now explore the worst-case bias properties of the other CRVE variance estimators. For these results we focus on individual coefficient estimates and their variance estimators. For some non-zero $k \times 1$ vector R , define the scalar parameter $\theta = R'\beta$. This includes individual coefficients and linear combinations. Its estimator is $\widehat{\theta} = R'\widehat{\beta}$. Under Assumption 1, $\widehat{\theta}$ is unbiased for θ and has exact variance $v^2 = \text{var}[\widehat{\theta}] = R'\mathbf{V}R$. The estimators of v^2 are $\widehat{v}_j^2 = R'\widehat{\mathbf{V}}_j R$ for $j = 1, \dots, 5$. Standard errors for $\widehat{\theta}$ are their square roots $\widehat{v}_j = \sqrt{R'\widehat{\mathbf{V}}_j R}$. To exclude degeneracy, some of our results require that the variance is strictly positive. For technical reasons, we also restrict the extent of clusterwise noninvertibility.

Assumption 2 $v^2 > 0$.

Assumption 3 *Partition $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2]$ and $\mathbf{X}_g = [\mathbf{X}_{1g}, \mathbf{X}_{2g}]$. Suppose that R loads only on \mathbf{X}_1 . Let $\dot{\mathbf{X}}_1$ (with g th component $\dot{\mathbf{X}}_{1g}$) be the residual from the linear regression of \mathbf{X}_1 onto \mathbf{X}_2 . For*

each $g = 1, \dots, G$, either (1) $\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g$ is invertible, or (2) $\dot{\mathbf{X}}'_{1g}\mathbf{X}_{2g} = 0$ and $\dot{\mathbf{X}}'_1\dot{\mathbf{X}}_1 - \dot{\mathbf{X}}'_{1g}\dot{\mathbf{X}}_{1g}$ is invertible.

Assumption 3 is essentially identical to the condition of Lemma 1 of Kolesár (2023). It requires that inference only concerns the “well-identified” variables \mathbf{X}_1 . Essentially, it states that after the controls are partialled out, the leave-cluster-out coefficients for \mathbf{X}_1 are uniquely defined. Assumption 3 holds when \mathbf{X}_2 are cluster-specific fixed effects, but excludes inference on the fixed effects.

It will be convenient to define sets of models. For fixed k and G let \mathcal{F} be the class of all $(\mathbf{X}, \boldsymbol{\Sigma})$ satisfying Assumptions 1-3. Let $\mathcal{F}^* \subset \mathcal{F}$ be the subset where \mathbf{X} satisfies clusterwise invertibility. The following shows that the variance estimators other than $\hat{\mathbf{V}}_5$ can be severely biased.

Theorem 2 *Under Assumptions 1-3,*

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\hat{\mathbf{v}}_1^2]}{\nu^2} = 0, \quad (15)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\hat{\mathbf{v}}_2^2]}{\nu^2} = 0, \quad (16)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\hat{\mathbf{v}}_3^2]}{\nu^2} \leq \left(\frac{G-1}{G}\right)^2 < 1, \quad (17)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\hat{\mathbf{v}}_4^2]}{\nu^2} = \frac{G-1}{G} < 1, \quad (18)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\hat{\mathbf{v}}_3^2]}{\nu^2} = 0, \quad (19)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\hat{\mathbf{v}}_4^2]}{\nu^2} = 0, \quad (20)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\hat{\mathbf{v}}_5^2]}{\nu^2} = 1. \quad (21)$$

Equation (15) shows that the expected value of the scaled CRVE₁ estimator can be arbitrarily close to zero. This means that there is some distribution under which variance estimation has arbitrarily large downward bias. In words, we say that CRVE₁ is *fully downwardly biased*. The proof of (15) focuses on models which satisfy $\boldsymbol{\Sigma} = \mathbf{I}_n$ and clusterwise invertibility, so the result (15) has nothing to do with heteroskedasticity, correlated errors, or invertibility failure. Rather, it is due to extreme regressor leverage (unbalanced regressors and/or cluster sizes). Equation (16) shows a similar result for CRVE₂. The difference between (15) and (16) is that the proof of the latter requires manipulation of the covariance matrices. Thus, CRVE₁ can be fully downward

biased due to regressor leverage alone, while the result for CRVE₂ requires non-i.i.d. errors.

Equations (17) and (18) show that the conventional jackknife estimators \hat{v}_3^2 and \hat{v}_4^2 violate the never-downward-biased property, even under clusterwise invertibility. The magnitude of the violation is small when G is large. Regardless, (17) and (18) show that the conventional interpretation⁷ of the results of Efron and Stein (1981) is incorrect, and that the never-downward-biased result of Bell and McCaffrey (2002) is not robust to non-i.i.d. errors. Equations (19) and (20) extend the analysis of the conventional jackknife estimators to the class of models allowing clusterwise noninvertibility but satisfying Assumption 3. The results show that under these conditions the conventional estimators \hat{v}_3^2 and \hat{v}_4^2 can be fully downward biased. The proof of these results focus on models which satisfy $\Sigma = I_n$. This means that the full downward bias of the conventional jackknife estimators is entirely a consequence of the deletion of noninvertible clusters, and does not require heteroskedasticity or correlated errors.

Equation (21) examines our recommended jackknife estimator \hat{v}_5^2 , allowing for general heteroskedasticity, correlated errors, and clusterwise noninvertibility. As implied by Theorem 1, \hat{v}_5^2 is never downward biased. Equation (21) extends this further and demonstrates that the infimum equals one, meaning that the inequality of Theorem 1 is sharp⁸.

Theorem 2 draws a stark contrast between the five variance estimators. The full downward bias of CRVE₁, CRVE₂, and the conventional jackknife estimators means that on average they can be “much too small” relative to the true variance. In contrast, the never-downward-biased property of \hat{v}_5^2 means that there is no situation where it is expected to be “too small”, even slightly. The model classes studied in Theorem 2 hold fixed the number of regressors k and number of clusters G , but allow the cluster sizes n_g , regressors \mathbf{X} , and covariance matrices Σ to vary freely. It is important to understand that the statements of Theorem 2 hold for all G . Thus the bias of CRVE₁, CRVE₂, and the conventional jackknife can be arbitrarily large in both very small and very large samples.

The worst-case downward bias in (15)-(20) is calculated by studying models with extreme leverage, arising when the regressor of interest has variation which is dominated by a single cluster. Intuitively, when a small number of clusters dominate the sample, standard variance estimators are highly biased towards zero. The least squares estimator overfits the dominating clusters, shrinking the residuals for these clusters relative to the true errors. This leads to downward estimation bias. Conventional fixes, such as the degrees of freedom adjustment of CRVE₁, are insufficient to counter the bias.

One suggestion has been for users to examine their model for features which lead to down-

⁷This distinction was recognized by Efron and Stein (1981). For a further discussion see Section 4.5 of Efron (1982).

⁸While (21) focuses on real-valued parameters, the proof of (21) extends to the full covariance matrix estimator, meaning that $\mathbb{E}[\hat{V}_5]$ can be arbitrarily close to V .

ward bias, such as high leverage points and heterogeneous cluster sizes. Conditional on this information, appropriate standard errors could be selected. We do not recommend this approach. This idea is analogous to selecting EHW standard errors instead of classical standard errors after a test for heteroskedasticity. In general, pre-test estimators have poor finite sample properties due to pre-test bias. It is generally preferred to use robust methods without pre-test selection. Therefore, our recommendation is for the jackknife to be the default method for variance estimation and standard error calculation. Another potential objection to Theorem 2 is that in any given application, \mathbf{X} is fixed, so what may be considered relevant is the worst-case bias over Σ , holding \mathbf{X} fixed. However, it is unclear how to calculate the latter bound, and we leave this investigation to future research.

Theorem 2 shows that the jackknife estimator is never downward biased. It is generically upward biased, which we now characterize. From the proof of Theorem 2 we can calculate

$$\mathbb{E}[\widehat{\mathbf{V}}_5] = \mathbf{V} + (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \mathbf{X}_g \mathbf{V}_{-g} \mathbf{X}'_g \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1} \quad (22)$$

where $\mathbf{V}_{-g} = \text{var}[\widehat{\beta}_{-g}]$ is the variance of the delete-one-cluster estimator. In well-behaved contexts, $\mathbf{V}_{-g} \approx \mathbf{V}$ and the bias component on the right side of (22) will be $O(\frac{1}{n} \|\mathbf{V}\|)$. Thus, in general, the proportionate upward bias will decrease with sample size n . In general, the above expression shows that the upward bias is larger when \mathbf{V}_{-g} is much larger than \mathbf{V} and $\mathbb{E}\|\mathbf{X}_g\|^4$ is large. These are both related to high leverage. Thus we expect $\widehat{\mathbf{V}}_5$ to have largest upward bias under high regressor leverage, which is exactly the context where conventional CRVE estimators are severely downward biased. Equation (22) leaves open the possibility of bias-correction for $\widehat{\mathbf{V}}_5$. We do not pursue this idea in this paper.

5 Cauchy Bound

Given a standard error \widehat{v}_j and a critical value c , a confidence interval for θ is

$$\widehat{C}_j(c) = \widehat{\theta} \pm c\widehat{v}_j. \quad (23)$$

The conventional⁹ critical value for a nominal $100(1 - \alpha)\%$ interval is $c = t_{G-1}^{1-\alpha/2}$, the $1 - \alpha/2$ quantile of the student t distribution with $G - 1$ degrees of freedom¹⁰. In finite samples, the actual coverage rate of the confidence interval (23) will deviate from its nominal level. We now present bounds on the exact coverage rate for intervals (23) constructed with different standard

⁹For example, that used in Stata.

¹⁰Or $n - k$ degrees of freedom in the absence of clustering.

errors. As these are finite sample results, we assume that the errors are normally distributed.

Assumption 4 \mathbf{e}_g is distributed $N(0, \Sigma_g)$.

Assumption 4 states that the error vectors \mathbf{e}_g are normally distributed with an unrestricted within-cluster covariance structure. As described after (3), this allows the covariance matrices to vary both with the regressors as well as the cluster, and thus allows both unconditional and conditional heteroskedasticity. Normality is a strong assumption, and is not meant to be taken literally. Rather, by studying coverage rates under this assumption we can gain insight into their behavior in finite samples without relying on asymptotic approximations. Alternatively, Assumption 4 could be replaced with the assumption that the cluster sums $\mathbf{X}'_g \mathbf{e}_g$ are asymptotically normal as the cluster sizes n_g diverge, as in Canay, Santos, and Shaikh (2021).

Under i.i.d. normal errors, it is well known that confidence intervals with classical standard errors have exact coverage $1 - \alpha$. What may be less well known (or at least less emphasized in standard curricula) is that this result does not extend to confidence intervals constructed with HC and CRVE standard errors. We now present the second main contribution of the paper, which provides a lower bound on the coverage rate of the jackknife confidence interval (when constructed with our recommended standard error \hat{v}_5). Let ζ denote a random variable with the Cauchy distribution.

Theorem 3 *Under Assumptions 1-4, for any $c \geq 1$,*

$$\mathbb{P}[\theta \in \hat{C}_5(c)] \geq \mathbb{P}[|\zeta| \leq c]. \quad (24)$$

Equation (24) shows that the interval $\hat{C}_5(c)$ has coverage probability which is uniformly bounded¹¹ away from zero. The lower bound is the Cauchy distribution. An important implication is that the finite sample coverage probability of the $\hat{C}_5(c)$ jackknife confidence interval has bounded distortion from its nominal level. Other than normality, Assumptions 1-4 are broad, including regression models with extreme leverage, within-cluster correlation, heteroskedasticity, and clusterwise noninvertibility. However, Theorem 3 relies on the normality assumption. It is unclear how the result extends to other contexts, or if a similar bound could be attained in a $G \rightarrow \infty$ asymptotic framework.

We contrast (24) with the coverage of confidence intervals constructed with other standard errors. To do so, we first state an intermediate result which may be of independent interest.

¹¹The bound (24) requires $c \geq 1$. This is not an important restriction for inference as all conventional critical values exceed 1.

Theorem 4 Under Assumptions 1-4, if for some variance estimator \hat{v}^2 and model class $\mathcal{F}_a \subset \mathcal{F}$,

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}_a} \frac{\mathbb{E}[\hat{v}^2]}{v^2} = 0, \quad (25)$$

then for $\hat{C}(c) = \hat{\theta} \pm c\hat{v}$ and any $0 \leq c < \infty$,

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}_a} \mathbb{P}[\theta \in \hat{C}(c)] = 0.$$

Theorem 4 shows that full downward bias of a variance estimator leads to an interval with zero coverage. This provides a simple method to demonstrate that an interval has zero coverage, or equivalently that the size of a hypothesis test equals one. Theorem 4 leads to the following characterization of confidence intervals constructed with \hat{v}_1 , \hat{v}_2 , \hat{v}_3 , and \hat{v}_4 .

Theorem 5 Under Assumptions 1-4, for $\hat{C}_j(c) = \hat{\theta} \pm c\hat{v}_j$ and any $0 \leq c < \infty$,

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \mathbb{P}[\theta \in \hat{C}_1(c)] = 0, \quad (26)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \mathbb{P}[\theta \in \hat{C}_2(c)] = 0, \quad (27)$$

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \mathbb{P}[\theta \in \hat{C}_3(c)] = 0, \quad (28)$$

and

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \mathbb{P}[\theta \in \hat{C}_4(c)] = 0. \quad (29)$$

Equation (26) shows that the worst-case coverage (the exact size) of the CRVE₁ confidence interval equals 0, so that coverage can be arbitrarily distorted from the nominal level, and this holds for *any* critical value c . This means that the confidence interval $\hat{C}_1(c)$ does not uniformly¹² achieve any desired coverage probability. The proof focuses on models which satisfy $\boldsymbol{\Sigma} = \mathbf{I}_n$ and clusterwise invertibility, so (26) is not due to heteroskedasticity, correlated errors, or invertibility failure. Rather, it is a consequence of extreme regressor leverage (unbalanced regressors and/or cluster sizes). Theorem 5 is similar to Theorem 4.2 of Preinerstorfer and Pötscher (2016), which establishes results similar to (26)-(29) in the context of heteroskedastic linear regression with HC₁-HC₃ standard errors. While Theorem 5 is stated under the normality assumption, normality is not essential (the theorem holds under Assumptions 1-3 alone), as the theorem extends to any enlarged model class which includes normality as a special case.

Equation (27) shows a similar result for the CRVE₂ confidence interval: the worst-case coverage of the CRVE₂ confidence interval equals 0. Again, this demonstrates that coverage can be

¹²Here, “uniformly” means over all regression designs \mathbf{X} .

arbitrarily distorted from the nominal level. The difference with (26) is that the proof of (27) requires the model class to include non-i.i.d. errors.

Equations (28) and (29) show that the conventional jackknife intervals also have worst-case coverage of 0 if the model class is broadened to include clusterwise noninvertibility. This is due to invertibility failure. These results show that the conventional (e.g., Stata) modification, which is explicitly intended to allow regressions with clusterwise noninvertibility, is not actually robust to clusterwise noninvertibility. These results hold under exactly the same conditions as Theorem 3, highlighting a key difference between the conventional jackknife standard errors \hat{v}_3 and \hat{v}_4 and our proposed standard error \hat{v}_5 .

The results (26)-(29) should not be surprising given Theorem 2, which showed that the variance estimators $\hat{v}_1^2 - \hat{v}_4^2$ can be arbitrarily downward biased. What is important about these results is that they show that these confidence intervals have no *a priori* guarantee that they are in any sense a confidence interval. Furthermore, the zero coverage rates of (26)-(29) cannot be fixed by simply using a larger critical value c , as these results hold for any finite c .

Returning to the confidence interval $\hat{C}_5(c)$, equation (24) bounds its smallest coverage probability in *any* regression satisfying Assumptions 1-4. For example, with the conventional $c = 1.96$ critical value, the bound (24) is 0.70. Thus, the finite sample coverage of \hat{C}_5 with $c = 1.96$ can never be less than 70%. Similarly, the finite sample size of a t -test using the standard error \hat{v}_5 and critical value $c = 1.96$ can never be greater than 30%.

Another implication of (24) is that the Cauchy distribution can be used for finite sample inference (substituting the Cauchy for student t critical values). Doing so will produce inferential statements (hypothesis tests and confidence intervals) with uniform size control. This uniformity holds over all regression designs X and error variances Σ . In practice, however, it is unlikely that researchers will use the Cauchy distribution for inference, as it is exceedingly conservative. For example, while the 5% normal critical value is 1.96, that for the Cauchy distribution is 12.7. It is difficult to imagine a user declaring a t -ratio equalling 10 to be “insignificant” simply because it is less than the Cauchy critical value. Instead, one practical message of Theorems 3-5 is that conventional EHW and CRVE confidence intervals can have arbitrary coverage distortion, while the jackknife interval has bounded distortion. This is a strong motivation for replacement of the EHW and CRVE standard errors by simple-to-calculate jackknife standard errors.

Another important feature of Theorem 3 is that it provides the first generally-applicable uniformly-valid confidence interval for clustered and heteroskedastic regression models. As discussed above, combining jackknife standard errors with Cauchy critical values is uniformly valid, regardless of the regressor or error variance structure. The only other proposals providing partial solutions are Ibragimov and Müller (2010, 2016) and Pötscher and Preinerstorfer (2025). The papers of Ibragimov and Müller achieve uniform size control by estimating the model sepa-

rately on q sub-samples, so as to produce q independent t -ratios. Such sub-sample estimation is not generally feasible, however. In the heteroskedastic regression model, Pötscher and Preinerstorfer (2025) propose numerically maximizing the largest exact critical value $c(\Sigma)$ over the class of feasible variance matrices, conditional on the regressors \mathbf{X} . Such maximization may be feasible when n is very small (their examples set $n = 25$) but is computationally infeasible for typical sample sizes, and is infeasible under clustered dependence.

The results of Theorems 4 and 5 are related to previous impossibility results. Most famously, Bahadur and Savage (1956) show that any confidence interval for the mean of an unknown distribution (without restrictions on the class of distributions) has an exact size of 0. Similar impossibility results are obtained (in different contexts) by Romano (2004), Hirano and Porter (2012), and Bertanha and Moreira (2020). The common feature of these results is the *impossibility* of confidence interval construction. In contrast, Theorem 5 shows that specific (commonly used) confidence intervals have an exact size of 0, while Theorem 3 shows that a properly constructed jackknife confidence interval has correct size. Therefore, our results are not impossibility results, and we do not describe them as such. Our results also bear resemblance to the following: Dufour (1997), who shows that when a parameter is locally almost unidentified, valid confidence intervals must be unbounded with positive probability; Dufour (2003), who shows that inference on the median of symmetric distributions can only be achieved via a sign test; and Coudin and Dufour (2009), who extend this analysis to median regression.

6 Satterthwaite Approximation

In this section we present a computationally attractive method for critical value construction with improved coverage relative to conventional methods. The method is based on the Satterthwaite (1946) approximation for a weighted sum of chi-squares, and is similar to those proposed by Bell and McCaffrey (2002), Imbens and Kolesár (2016), and Young (2016) for HC_1 and HC_2 confidence intervals. Our proposed Satterthwaite $100(1 - \alpha)\%$ confidence interval for θ is

$$\tilde{C}_5 = \hat{\theta} \pm \frac{t_K^{1-\alpha/2} \hat{v}_5}{a} \quad (30)$$

where \hat{v}_5 is the jackknife standard error, $t_K^{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the student t distribution with K degrees of freedom, and a is a scale adjustment. The degrees of freedom K and scale adjustment a are sample-dependent and given below by the formula (34)-(39).

Our proposed Satterthwaite p-value for a test of $\theta = \theta_0$ can be calculated using either the

student t or the F distribution. Let $T = |\hat{\theta} - \theta| / \hat{v}_5$. Our proposed Satterthwaite p-value is

$$p = 2(1 - G(aT; K)) = 1 - F(a^2 T^2; 1, K) \quad (31)$$

where $G(x; K)$ is the student t distribution with K degrees of freedom and $F(x; 1, K)$ is the F distribution with degrees of freedom $(1, K)$. As for the confidence interval (30), the p-value (31) makes both a degrees of freedom adjustment (through K) and a scale adjustment (through a).

The foundation is the following representation for the distribution of the jackknife t -ratio.

Theorem 6 *Under Assumptions 1-4,*

$$\frac{\hat{\theta} - \theta}{\hat{v}_5} = \frac{\xi_0}{\sqrt{\sum_{g=1}^G \lambda_g \xi_g^2}} \quad (32)$$

where ξ_0, \dots, ξ_G are $N(0, 1)$, ξ_1, \dots, ξ_G are mutually independent, λ_g are the eigenvalues of

$$\mathbf{D} = (\mathbf{S} + \mathbf{U}\mathbf{X}'\boldsymbol{\Sigma}\mathbf{X}\mathbf{U}' - \mathbf{V}\mathbf{V}' - \mathbf{U}\mathbf{V}') / v^2, \quad (33)$$

$v^2 = R'(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\boldsymbol{\Sigma}\mathbf{X})^{-1}(\mathbf{X}'\mathbf{X})^{-1}R$, \mathbf{U} and \mathbf{V} are the $G \times k$ matrices

$$\mathbf{U} = \begin{bmatrix} \mathbf{U}'_1 \\ \vdots \\ \mathbf{U}'_G \end{bmatrix}, \quad \mathbf{V} = \begin{bmatrix} \mathbf{V}'_1 \\ \vdots \\ \mathbf{V}'_G \end{bmatrix}$$

with rows $\mathbf{V}_g = \mathbf{X}'_g \boldsymbol{\Sigma}_g \mathbf{T}_g$ and $\mathbf{U}_g = (\mathbf{X}'\mathbf{X} - \mathbf{X}'_g \mathbf{X}_g)^- \mathbf{X}'_g \mathbf{X}_g (\mathbf{X}'\mathbf{X})^{-1} R$, where $\mathbf{T}_g = \mathbf{Z}_g + \mathbf{X}_g \mathbf{U}_g$, and \mathbf{S} is the $G \times G$ diagonal matrix with diagonal elements $S_g = \mathbf{T}'_g \boldsymbol{\Sigma}_g \mathbf{T}_g$.

Theorem 6 shows that the t -ratio can be written as the ratio of a normal random variable to the square root of a weighted sum of chi-square random variables. The weights are a computable function of the regressors \mathbf{X} and covariance matrices $\boldsymbol{\Sigma}_g$. As the covariance matrices are unknown, we approximate the above distribution by replacing the unknown $\boldsymbol{\Sigma}_g$ with a reference choice $\boldsymbol{\Sigma}_g^0 \sigma^2$ for arbitrary scale σ^2 . Following the previous literature, we recommend $\boldsymbol{\Sigma}_g^0 = \mathbf{I}_{n_g}$, but any other choice can be implemented. The choice $\boldsymbol{\Sigma}_g^0 = \mathbf{I}_{n_g}$ has the special implication that in (32), the random variables ξ_0 and ξ_1^2, \dots, ξ_G^2 are independent. For other choices of $\boldsymbol{\Sigma}_g^0$ this independence does not generally hold.

The Satterthwaite approximation simplifies the distribution (32). It replaces the weighted sum of chi-squares in the denominator of (32) by $a^2 \chi_K^2 / K$, where χ_K^2 is a single chi-square random variable with K degrees of freedom, and a and K are free parameters selected so that the

first two moments of $\sum_{g=1}^G \lambda_g \xi_g^2$ and $a^2 \chi_K^2 / K$ match. This matching is achieved by setting

$$a = \sqrt{\sum_{g=1}^G \lambda_g} \quad (34)$$

and

$$K = \frac{\left(\sum_{g=1}^G \lambda_g\right)^2}{\sum_{g=1}^G \lambda_g^2}. \quad (35)$$

Furthermore, when $\Sigma_g^0 \neq \mathbf{I}_{n_g}$, the variables ξ_0 and χ_K^2 are approximated by independent versions. With these simplifications, (32) equals T_K / a , where T_K is a student t random variable with K degrees of freedom. This leads to the confidence interval and p-value (30) and (31). An important limitation is that the adjustment is explicitly limited to inference on real-valued parameters. It does not apply to tests of joint hypotheses and (to my knowledge) has not been extended to joint tests. If feasible, this would be a useful extension.

In what sense is the Satterthwaite approximation (30)-(31) an approximation? Take the reference model $\Sigma_g^0 = \mathbf{I}_{n_g}$. We calculate¹³ that the difference between the coverage probabilities is bounded by the uniform absolute difference between the distributions of $\sum_{g=1}^G \lambda_g \xi_g^2$ and $a^2 \chi_K^2 / K$. Khuri (1995) showed that this difference converges to 0 as $G \rightarrow \infty$. This is consistent with the extensive simulations of Bodenham and Adams (2016), who showed that this difference is numerically less than 0.01 when $G \geq 100$.

We can also show that the adjustments vanish in well-balanced large samples, so the confidence interval and p-value (30) and (31) converge to those based on the standard normal. Let $\|\mathbf{A}\|$ denote the spectral matrix norm.

Theorem 7 *If $\left\|\left(\frac{1}{n} \mathbf{X}' \mathbf{X}\right)^{-1}\right\| = O(1)$ and $n^{-1} \max_{g \leq G} \|\mathbf{X}_g\|^2 = o(1)$ as $n \rightarrow \infty$, and the coefficients (34) and (35) are calculated with $\Sigma_g^0 = \mathbf{I}_{n_g}$, then $a \rightarrow 1$ and $K \rightarrow \infty$ as $n \rightarrow \infty$.*

Theorem 7 shows that under mild conditions (the design matrix is uniformly non-singular and no single cluster dominates the sample) the scale adjustment coefficient a converges to 1, and the degrees of freedom parameter K diverges to infinity. Consequently, the interval (30) collapses to that based on the standard normal distribution in well-balanced large samples.

¹³Let the distribution functions of ξ_0^2 , $\sum_{g=1}^G \lambda_g \xi_g^2$, and $a^2 \chi_K^2 / K$ be denoted by F , G , and G_K , with associated density functions f , g , and g_K . Let $T^2 = \xi_0^2 / \sum_{g=1}^G \lambda_g \xi_g^2$ and $T_K^2 = \xi_0^2 / a^2 \chi_K^2 / K$. Let $\delta = \sup_x |G(x) - G_K(x)|$. Then $|\mathbb{P}[T^2 \leq x] - \mathbb{P}[T_K^2 \leq x a^2]| = \left| \int F(xq) (g(q) - g_K(q)) dq \right| = x \left| \int f(xq) (G(q) - G_K(q)) dq \right| \leq \delta x \int f(xq) dq = \delta$.

7 Computation

The expressions (34)-(35) are written in terms of the eigenvalues of \mathbf{D} . In practice, computation of the eigenvalues is not necessary. Alternative computational expressions are as follows. Let D_{gh} denote the gh th element of \mathbf{D} . Then (34)-(35) equal

$$a = \sqrt{\text{tr}[\mathbf{D}]} = \sqrt{\sum_{g=1}^G D_{gg}} \quad (36)$$

$$K = \frac{(\text{tr}[\mathbf{D}])^2}{\text{tr}[\mathbf{D}\mathbf{D}]} = \frac{\left(\sum_{g=1}^G D_{gg}\right)^2}{\sum_{g=1}^G \sum_{h=1}^G D_{gh}^2}. \quad (37)$$

These expressions are computationally preferred to (34)-(35), and are convenient for small G .

For large G , explicit calculation of the $G \times G$ matrix \mathbf{D} should be avoided. In this case we can replace $\text{tr}[\mathbf{D}]$ and $\text{tr}[\mathbf{D}\mathbf{D}]$ in (36)-(37) with matrix operations. We focus on the reference choice $\Sigma_g^0 = \mathbf{I}_{n_g}$, in which case $v^2 = R'(\mathbf{X}'\mathbf{X})^{-1}R$, $\mathbf{V}_g^0 = \mathbf{X}'_g\mathbf{T}_g$, and $S_g = \mathbf{T}'_g\mathbf{T}_g$. Then, expressions¹⁴ for the components needed for (36)-(37) are

$$\text{tr}[\mathbf{D}] = (\text{tr}[\mathbf{S}] - \text{tr}[\mathbf{V}'\mathbf{U}]) / v^2, \quad (38)$$

$$\begin{aligned} \text{tr}[\mathbf{D}\mathbf{D}] = & (\text{tr}[\mathbf{S}\mathbf{S}] - 2\text{tr}(\mathbf{V}'\mathbf{S}\mathbf{U}) + \text{tr}(\mathbf{U}'\mathbf{U}\mathbf{X}'\mathbf{X}\mathbf{U}'\mathbf{U}\mathbf{X}'\mathbf{X}) \\ & - 4\text{tr}(\mathbf{U}'\mathbf{U}\mathbf{X}'\mathbf{X}\mathbf{U}'\mathbf{V}) + 2\text{tr}(\mathbf{U}'\mathbf{V}\mathbf{U}'\mathbf{V}) + 2\text{tr}(\mathbf{V}'\mathbf{V}\mathbf{U}'\mathbf{U})) / v^4. \end{aligned} \quad (39)$$

The expressions (38)-(39) can be evaluated without explicit computation of \mathbf{D} . Numerically, we find that the expressions (38)-(39) are (roughly) computationally faster than (36)-(37) if $G > k$, while (36)-(37) are computationally faster when $G < k$. The derivation of (38)-(39) is presented in the Online Appendix. For clarity, we summarize our recommended procedure for variance estimation, standard errors, confidence intervals, and p-value construction.

1. Calculate the jackknife variance $\widehat{\mathbf{V}}_5$ from (10) using the delete-cluster estimator (11) and the Moore-Penrose generalized inverse.
2. For any linear coefficient $\theta = R'\beta$:

(a) Set $\widehat{\theta} = R'\widehat{\beta}$.

(b) Calculate its standard error as $\widehat{v}_5 = \sqrt{R'\widehat{\mathbf{V}}_5R}$.

¹⁴The expressions (38)-(39) use the assumption that the generalized inverse \mathbf{A}^- is a reflective generalized inverse (and thus satisfies $\mathbf{A}^-\mathbf{A}\mathbf{A}^- = \mathbf{A}^-$) which is true of the Moore-Penrose inverse but not all generalized inverses.

- (c) Calculate the reference adjustment parameters a and K using $\Sigma_g^0 = I_{n_g}$ with the expressions (36)-(37) if $G \leq k$ and (38)-(39) if $G > k$.
- (d) Calculate the critical value $t_K^{1-\alpha/2}$ from the student t distribution with degrees of freedom K .
- (e) Calculate the confidence interval \tilde{C}_5 using (30).
- (f) Calculate the p-value for a test of $\theta = \theta_0$ using (31).

The adjustment parameters a and K are specific to the linear combination R , and hence specific to each individual coefficient. Therefore, the above steps need to be repeated for each element of β in order to calculate intervals and p-values for all coefficients. In the above procedure, we use the scale adjustment a to adjust the intervals and p-values via (30) and (31). A computationally equivalent method is to adjust the standard errors, by calculating and reporting $\hat{v}_5^* = \hat{v}_5/a$, say. I do not recommend this second method, however, as the adjusted standard errors \hat{v}_5^* do not satisfy the “never downward biased” property of \hat{v}_5 . As the standard error is of first-order importance, and many users only examine coefficient estimates and standard errors, it is prudent to report the standard errors without adjustment, along with the adjusted intervals and p-values. This will result in more reliable interpretation of empirical results.

The adjusted confidence intervals and p-values as described above are available in Stata (version 19.5). To be specific, the syntax for linear regression of y on x clustered by a variable $cvar$ is `regress y x, vce(hc3 cvar, dfadjust hansen)`.

It is common in panel/clustering contexts to include cluster-specific fixed effects. This possibility is allowed in our framework either by explicit inclusion of fixed effect dummy variables in the regressor matrix \mathbf{X} , or by specifying \mathbf{Y}_g and \mathbf{X}_g as within-transformed. In the latter case, equation (1) represents the model *after* the within transformation has been applied. Note that in this case, the covariance matrix (3) is that of the within-transformed errors, not the original equation errors. As (3) is unstructured and allows Σ_g to be singular, this is without loss of generality. When the fixed effects are applied at the same level as the error clustering, our recommendation is to first apply the within transformation and then apply least squares estimation to the within transformed variables. This is both computationally and theoretically preferred. It is theoretically preferred because the model after the within transformation is clusterwise invertible, so inference on any linear coefficient is valid under our theory. In contrast, in the model with the included fixed effect dummy variables, only inference on linear coefficients satisfying Assumption 3 is justified under our theory.

8 Simulation Evidence

We present a simulation experiment to investigate performance. The experiment concerns the clustered linear regression (1)-(3). The goal is 95% confidence intervals for the slope coefficients. Our baseline model is the simple regression $Y_g = \mathbf{1}_g \alpha + \mathbf{X}_g \beta + \mathbf{e}_g$ with \mathbf{X}_g a single $n_g \times 1$ regressor. In our baseline specification, the cluster sizes are homogeneous with $n_g = 10$ observations per cluster. Within this model we consider six designs which vary the distributions of the regressors \mathbf{X}_g and the heteroskedasticity of the equation errors. Let \mathbf{I}_{n_g} denote the $n_g \times n_g$ identity matrix, $\mathbf{1}_g = (1, 1, \dots, 1)'$ the $n_g \times 1$ vector of ones, and $\mathbf{h}_g = (1, -1, 1, -1, \dots)'$ the $n_g \times 1$ vector containing alternating ± 1 .

The designs for the regressors are:

1. Normal with Clustered Dependence: $\mathbf{X}_g \sim N\left(\mathbf{1}_g, \mathbf{I}_{n_g} + \mathbf{1}_g \mathbf{1}_g'\right)$.
2. LogNormal with Clustered Dependence: $\mathbf{X}_g \sim \exp\left(N\left(0, \mathbf{I}_{n_g} + \mathbf{1}_g \mathbf{1}_g'\right)\right)$.
3. Dummy: $X_{ig} = \begin{cases} 10, & \text{if } g = 1 \text{ and } i \leq 2, \text{ or } g = 2 \text{ and } i = 1 \\ 1, & \text{otherwise} \end{cases}$.

In the first design the regressors have a cluster-level random effects covariance matrix. In the second design the regressors are log-normally distributed with the same cluster-level random effects covariance matrix. In the third design the regressor is binary, taking the value 10 for three observations and the value 1 for all other observations. As documented by Chesher and Jewitt (1987) for the LogNormal design, and Imbens and Kolesár (2016) for the Dummy designs, these two designs produce highly leveraged regressor matrices and standard robust covariance matrix estimators can be highly biased.

The designs for the errors are:

1. Normal with Clustered Dependence: $\mathbf{e}_g = \mathbf{u}_g \sim N\left(0, \mathbf{I}_{n_g} + \mathbf{1}_g \mathbf{1}_g' + \mathbf{h}_g \mathbf{h}_g'\right)$.
2. Heteroskedastic: $\mathbf{e}_g = \mathbf{X}_g \odot \mathbf{u}_g$ where \odot denotes Hadamard product.

In the first design, the errors are normally distributed with a cluster-level factor covariance matrix. The second design adds conditional heteroskedasticity. The three regressor designs and two error designs combine for a total of six designs. For each design, the number of clusters G is varied among $\{6, 12, 40, 100\}$. As there are 10 observations per cluster, the associated total sample sizes are $\{60, 120, 400, 1000\}$.

For each simulation replication we estimate the coefficients by least squares. We calculate the five standard errors described in the text: $CRVE_1(\hat{\nu}_1)$, $CRVE_2(\hat{\nu}_2)$, the two conventional

jackknife estimators (\hat{v}_3 and \hat{v}_4), and our recommended estimator \hat{v}_5 . We calculate confidence intervals using eight methods. The first five are conventional, based on the five standard errors and conventional student t critical values. Thus, given each standard error \hat{v}_j , we form the confidence interval $\hat{\beta} \pm t_{G-1}^{0.975} \hat{v}_j$ where $t_{G-1}^{0.975}$ is the 0.975 quantile of the t_{G-1} distribution. We use the $t_{G-1}^{0.975}$ critical value as this is the current implementation in Stata for cluster-robust inference. The next is the Bell and McCaffrey (2002) adjusted t interval (BM) based on the CRVE₂ standard error. It equals $\hat{\beta} \pm t_K^{0.975} \hat{v}_2$ where K is calculated similar to (35) under the reference model $\Sigma_g^0 = \mathbf{I}_{n_g}$.

The next interval is the restricted wild cluster bootstrap symmetric percentile- t interval, using the standard errors \hat{v}_5 and 999 bootstrap replications. The wild clustered bootstrap is proposed and studied by Cameron, Gelbach, and Miller (2008), Djogbenou, MacKinnon, and Nielsen (2019), Canay, Santos, and Shaikh (2021), and MacKinnon, Nielsen and Webb (2023b). The latter paper provides considerable simulation evidence regarding the performance of several implementations, and their evidence strongly supports using this implementation, which they call WCR-V. We follow this implementation, though using $N(0, 1)$ auxiliary bootstrap errors¹⁵, and constructing confidence intervals through test inversion. We provide a thorough description of our implementation in the Online Appendix.

Our final interval is our Satterthwaite interval (30) using the reference model $\Sigma_g^0 = \mathbf{I}_{n_g}$.

We compute the actual coverage probability of these nominal 95% intervals by simulation with 20,000 replications. These estimates are precise, as their standard errors are all less than 0.003. We also compute the average length of these intervals, though we only report this length for the final four confidence intervals (conventional with jackknife \hat{v}_5 , Bell-McCaffrey, wild bootstrap, and Satterthwaite).

We report the results for the baseline regression model in Table 1. The top block reports the results for $G = 6$. The first five columns are the coverage probabilities of the conventional confidence intervals. We can see that the conventional CRVE₁ confidence interval has substantial under-coverage in most designs. The worst case is LogNormal regressors with heteroskedastic error variances, where the interval has only 57% coverage. The CRVE₂ confidence interval has slightly better coverage, but still substantially undercovers in most designs. The jackknife \hat{v}_5 confidence interval has better coverage, but undercovers (as low as 85%) in the leveraged/heteroskedastic designs.

The next column is the coverage probability of the BM confidence interval, which has improved coverage relative to the conventional methods. Its coverage rates are somewhat sensitive

¹⁵MacKinnon, Nielsen and Webb (2023b) recommend Rademacher auxiliary errors. This works poorly in our simulation design due to the small cluster sizes and heavy regressor leverage. Consequently, we use normal auxiliary errors, which has much better performance.

Table 1: Baseline Regression Model. Coverage of Nominal 95% Confidence Intervals for β

		Coverage Probability							Interval Length				
	Cr. Value	Conventional t_{G-1}					BM	Wild	Satt	Jack	BM	Wild	Satt
	St. Error	\hat{v}_1	\hat{v}_2	\hat{v}_3	\hat{v}_4	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5
$G = 6$ Clustered	Normal	0.91	0.93	0.95	0.95	0.96	0.95	0.91	0.96	1.4	1.3	0.5	1.4
	LogNorm	0.85	0.91	0.95	0.96	0.96	0.99	0.95	0.99	0.7	0.8	0.3	1.1
	Dummy	0.86	0.89	0.93	0.93	0.94	1.00	0.96	1.00	0.8	2.0	0.4	2.2
Hetero.	Normal	0.89	0.91	0.93	0.93	0.94	0.93	0.90	0.94	3.2	2.9	2.7	3.2
	LogNorm	0.57	0.70	0.86	0.87	0.89	0.90	0.92	0.93	5.8	6.3	16	9.1
	Dummy	0.70	0.77	0.84	0.84	0.85	0.95	0.95	0.95	7.2	19	48	21
$G = 12$ Clustered	Normal	0.92	0.93	0.95	0.95	0.95	0.95	0.92	0.95	0.8	0.8	0.2	0.8
	LogNorm	0.84	0.89	0.94	0.94	0.94	0.99	0.95	0.99	0.3	0.4	0.0	0.5
	Dummy	0.75	0.81	0.88	0.88	0.89	1.00	0.95	1.00	0.6	2.2	0.5	2.2
Hetero.	Normal	0.91	0.92	0.93	0.93	0.94	0.94	0.92	0.94	1.9	1.8	1.2	1.9
	LogNorm	0.61	0.73	0.86	0.86	0.88	0.92	0.93	0.94	4.2	5.3	11	7.5
	Dummy	0.64	0.74	0.82	0.82	0.83	0.95	0.94	0.95	6.2	20	51	22
$G = 40$ Clustered	Normal	0.94	0.94	0.95	0.95	0.95	0.95	0.94	0.95	0.4	0.4	0.0	0.4
	LogNorm	0.86	0.90	0.93	0.93	0.93	0.97	0.94	0.98	0.1	0.1	0.0	0.2
	Dummy	0.65	0.75	0.83	0.83	0.84	1.00	0.94	0.99	0.6	2.2	0.5	2.3
Hetero.	Normal	0.93	0.94	0.94	0.94	0.94	0.94	0.94	0.94	1.0	1.0	0.3	1.0
	LogNorm	0.70	0.79	0.88	0.88	0.89	0.93	0.93	0.95	3.0	3.6	5.4	4.9
	Dummy	0.60	0.71	0.81	0.81	0.81	0.95	0.95	0.95	5.7	22	52	22
$G = 100$ Clustered	Normal	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.2	0.2	0.0	0.2
	LogNorm	0.89	0.91	0.93	0.93	0.93	0.97	0.93	0.97	0.1	0.1	0.0	0.1
	Dummy	0.61	0.72	0.81	0.81	0.81	0.98	0.94	0.97	0.6	2.3	0.5	2.3
Hetero.	Normal	0.94	0.94	0.95	0.95	0.95	0.94	0.94	0.95	0.6	0.6	0.1	0.6
	LogNorm	0.76	0.83	0.89	0.89	0.90	0.93	0.93	0.95	2.3	2.6	3.2	3.4
	Dummy	0.59	0.70	0.80	0.80	0.80	0.95	0.95	0.95	5.6	22	52	22

to the design, being excessively conservative (100% coverage) in some designs and undercovering (90%) in others. The next column is the coverage probability of the wild bootstrap. The coverage rates are generally excellent when compared to conventional methods. The intervals undercover (90%) in some designs. There are no cases of excessive conservative coverage. The following column is the coverage probability of the Satterthwaite interval. The coverage rates are excellent, uniformly exceeding 93%. The intervals are excessively conservative (100%) in some designs.

It is instructive to use the results to separate the effect of the jackknife standard error from the effect of the Satterthwaite correction. Take the LogNormal heteroskedastic design. Moving from $CRVE_1$ to jackknife standard errors improves the coverage rates from 57% to 89%, and the Satterthwaite correction improves the coverage from 89% to 93%. Take the Dummy het-

Table 2: Regression on Dummy. Coverage of Nominal 95% Confidence Intervals for β

		Coverage Probability							Interval Length				
	Cr. Value	Conventional t_{G-1}					BM	Wild	Satt	Jack	BM	Wild	Satt
	St. Error	\hat{v}_1	\hat{v}_2	\hat{v}_3	\hat{v}_4	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5
$G = 6$ Clustered	Normal	0.90	0.93	0.93	0.94	0.96	0.95	0.91	0.96	1.4	1.3	0.5	1.4
	LogNorm	0.85	0.91	0.92	0.93	0.96	0.99	0.95	0.99	0.7	0.8	0.3	1.1
	Dummy	0.81	0.86	0.77	0.80	0.91	1.00	0.97	0.99	0.6	2.1	0.4	2.1
Hetero.	Normal	0.89	0.91	0.90	0.91	0.94	0.93	0.90	0.94	3.2	2.9	2.8	3.2
	LogNorm	0.58	0.70	0.79	0.81	0.89	0.90	0.92	0.93	5.9	6.5	16	9.1
	Dummy	0.71	0.78	0.67	0.69	0.85	0.96	0.95	0.96	6.9	22	48	22
$G = 12$ Clustered	Normal	0.92	0.93	0.94	0.94	0.95	0.95	0.92	0.95	0.8	0.8	0.2	0.8
	LogNorm	0.83	0.89	0.92	0.92	0.95	0.99	0.94	0.99	0.3	0.4	0.0	0.5
	Dummy	0.73	0.81	0.73	0.74	0.87	1.00	0.96	0.99	0.5	2.1	0.4	2.1
Hetero.	Normal	0.91	0.92	0.93	0.93	0.94	0.94	0.92	0.94	1.9	1.9	1.2	1.9
	LogNorm	0.61	0.73	0.83	0.84	0.88	0.92	0.93	0.94	4.3	5.3	11	7.5
	Dummy	0.66	0.74	0.66	0.67	0.82	0.95	0.95	0.95	5.9	22	48	22
$G = 40$ Clustered	Normal	0.94	0.94	0.95	0.95	0.95	0.95	0.94	0.95	0.4	0.4	0.0	0.4
	LogNorm	0.86	0.90	0.93	0.93	0.93	0.98	0.94	0.98	0.1	0.1	0.0	0.2
	Dummy	0.67	0.77	0.70	0.71	0.85	0.99	0.95	0.98	0.5	2.1	0.4	2.1
Hetero.	Normal	0.93	0.94	0.94	0.94	0.94	0.94	0.94	0.94	1.0	1.0	0.3	1.0
	LogNorm	0.70	0.79	0.87	0.88	0.89	0.93	0.93	0.95	3.0	3.6	5.5	4.9
	Dummy	0.62	0.72	0.65	0.65	0.81	0.95	0.95	0.95	5.4	23	48	23
$G = 100$ Clustered	Normal	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.2	0.2	0.0	0.2
	LogNorm	0.89	0.91	0.93	0.93	0.93	0.97	0.93	0.97	0.1	0.1	0.0	0.1
	Dummy	0.66	0.76	0.70	0.70	0.84	0.98	0.95	0.97	0.5	2.1	0.4	2.1
Hetero.	Normal	0.94	0.94	0.94	0.94	0.95	0.94	0.94	0.95	0.6	0.6	0.1	0.6
	LogNorm	0.76	0.83	0.89	0.89	0.90	0.93	0.93	0.95	2.3	2.6	3.2	3.4
	Dummy	0.61	0.71	0.65	0.65	0.80	0.95	0.95	0.95	5.3	23	49	23

eroskedastic design. Moving from CRVE₁ to jackknife improves the coverage rates from 70% to 85%, and the Satterthwaite further improves the coverage to 95%. In general, we see that the largest effect is due to the use of jackknife standard errors, with the Satterthwaite correction a secondary, but important, effect.

The final four columns report the average length of four confidence interval methods: conventional with jackknife \hat{v}_5 , Bell-McCaffrey, wild bootstrap, and Satterthwaite. We focus on these four as they are the only methods with reasonable coverage rates. The lengths vary as expected. In cases where the confidence intervals have similar coverage rates, then the interval lengths are also similar. In cases where the confidence intervals have differing coverage rates, the interval lengths differ similarly. In most cases the lengths of the most conservative (Satterthwaite) interval is not much longer than the BM and wild bootstrap intervals. If we compare the

Table 3: Regression on Dummy. Coverage of Nominal 95% Confidence Intervals for γ

		Coverage Probability								Interval Length			
	Cr. Value	Conventional t_{G-1}					BM	Wild	Satt	Jack	BM	Wild	Satt
	St. Error	\hat{v}_1	\hat{v}_2	\hat{v}_3	\hat{v}_4	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5
$G = 6$ Clustered	Normal	0.63	0.66	0.68	0.68	1.00	0.73	1.00	1.00	6.1	3.0	6.5	11
	LogNorm	0.64	0.66	0.66	0.66	1.00	0.71	1.00	1.00	5.9	2.9	6.1	11
	Dummy	0.64	0.66	0.65	0.65	1.00	0.76	1.00	1.00	5.8	3.0	5.9	11
Hetero.	Normal	0.70	0.73	0.74	0.75	1.00	0.79	0.98	1.00	8.4	4.5	15	14
	LogNorm	0.72	0.77	0.79	0.80	1.00	0.81	0.99	1.00	18	12	111	34
	Dummy	0.86	0.91	0.85	0.87	1.00	0.95	1.00	1.00	14	11	42	27
$G = 12$ Clustered	Normal	0.48	0.49	0.51	0.51	1.00	0.52	1.00	1.00	4.4	1.6	5.0	14
	LogNorm	0.47	0.47	0.48	0.48	1.00	0.50	1.00	1.00	4.4	1.6	4.9	15
	Dummy	0.47	0.49	0.48	0.48	1.00	0.77	1.00	1.00	4.4	2.8	4.9	13
Hetero.	Normal	0.58	0.60	0.61	0.62	1.00	0.62	1.00	1.00	6.1	2.6	12	19
	LogNorm	0.62	0.67	0.71	0.72	1.00	0.69	1.00	1.00	15	9	121	48
	Dummy	0.77	0.86	0.79	0.80	1.00	0.97	1.00	1.00	11	14	35	33
$G = 40$ Clustered	Normal	0.27	0.28	0.28	0.28	1.00	0.28	1.00	1.00	3.6	0.8	4.1	19
	LogNorm	0.26	0.26	0.26	0.26	1.00	0.27	1.00	1.00	3.6	0.7	4.1	19
	Dummy	0.29	0.33	0.30	0.31	1.00	0.84	1.00	1.00	3.7	3.4	4.2	15
Hetero.	Normal	0.38	0.38	0.39	0.39	1.00	0.39	1.00	1.00	4.8	1.3	10	25
	LogNorm	0.49	0.54	0.59	0.59	1.00	0.55	1.00	1.00	14	6.3	132	67
	Dummy	0.68	0.82	0.74	0.74	1.00	0.99	1.00	1.00	9.5	24	31	38
$G = 100$ Clustered	Normal	0.17	0.17	0.17	0.17	1.00	0.18	1.00	1.00	3.4	0.5	3.8	20
	LogNorm	0.16	0.16	0.16	0.16	1.00	0.16	1.00	1.00	3.4	0.5	3.8	20
	Dummy	0.22	0.27	0.24	0.24	1.00	0.84	1.00	1.00	3.5	3.8	3.9	15
Hetero.	Normal	0.26	0.26	0.26	0.26	1.00	0.26	1.00	1.00	4.5	0.8	9.0	27
	LogNorm	0.42	0.47	0.52	0.52	1.00	0.47	1.00	1.00	13	4.6	120	72
	Dummy	0.66	0.81	0.72	0.73	1.00	0.99	1.00	1.00	9.1	33	31	39

wild bootstrap with the Satterthwaite intervals, we can see considerable differences, with cases where the wild bootstrap has shorter intervals, and other cases (the leveraged/heteroskedastic designs) where the wild bootstrap produces much longer intervals. We do not have a good explanation for this phenomenon, and believe it is an interesting topic for future research.

The following blocks are for $G = 12$, $G = 40$, and $G = 100$. Qualitatively, the results are similar to the $G = 6$ case. For some designs and methods the coverage probabilities improve slightly as G increases, but in other cases the coverage probabilities worsen. The best coverage is obtained by the Satterthwaite interval, with coverage uniformly exceeding 94%. The BM and wild bootstrap have coverage uniformly exceeding 92%.

We expand the analysis by examining a model with clusterwise noninvertibility. This is $\mathbf{Y}_g = \mathbf{1}_g \alpha + \mathbf{X}_g \beta + \mathbf{D}_g \gamma + \mathbf{e}_g$ with \mathbf{D}_g a dummy indicator for the first cluster. This model is cluster-level

Table 4: Heterogeneous Cluster Sizes. Coverage of Nominal 95% Confidence Intervals for β

		Coverage Probability								Interval Length			
	Cr. Value	Conventional t_{G-1}					BM	Wild	Satt	Jack	BM	Wild	Satt
	St. Error	\hat{v}_1	\hat{v}_2	\hat{v}_3	\hat{v}_4	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5
$G = 6$ Clustered	Normal	0.75	0.87	0.95	0.96	0.97	0.97	0.93	0.98	1.7	1.6	0.9	2.0
	LogNorm	0.79	0.89	0.95	0.96	0.97	0.99	0.95	0.99	1.0	1.2	0.8	1.5
	Dummy	0.77	0.84	0.90	0.91	0.92	0.99	0.95	0.99	0.6	1.7	0.3	1.8
Hetero.	Normal	0.74	0.82	0.89	0.89	0.91	0.93	0.91	0.93	4.1	4.5	6.1	5.1
	LogNorm	0.48	0.63	0.86	0.86	0.88	0.89	0.93	0.91	6.3	7.4	21	9.6
	Dummy	0.69	0.76	0.83	0.83	0.85	0.95	0.95	0.95	7.0	20	49	21
$G = 12$ Clustered	Normal	0.67	0.83	0.95	0.95	0.96	0.96	0.93	0.98	1.1	1.1	0.5	1.5
	LogNorm	0.75	0.86	0.95	0.95	0.96	0.99	0.95	0.99	0.5	0.7	0.2	0.8
	Dummy	0.70	0.80	0.88	0.89	0.90	0.99	0.95	0.99	0.5	1.8	0.3	1.9
Hetero.	Normal	0.68	0.79	0.87	0.87	0.89	0.92	0.91	0.93	2.9	3.2	4.0	3.9
	LogNorm	0.50	0.63	0.85	0.85	0.86	0.90	0.93	0.93	4.7	6.2	15	8.4
	Dummy	0.64	0.73	0.81	0.82	0.82	0.95	0.95	0.95	6.0	21	49	21
$G = 40$ Clustered	Normal	0.54	0.78	0.96	0.96	0.96	0.96	0.94	0.99	0.8	0.8	0.4	1.1
	LogNorm	0.73	0.86	0.95	0.96	0.96	0.99	0.95	0.99	0.2	0.3	0.0	0.4
	Dummy	0.64	0.77	0.87	0.87	0.88	0.99	0.95	0.99	0.5	1.8	0.3	1.9
Hetero.	Normal	0.62	0.74	0.86	0.86	0.86	0.89	0.91	0.93	2.4	2.5	3.1	3.2
	LogNorm	0.55	0.68	0.85	0.86	0.86	0.92	0.93	0.94	3.5	4.7	9.0	6.6
	Dummy	0.60	0.71	0.81	0.81	0.81	0.95	0.95	0.95	5.6	22	49	22
$G = 100$ Clustered	Normal	0.46	0.76	0.97	0.97	0.97	0.97	0.95	0.99	0.8	0.7	0.3	1.0
	LogNorm	0.71	0.85	0.96	0.96	0.96	0.98	0.94	0.99	0.2	0.2	0.0	0.3
	Dummy	0.64	0.76	0.87	0.87	0.87	0.99	0.94	0.99	0.5	1.8	0.3	1.9
Hetero.	Normal	0.57	0.71	0.84	0.84	0.84	0.87	0.91	0.92	2.2	2.3	2.8	3.0
	LogNorm	0.57	0.69	0.86	0.86	0.86	0.92	0.93	0.94	3.0	3.8	6.4	5.4
	Dummy	0.59	0.70	0.80	0.80	0.80	0.95	0.95	0.95	5.4	22	49	22

treatment with a single treated cluster. The regression is clusterwise noninvertible as the least squares estimator is undefined when the first cluster is omitted. In this model we examine confidence intervals for both β and γ . For this model we consider the same six designs as in the baseline model for the distributions of the regressors and error variances.

Table 2 presents the results for confidence intervals for β . In general, the results are similar to those of Table 1 but with one important difference. In Table 2 we see a meaningful divergence in performance between the intervals based on conventional jackknife standard errors \hat{v}_3 and \hat{v}_4 and those based on our recommended jackknife standard error \hat{v}_5 . In Table 1, these three methods were nearly identical; in Table 2 we can see that the two conventional intervals exhibit substantial undercoverage in many designs. This difference is due to the treatment of noninvertible clusters. This strongly supports our recommendation for \hat{v}_5 over \hat{v}_3 and \hat{v}_4 .

Table 5: LogNormal Errors. Coverage of Nominal 95% Confidence Intervals for β

		Coverage Probability							Interval Length				
	Cr. Value	Conventional t_{G-1}					BM	Wild	Satt	Jack	BM	Wild	Satt
	St. Error	\hat{v}_1	\hat{v}_2	\hat{v}_3	\hat{v}_4	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5	\hat{v}_5	\hat{v}_2	\hat{v}_5	\hat{v}_5
$G = 6$ Clustered	Normal	0.93	0.94	0.96	0.96	0.97	0.96	0.93	0.97	1.2	1.1	0.5	1.2
	LogNorm	0.86	0.92	0.96	0.96	0.97	0.99	0.96	0.99	0.6	0.7	0.3	1.0
	Dummy	0.85	0.89	0.92	0.92	0.94	1.00	0.97	1.00	0.6	1.7	0.5	1.8
Hetero.	Normal	0.84	0.86	0.89	0.89	0.91	0.89	0.86	0.90	2.6	2.4	2.7	2.7
	LogNorm	0.55	0.68	0.84	0.85	0.86	0.88	0.91	0.91	4.9	5.3	16	8
	Dummy	0.65	0.72	0.79	0.80	0.81	0.94	0.93	0.94	6.0	15	50	17
$G = 12$ Clustered	Normal	0.94	0.95	0.96	0.96	0.97	0.96	0.94	0.97	0.7	0.7	0.2	0.7
	LogNorm	0.86	0.91	0.95	0.95	0.96	0.99	0.96	0.99	0.3	0.3	0.0	0.4
	Dummy	0.74	0.80	0.86	0.86	0.87	1.00	0.96	1.00	0.5	1.8	0.5	1.0
Hetero.	Normal	0.86	0.87	0.88	0.88	0.89	0.89	0.87	0.89	1.6	1.6	1.1	1.6
	LogNorm	0.60	0.70	0.83	0.83	0.84	0.89	0.90	0.92	3.5	4.4	10	6.2
	Dummy	0.59	0.68	0.77	0.78	0.78	0.94	0.93	0.94	5.1	17	50	18
$G = 40$ Clustered	Normal	0.95	0.96	0.96	0.96	0.96	0.96	0.95	0.96	0.4	0.3	0.0	0.3
	LogNorm	0.88	0.91	0.94	0.94	0.94	0.98	0.94	0.98	0.1	0.1	0.0	0.1
	Dummy	0.62	0.71	0.79	0.80	0.80	1.00	0.93	0.99	0.5	1.9	0.5	1.9
Hetero.	Normal	0.89	0.90	0.90	0.90	0.91	0.90	0.90	0.90	0.9	0.9	0.3	0.9
	LogNorm	0.68	0.76	0.85	0.85	0.85	0.90	0.90	0.93	2.5	3.1	5.4	4.1
	Dummy	0.56	0.66	0.76	0.76	0.76	0.94	0.94	0.94	4.7	18	53	19
$G = 100$ Clustered	Normal	0.96	0.96	0.96	0.96	0.96	0.96	0.95	0.96	0.2	0.2	0.0	0.2
	LogNorm	0.90	0.92	0.94	0.94	0.94	0.97	0.93	0.97	0.0	0.1	0.0	0.1
	Dummy	0.58	0.68	0.77	0.77	0.77	0.99	0.92	0.98	0.5	1.9	0.5	1.9
Hetero.	Normal	0.92	0.92	0.92	0.92	0.92	0.92	0.92	0.92	0.6	0.6	0.1	0.6
	LogNorm	0.75	0.81	0.87	0.87	0.88	0.91	0.91	0.93	2.0	2.3	3.0	2.9
	Dummy	0.55	0.65	0.76	0.76	0.76	0.94	0.94	0.94	4.6	19	53	19

Next, we examine Table 3, which presents the results for the dummy variable coefficient γ in the regression model $Y_g = \mathbf{1}_g\alpha + \mathbf{X}_g\beta + \mathbf{D}_g\gamma + \mathbf{e}_g$. This is a treacherous context, as we are essentially making inference for a coefficient based on a single cluster. The results in Table 3 reveal that the conventional inference methods fail, and worsen as the sample size increases. The Bell-McCaffrey adjusted interval similarly fails. The exceptions (which generally produce coverage rates of 100%) are the three intervals (conventional, wild bootstrap, and Satterthwaite) using the jackknife \hat{v}_5 standard error. The reason why the intervals based on the jackknife standard error \hat{v}_5 have coverage rates of 100% is because in this context \hat{v}_5 tends to approximately equal the coefficient estimate $\hat{\gamma}$, so the confidence interval always covers the true value of 0. Coverage is conservative, but at least there is no tendency towards false significance. We can also see that the average lengths of these three feasible intervals are different from one another, with hetero-

geneous rankings. Overall, the message of Table 3 is that in the extreme context of inference on a dummy indicator for a single treated cluster, the Bell-McCaffrey interval is not robust, the wild bootstrap and Satterthwaite intervals are robust, but can be highly conservative.

We next explore the impact of heterogeneous cluster sizes. In Tables 1-3, cluster sizes were all set at $n_g = 10$. In our next experiment we set two clusters to have $2 + 4G$ observations, with the remainder with 2 observations each. This means that the total number of observations is the same as before, but the samples are dominated by two large clusters. Otherwise, the experiment follows the baseline model. The results are presented in Table 4. Most inference methods are adversely affected, with coverage rates meaningfully lower in Table 4 relative to Table 1. Generally good coverage is obtained by the Satterthwaite confidence interval. Interval lengths follow similar patterns as in Table 1.

Our inference theory is developed under the assumption of normally distributed errors, which raises the question if the coverage rates are sensitive to departures to normality. To explore this question, we repeat the analysis of the baseline model, but with the errors drawn from a LogNormal distribution, centered and scaled to have a mean of zero and variance of one. We sample the regression errors from this skewed distribution, with the same cluster covariances as in the baseline model. The results are displayed in Table 5. Comparing Tables 1 and 5, we can see that the coverage rates of most methods worsen in Table 5. The worst case coverage of the Bell-McCaffrey interval is 88%, of the wild bootstrap is 86%, and of the Satterthwaite interval is 89%. These coverage rates improve, however, as G increases.

It is instructive to contrast the coverage rates of Tables 1-5 with the lower bound of Theorem 3. Recall, the latter states under normal errors and inference on the slope β , the coverage rates of a nominal 95% interval using the jackknife \hat{v}_5 can never be lower than 70%. Indeed, the coverage rates for the jackknife interval uniformly exceed this bound; in fact, the coverage rates are never lower than 80% under normal errors. In contrast, the intervals using \hat{v}_1 have coverage rates as low as 46% for β , and those using \hat{v}_2 have coverage rates as low as 63%.

In unreported simulations, we explored the impact of inclusion of a large number of extraneous regressors. Cattaneo, Jansson and Newey (2018) showed that the downward bias of conventional variance estimators is increasing in the number of extraneous regressors, while jackknife variance estimators can be upward biased. While our theory allows for an arbitrary number of regressors under normal errors, it is unclear how the results extend to the case of nonnormal errors. Our simulation results are qualitatively similar to those in the reported tables. We also explored the performance of alternative inference methods, including the pairs clustered bootstrap and alternative wild bootstrap methods. None of these alternatives had good coverage accuracy, and so are not reported. We also explored the performance of the Satterthwaite interval calculated with alternative reference covariance matrix models (for ex-

ample, $\Sigma_g = \mathbf{X}_g \mathbf{X}_g'$). These produced qualitatively similar results, but more conservative, than the reported Satterthwaite interval.

In summary, we are able to draw the following conclusions from the simulation evidence. First, the standard error which produces intervals with the best coverage rates is \hat{v}_5 . Second, the simple interval $\hat{\beta} \pm t_{G-1}^{1-\alpha/2} \hat{v}_5$ has good coverage rates in many contexts, but can undercover in extreme designs. Third, the adjusted interval $\hat{\beta} \pm t_K^{1-\alpha/2} \hat{v}_5 / a$ has excellent (but sometimes conservative) coverage in all contexts examined. Fourth, the restricted wild bootstrap with jackknife standard errors also is an excellent option for inference. Fifth, issues such as clusterwise invertibility should not be handled by *ad hoc* computational implementations, but rather by methods justified by theoretical insight. In particular, the jackknife should be implemented without the discarding of iterations with noninvertible design matrices.

9 Empirical Illustration

We illustrate the applicability of the methods with an empirical example. We follow Canay, Santos, and Shaikh (2021) by revisiting an application by Meng, Qian, and Yared (2015) into the causes of the Chinese Great Famine between 1958 and 1960. Their regressions (Table 2 of Meng-Qian-Yared) take the form $Y = Z_1 \beta_1 + Z_2 \beta_2 + W' \gamma + e$ for $G = 19$ provinces between 1953 and 1982, where Y equals the log of deaths in the province, Z_1 equals the log of predicted grain production, Z_2 equals the product of Z_1 and an indicator for a famine year, and W are other controls. The focus is on the coefficient sum $\beta_1 + \beta_2$. The authors report six specifications which vary the sample period (1953-1982 vs 1953-1965), the provinces (19 vs 23), and replacing predicted with reported grain production. Canay, Santos, and Shaikh (2021) use this application to illustrate hypothesis testing using the cluster wild bootstrap. In contrast, we are interested in standard error calculation, confidence interval construction, and hypothesis testing. Following these authors, we cluster by province.

We estimate the same six regression specifications as Meng, Qian, and Yared (2015) and focus on the coefficient sum $\beta_1 + \beta_2$. In Table 6 we report the least squares estimates $\hat{\beta}_1 + \hat{\beta}_2$ plus three standard errors: $CRVE_1$, $CRVE_2$, and jackknife. What you can see from the table is that in some of the specifications there are considerable differences between the three standard errors, and in particular between the jackknife and the other two. The discrepancies between the $CRVE_1$ and jackknife standard errors range from 10% (in specification #1) to 66% (in specification #3). These are large and substantial differences.

We next construct 95% adjusted confidence intervals for the coefficient sum $\beta_1 + \beta_2$. We start by calculating the adjustment coefficients K and a for each of the six specifications, and report these coefficients in Table 6. The values for the adjusted degrees of freedom K range between 4

and 6, which are all small. The values for the scale adjustment a range between 1.17 and 1.26. Together, these are used to construct the confidence intervals, which are reported in the table. Take, for example, specification #1, where $\hat{\beta}_1 + \hat{\beta}_2 = 0.141$, $\hat{v}_5 = 0.066$, $K = 4.18$, and $a = 1.21$. The 95% critical value from the t distribution with $K = 4.18$ degrees of freedom is 2.73. The 95% confidence interval is therefore $0.141 \pm 2.73 \times 0.066/1.21 = [-0.01, 0.29]$. The confidence intervals are wide, indicating uncertainty about the value of the coefficient sum. The intervals do not vary greatly across the six specifications, indicating that the result is reasonably robust to the specification.

Table 6: China's Great Famine, 1959-1961

	Dependent variable: log deaths in year $t + 1$					
	Constructed grain production				Reported grain production	
	19 provinces		23 provinces		19 provinces	
	1953-1982	1953-1965	1953-1982	1953-1965	1953-1982	1953-1965
	(1)	(2)	(3)	(4)	(5)	(6)
$\hat{\beta}_1 + \hat{\beta}_2$	0.141	0.098	0.115	0.094	0.113	0.089
CRVE ₁	(0.060)	(0.053)	(0.037)	(0.037)	(0.063)	(0.059)
CRVE ₂	[0.061]	[0.056]	[0.039]	[0.037]	[0.068]	[0.063]
Jackknife	<0.066>	<0.066>	<0.061>	<0.052>	<0.079>	<0.073>
K	4.18	3.95	5.34	5.03	5.18	5.97
a	1.21	1.26	1.21	1.23	1.18	1.17
Interval	[-0.01, 0.29]	[-0.05, 0.24]	[-0.01, 0.24]	[-0.01, 0.20]	[-0.06, 0.28]	[-0.06, 0.24]
p-value	0.058	0.135	0.069	0.076	0.151	0.200
CSS p-value	0.061	0.072	0.029	0.030	0.141	0.171

Notes: All regressions include log total population, log urban population, and year fixed effects. Standard errors are clustered by province. The CSS p-value is that reported in Canay, Santos, and Shaikh (2021).

We also construct and report adjusted p-values for t -tests of the hypothesis $\beta_1 + \beta_2 = 0$. The t -statistic using the jackknife standard error is $0.141/0.066 = 2.16$. The adjusted p-value is $1 - F(1.21^2 \times 2.16^2; 1, 4.18) = 0.058$. None of the six p-values are significant at the 5% level. This contrasts with the p-values reported by Meng, Qian, and Yared (2015), which were all statistically significant, some greatly so. Several of our p-values are similar to those calculated by the wild cluster bootstrap as reported by Canay, Santos, and Shaikh (2021), which we report in the final row of the table. For example, for the baseline specification #1, our p-value of 0.058 is nearly identical to their wild studentized p-value of 0.061. However, in contrast to bootstrap and randomization methods, our confidence intervals and p-values can be calculated by Stata's regress command without specialized programs.

10 Conclusion

Heteroskedasticity-consistent and cluster-robust standard errors are routinely reported in applied econometric practice. It is prudent to coalesce on simple yet well-behaved methods which produce reliable inference across reasonable estimation settings. It is our contention that jackknife variance estimators are superior to conventional (EHW and CRVE₁) estimators, based on our analysis of worst-case downward bias and confidence interval coverage. Consequently, we recommend jackknife variance estimation for empirical practice. We also recommend that conventional statistical software calculate default confidence intervals and p-values using the Satterthwaite adjustment, as it is computationally convenient, more conservative than student t inference, and more reliable in finite samples.

A Appendix: Technical Proofs

Let $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denote the smallest and largest eigenvalue of a Hermitian matrix \mathbf{A} , $\|\mathbf{A}\| = (\lambda_{\max}(\mathbf{A}'\mathbf{A}))^{1/2}$ denote the spectral matrix norm, and $\|\mathbf{A}\|_F = (\text{tr}[\mathbf{A}'\mathbf{A}])^{1/2}$ denote the Frobenius matrix norm.

Proof of Theorem 1: The estimator (11) with any generalized inverse is a least-squares minimizer and solves the first order condition

$$(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)\widehat{\beta}_{-g} = (\mathbf{X}'\mathbf{Y} - \mathbf{X}'_g\mathbf{Y}_g).$$

Pre-multiplying by $(\mathbf{X}'\mathbf{X})^{-1}$, rearranging, and using (12), we obtain

$$\begin{aligned}\widehat{\beta}_{-g} &= (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g\mathbf{Y}_g + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g\mathbf{X}_g\widehat{\beta}_{-g} \\ &= \widehat{\beta} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g(\mathbf{Y}_g - \mathbf{X}_g\widehat{\beta}_{-g}) \\ &= \widehat{\beta} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g\widehat{\mathbf{e}}_{-g},\end{aligned}$$

which is (13) as claimed. By definition (12) and model (1), the prediction errors equal $\widehat{\mathbf{e}}_{-g} = \mathbf{e}_g - \mathbf{X}_g(\widehat{\beta}_{-g} - \beta)$. Squaring and expanding,

$$\widehat{\mathbf{e}}_{-g}\widehat{\mathbf{e}}'_{-g} = \mathbf{e}_g\mathbf{e}'_g - \mathbf{e}_g(\widehat{\beta}_{-g} - \beta)' \mathbf{X}'_g - \mathbf{X}_g(\widehat{\beta}_{-g} - \beta)\mathbf{e}'_g + \mathbf{X}_g(\widehat{\beta}_{-g} - \beta)(\widehat{\beta}_{-g} - \beta)' \mathbf{X}'_g \quad (40)$$

$$\geq \mathbf{e}_g\mathbf{e}'_g - \mathbf{e}_g(\widehat{\beta}_{-g} - \beta)' \mathbf{X}'_g - \mathbf{X}_g(\widehat{\beta}_{-g} - \beta)\mathbf{e}'_g. \quad (41)$$

The inequality holds because the final term in (40) is positive semi-definite.

The first term in (41) has expectation Σ_g . Observe that \mathbf{e}_g is independent of $\widehat{\beta}_{-g} - \beta$ and

mean zero, so the expectation of the second and third terms in (41) equals zero. We deduce that

$$\mathbb{E} \left[\widehat{\boldsymbol{e}}_{-g} \widehat{\boldsymbol{e}}'_{-g} \right] \geq \mathbb{E} \left[\boldsymbol{e}_g \boldsymbol{e}'_g \right] = \boldsymbol{\Sigma}_g. \quad (42)$$

Using expression (14) and inequality (42),

$$\begin{aligned} \mathbb{E} [\widehat{V}_5] &= (\boldsymbol{X}'\boldsymbol{X})^{-1} \left(\sum_{g=1}^G \boldsymbol{X}'_g \mathbb{E} \left[\widehat{\boldsymbol{e}}_{-g} \widehat{\boldsymbol{e}}'_{-g} \right] \boldsymbol{X}_g \right) (\boldsymbol{X}'\boldsymbol{X})^{-1} \\ &\geq (\boldsymbol{X}'\boldsymbol{X})^{-1} \left(\sum_{g=1}^G \boldsymbol{X}'_g \boldsymbol{\Sigma}_g \boldsymbol{X}_g \right) (\boldsymbol{X}'\boldsymbol{X})^{-1} = \boldsymbol{V}. \end{aligned}$$

This is the stated result. \blacksquare

Proof of Theorem 2, equation (15): Set $\boldsymbol{Z}_g = (\boldsymbol{X}'\boldsymbol{X})^{-1} \boldsymbol{X}_g R$. Without loss of generality, normalize $\boldsymbol{X}'\boldsymbol{X} = \boldsymbol{I}_k$ and $R'R = 1$. Set $c_g = \boldsymbol{Z}'_g \boldsymbol{Z}_g$ and observe that $\sum_{g=1}^G c_g = 1$. Assume that $\boldsymbol{\Sigma}_g = \boldsymbol{I}_{n_g}$. Under these conditions, $\nu^2 = 1$. In the model class \mathcal{F}^* the value of c_1 is freely varying in $(0,1)$. The CRVE₁ estimator is

$$\widehat{\nu}_1^2 = d \sum_{g=1}^G \boldsymbol{Z}'_g \widehat{\boldsymbol{e}}_g \widehat{\boldsymbol{e}}'_g \boldsymbol{Z}_g \quad (43)$$

where $d = G(n-1)/(G-1)(n-k)$. We can calculate that under these conditions

$$\mathbb{E} \left[\widehat{\boldsymbol{e}}_g \widehat{\boldsymbol{e}}'_g \right] = \boldsymbol{M}_g. \quad (44)$$

Observe that since $\boldsymbol{I}_k - RR'$ is a projection matrix, it is positive semi-definite, and then

$$\boldsymbol{X}_g \boldsymbol{X}'_g = \boldsymbol{Z}_g \boldsymbol{Z}'_g + \boldsymbol{X}_g (\boldsymbol{I}_k - RR') \boldsymbol{X}'_g \geq \boldsymbol{Z}_g \boldsymbol{Z}'_g.$$

It follows that

$$\boldsymbol{M}_g = \boldsymbol{I}_{n_g} - \boldsymbol{X}_g \boldsymbol{X}'_g \leq \boldsymbol{I}_{n_g} - \boldsymbol{Z}_g \boldsymbol{Z}'_g. \quad (45)$$

Taking expectations of (43) using (44), (45), $\boldsymbol{Z}'_g \boldsymbol{Z}_g = c_g$, and $\sum_{g=1}^G c_g = 1$, we find that

$$\begin{aligned} \frac{\mathbb{E} [\widehat{\nu}_1^2]}{\nu^2} &\leq d \sum_{g=1}^G \boldsymbol{Z}'_g \left(\boldsymbol{I}_{n_g} - \boldsymbol{Z}_g \boldsymbol{Z}'_g \right) \boldsymbol{Z}_g \\ &= d \left(1 - \sum_{g=1}^G c_g^2 \right) \\ &\leq d (1 - c_1^2). \end{aligned} \quad (46)$$

The assumptions are a special case of the model class \mathcal{F}^* . Inequality (46) implies that the

left side of (15) is weakly smaller than the infimum of the right-hand side of (46) over c_1 . Thus

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\widehat{v}_1^2]}{v^2} \leq d \inf_{0 < c_1 < 1} (1 - c_1^2) = 0. \quad (47)$$

The rightmost equality in (47) is attained as $c_1 \rightarrow 1$. Since the left side of (47) is non-negative, the equation hold as an equality. This verifies (15). ■

Proof of Theorem 2, equation (16): We use the same normalizations and notation as in the proof of (15). We replace the assumption on the covariance matrices with the assumption that $\boldsymbol{\Sigma}_1 = \mathbf{I}_{n_1}$ and $\boldsymbol{\Sigma}_g = 0$ for $g > 1$, which is extreme heteroskedasticity. Under these conditions, you can calculate that $\mathbf{V} = \mathbf{X}'_1 \mathbf{X}_1$ and $v^2 = c_1$. It will be convenient to define

$$\delta = \max_{2 \leq g \leq G} \|\mathbf{X}_g\|^2 \geq \frac{1 - c_1}{G - 1}. \quad (48)$$

The inequality in (48) holds because

$$\delta \geq \sum_{g=2}^G \lambda_{\max}(\mathbf{X}'_g \mathbf{X}_g) \geq \lambda_{\max}(\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1) = 1 - \lambda_{\min}(\mathbf{X}'_1 \mathbf{X}_1) \geq 1 - c_1.$$

The CRVE₂ estimator for v^2 can be written as

$$\widehat{v}_2^2 = \sum_{g=1}^G \mathbf{Z}'_g \mathbf{M}_g^{-1/2} \widehat{\boldsymbol{e}}_g \widehat{\boldsymbol{e}}'_g \mathbf{M}_g^{-1/2} \mathbf{Z}_g. \quad (49)$$

We can calculate that for $g = 1$,

$$\mathbb{E}[\widehat{\boldsymbol{e}}_1 \widehat{\boldsymbol{e}}'_1] = \mathbf{I}_{n_1} - 2\mathbf{X}_1 \mathbf{X}'_1 + \mathbf{X}_1 \mathbf{X}'_1 \mathbf{X}_1 \mathbf{X}'_1 = \mathbf{M}_1 \mathbf{M}_1, \quad (50)$$

and for $g > 1$,

$$\mathbb{E}[\widehat{\boldsymbol{e}}_g \widehat{\boldsymbol{e}}'_g] = \mathbf{X}_g \mathbf{X}'_1 \mathbf{X}_1 \mathbf{X}'_g. \quad (51)$$

Taking expectations of (49) and using (50)-(51)

$$\begin{aligned} \mathbb{E}[\widehat{v}_2^2] &= \sum_{g=1}^G \mathbf{Z}'_g \mathbf{M}_g^{-1/2} \mathbb{E}[\widehat{\boldsymbol{e}}_g \widehat{\boldsymbol{e}}'_g] \mathbf{M}_g^{-1/2} \mathbf{Z}_g \\ &= \mathbf{Z}'_1 \mathbf{M}_1^{-1/2} \mathbf{M}_1 \mathbf{M}_1 \mathbf{M}_1^{-1/2} \mathbf{Z}_1 + \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1/2} \mathbf{X}_g \mathbf{X}'_1 \mathbf{X}_1 \mathbf{X}'_g \mathbf{M}_g^{-1/2} \mathbf{Z}_g. \end{aligned} \quad (52)$$

The first term on the right side of (52) equals

$$\mathbf{Z}'_1 \mathbf{M}_1 \mathbf{Z}_1 \leq \mathbf{Z}'_1 (\mathbf{I}_{n_1} - \mathbf{Z}_1 \mathbf{Z}'_1) \mathbf{Z}_1 = c_1 - c_1^2, \quad (53)$$

where the inequality is (45). We observe that for $g > 1$

$$\|\mathbf{M}_g^{-1}\| = \frac{1}{\lambda_{\min}(\mathbf{M}_g)} = \frac{1}{1 - \lambda_{\max}(\mathbf{X}'_g \mathbf{X}_g)} = \frac{1}{1 - \|\mathbf{X}_g\|^2} \leq \frac{1}{1 - \delta}. \quad (54)$$

Using the fact $\mathbf{X}'_1 \mathbf{X}_1 \leq \mathbf{X}' \mathbf{X} = \mathbf{I}_k$, the quadratic inequality, (48), and (54), the second term on the right side of (52) is smaller than

$$\sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1/2} \mathbf{X}_g \mathbf{X}'_g \mathbf{M}_g^{-1/2} \mathbf{Z}_g \leq \frac{\delta}{(1 - \delta)} \sum_{g=2}^G \mathbf{Z}'_g \mathbf{Z}_g = \frac{\delta(1 - c_1)}{(1 - \delta)}. \quad (55)$$

Combining (52), (53), (55), and $v^2 = c_1$ we find that

$$\frac{\mathbb{E}[\widehat{v}_2^2]}{v^2} \leq \frac{c_1 - c_1^2}{c_1} + \frac{\delta(1 - c_1)}{c_1(1 - \delta)}. \quad (56)$$

The assumptions we have made are a special case of the model class \mathcal{F}^* . Therefore, the left side of (16) is weakly smaller than the infimum of (56) over (c_1, δ) . Hence

$$\begin{aligned} \inf_{(\mathbf{X}, \Sigma) \in \mathcal{F}^*} \frac{\mathbb{E}[\widehat{v}_2^2]}{v^2} &\leq \inf_{c_1, \delta} \left(1 - c_1 + \frac{\delta(1 - c_1)}{c_1(1 - \delta)} \right) \\ &= \inf_{c_1} \left(1 - c_1 + \frac{(1 - c_1)^2}{c_1(G - 2 + c_1)} \right) = 0. \end{aligned} \quad (57)$$

The first equality in the second line is obtained by setting δ at its lower bound $(1 - c_1)/(G - 1)$ from (48). The final equality is attained as $c_1 \rightarrow 1$. Equation (57) implies (16), as claimed. \blacksquare

Proof of Theorem 2, equation (17): To show (17) we calculate an upper bound in the example used in the proof of (16). We adopt the assumptions made therein. The model class \mathcal{F}^* imposes clusterwise invertibility. As shown by equation (19) of MacKinnon, Nielsen, and Webb (2023b), clusterwise invertibility implies

$$R'(\widehat{\beta} - \widehat{\beta}_{-g}) = R'(\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}'_g \mathbf{M}_g^{-1} \widehat{\mathbf{e}}_g = \mathbf{Z}'_g \mathbf{M}_g^{-1} \widehat{\mathbf{e}}_g. \quad (58)$$

Consequently,

$$\begin{aligned}\hat{v}_3^2 &= \left(\frac{G-1}{G}\right) \sum_{g=1}^G R'(\hat{\beta}_{-g} - \hat{\beta})(\hat{\beta}_{-g} - \hat{\beta})' R - (G-1) R'(\hat{\beta} - \bar{\beta})(\hat{\beta} - \bar{\beta})' R \\ &= \left(\frac{G-1}{G}\right)^2 \mathbf{Z}'_1 \mathbf{M}_1^{-1} \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1' \mathbf{M}_1^{-1} \mathbf{Z}_1\end{aligned}\quad (59)$$

$$+ \left(\frac{G-1}{G}\right) \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \hat{\mathbf{e}}_g \hat{\mathbf{e}}_g' \mathbf{M}_g^{-1} \mathbf{Z}_g \quad (60)$$

$$- 2 \left(\frac{G-1}{G^2}\right) \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \hat{\mathbf{e}}_g \hat{\mathbf{e}}_1' \mathbf{M}_1^{-1} \mathbf{Z}_1 \quad (61)$$

$$- \left(\frac{G-1}{G^2}\right) \left(\sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \hat{\mathbf{e}}_g \right) \left(\sum_{g=2}^G \hat{\mathbf{e}}_g' \mathbf{M}_g^{-1} \mathbf{Z}_g \right). \quad (62)$$

Using (50), the expectation of (59) equals $((G-1)/G)^2$ times

$$\mathbf{Z}'_1 \mathbf{M}_1^{-1} \mathbb{E}[\hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1'] \mathbf{M}_1^{-1} \mathbf{Z}_1 = \mathbf{Z}'_1 \mathbf{M}_1^{-1} \mathbf{M}_1 \mathbf{M}_1 \mathbf{M}_1^{-1} \mathbf{Z}_1 = \mathbf{Z}'_1 \mathbf{Z}_1 = c_1. \quad (63)$$

Using (51) and (54) the expectation of (60) equals $((G-1)/G)$ times

$$\sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \mathbb{E}[\hat{\mathbf{e}}_g \hat{\mathbf{e}}_g'] \mathbf{M}_g^{-1} \mathbf{Z}_g = \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \mathbf{X}_g \mathbf{X}'_1 \mathbf{X}_1 \mathbf{X}'_g \mathbf{M}_g^{-1} \mathbf{Z}_g \leq \frac{(1-c_1)\delta}{(1-\delta)^2}, \quad (64)$$

where the inequality follows by similar same steps as for (55).

We calculate that for $g \neq 1$

$$\mathbb{E}[\hat{\mathbf{e}}_g \hat{\mathbf{e}}_1'] = -\mathbf{X}_g (\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1) \mathbf{X}'_1. \quad (65)$$

Using the Woodbury identity,

$$\begin{aligned}\mathbf{X}'_1 \mathbf{M}_1^{-1} \mathbf{X}_1 &= \mathbf{X}'_1 \mathbf{X}_1 + \mathbf{X}'_1 \mathbf{X}_1 (\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1)^{-1} \mathbf{X}'_1 \mathbf{X}_1 \\ &= -\mathbf{I}_k + (\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1)^{-1} \leq (\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1)^{-1}.\end{aligned}\quad (66)$$

Combining (65), (66), and (54), the expectation of (61) equals $2(G-1)/G^2$ times

$$\begin{aligned}R' \sum_{g=2}^G \mathbf{X}'_g \mathbf{M}_g^{-1} \mathbf{X}_g (\mathbf{I}_k - \mathbf{X}'_1 \mathbf{X}_1) \mathbf{X}'_1 \mathbf{M}_1^{-1} \mathbf{X}_1 R &\leq \sum_{g=2}^G R' \mathbf{X}'_g \mathbf{M}_g^{-1} \mathbf{X}_g R \\ &\leq \frac{1}{(1-\delta)} \sum_{g=2}^G \mathbf{Z}'_g \mathbf{Z}_g = \frac{(1-c_1)}{(1-\delta)}.\end{aligned}\quad (67)$$

Together, (59)-(64), (67), $v^2 = c_1$, and the fact that (62) is non-positive, imply

$$\frac{\mathbb{E}[\hat{v}_3^2]}{v^2} \leq \left(\frac{G-1}{G}\right)^2 + \left(\frac{G-1}{G}\right) \frac{(1-c_1)\delta}{c_1(1-\delta)^2} + 2\left(\frac{G-1}{G^2}\right) \frac{(1-c_1)}{c_1(1-\delta)}. \quad (68)$$

As in the proof of (16), the left side of (17) is weakly smaller than the infimum of (68) over (c_1, δ) . By similar steps as in (57),

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}^*} \frac{\mathbb{E}[\hat{v}_3^2]}{v^2} \leq \inf_{c_1} \left(\left(\frac{G-1}{G}\right)^2 + \frac{(G-1)^2}{G} \frac{(1-c_1)^2}{c_1(G-2+c_1)^2} + 2\left(\frac{G-1}{G}\right) \frac{(1-c_1)}{c_1(G-2+c_1)} \right) = \left(\frac{G-1}{G}\right)^2.$$

The final equality is attained as $c_1 \rightarrow 1$. This is (17), as claimed. \blacksquare

Proof of Theorem 2, equation (18): Since the model is clusterwise invertible, $\hat{v}_4^2 = \left(\frac{G-1}{G}\right) \hat{v}_5^2$. The result follows from (21), which we establish below. \blacksquare

Proof of Theorem 2, equations (19) and (20): We calculate an upper bound for the left side of (19)-(20) assuming a single cluster ($g = 1$) is noninvertible and the remaining clusters ($g > 1$) are invertible. The noninvertibility is due to the inclusion of a fixed effect dummy variable (only for this single cluster), and it is assumed that R does not load on this variable. Thus Assumption 3 holds. Otherwise, we adopt the assumptions used in the proof of (15), the notation (48), and inequalities (48)-(54).

Because cluster $g = 1$ is noninvertible, it is discarded from the calculation of \hat{v}_3^2 and \hat{v}_4^2 . Therefore the latter equals

$$\begin{aligned} \hat{v}_4^2 &= \left(\frac{G-2}{G-1}\right) \sum_{g=2}^G R' (\hat{\beta}_{-g} - \hat{\beta}) (\hat{\beta}_{-g} - \hat{\beta})' R \\ &= \left(\frac{G-2}{G-1}\right) \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \hat{\mathbf{e}}_g \hat{\mathbf{e}}'_g \mathbf{M}_g^{-1} \mathbf{Z}_g \\ &\leq \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \hat{\mathbf{e}}_g \hat{\mathbf{e}}'_g \mathbf{M}_g^{-1} \mathbf{Z}_g, \end{aligned}$$

the equality using (58) (which holds for $g \geq 2$ because these are clusterwise invertible), and the inequality is $(G-2)/(G-1) \leq 1$. Taking expectations, using $v^2 = 1$, (44), and then (54)

$$\frac{\mathbb{E}[\hat{v}_4^2]}{v^2} \leq \sum_{g=2}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \mathbf{Z}_g \leq \frac{1}{1-\delta} \sum_{g=2}^G \mathbf{Z}'_g \mathbf{Z}_g = \frac{1-c_1}{1-\delta}. \quad (69)$$

The assumptions we have made are a special case of the model class \mathcal{F} . Therefore, the left side

of (18) is weakly smaller than the infimum of (69) over (c_1, δ) . Hence

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\widehat{v}_4^2]}{v^2} \leq \inf_{c_1, \delta} \left(\frac{1 - c_1}{1 - \delta} \right) = \inf_{c_1} \frac{(1 - c_1)(G - 1)}{G - 2 + c_1} = 0.$$

This implies (18) as claimed. Since $\widehat{v}_3^2 \leq \widehat{v}_4^2$ we find that (19) holds as well. \blacksquare

Proof of Theorem 2, equation (21): Theorem 1 implies

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\widehat{v}_5^2]}{v^2} = \inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{R' \mathbb{E}[\widehat{\mathbf{V}}_5] R}{R' \mathbf{V} R} \geq \inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{R' \mathbf{V} R}{R' \mathbf{V} R} = 1. \quad (70)$$

To show this is a strict equality we calculate an upper bound for the left side of (70) in the context of the example from the proof of (17). Using the calculations therein, including (63) and (64)

$$\widehat{v}_5^2 = \sum_{g=1}^G \mathbf{Z}'_g \mathbf{M}_g^{-1} \widehat{\mathbf{e}}_g \widehat{\mathbf{e}}'_g \mathbf{M}_g^{-1} \mathbf{Z}_g.$$

Using (63) and (64), its expectation satisfies $\mathbb{E}[\widehat{v}_5^2]/v^2 \leq 1 + (1 - c_1)\delta/[(1 - \delta)^2 c_1]$. This implies

$$\inf_{(\mathbf{X}, \boldsymbol{\Sigma}) \in \mathcal{F}} \frac{\mathbb{E}[\widehat{v}_5^2]}{v^2} \leq \inf_{c_1, \delta} \left[1 + \frac{(1 - c_1)\delta}{(1 - \delta)^2 c_1} \right] = \inf_{c_1} \left[1 + \frac{(1 - c_1)^2 (G - 1)}{(G - 2 + c_1)^2 c_1} \right] = 1.$$

Combined with (70) this yields (21) as stated. \blacksquare

Proof of Theorem 3: Define the delete-one-cluster operator $\widetilde{\mathbf{M}}$, which is the $n \times n$ matrix with g j th block

$$\widetilde{\mathbf{M}}_{gj} = \begin{cases} \mathbf{I}_{n_g} & g = j \\ -\mathbf{X}_g (\mathbf{X}' \mathbf{X} - \mathbf{X}'_g \mathbf{X}_g)^{-1} \mathbf{X}'_j & g \neq j. \end{cases}$$

This operator has the algebraic property that it creates delete-one-cluster prediction errors and is the jackknife analog of the least squares annihilation matrix. Define $\mathbf{Z}_g = \mathbf{X}_g (\mathbf{X}' \mathbf{X})^{-1} R$,

$$\widetilde{\mathbf{Z}} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{0}_{n_1 \times 1} & \cdots & \mathbf{0}_{n_1 \times 1} \\ \mathbf{0}_{n_2 \times 1} & \mathbf{Z}_2 & \cdots & \mathbf{0}_{n_2 \times 1} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{n_G \times 1} & \mathbf{0}_{n_G \times 1} & \cdots & \mathbf{Z}_G \end{bmatrix}, \quad (71)$$

and

$$\mathbf{B} = \widetilde{\mathbf{M}}' \widetilde{\mathbf{Z}}. \quad (72)$$

The matrices $\bar{\mathbf{Z}}$ and \mathbf{B} are $n \times G$. We now show that under Assumption 3,

$$\mathbf{B}'\mathbf{X} = 0. \quad (73)$$

The left side of (73) is a $G \times 1$ vector with g th element

$$\begin{aligned} & R'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g\mathbf{X}_g \left[\mathbf{I}_{n_g} - (\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)^{-} (\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g) \right] \\ &= R'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'_g\mathbf{X}_g + R'(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)^{-}(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g) \\ &\quad - R'(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)^{-}(\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g) \\ &= R' \left[\mathbf{I}_{n_g} - (\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g)^{-} (\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g) \right]. \end{aligned} \quad (74)$$

The final equality uses the generalized inverse property $\mathbf{A}\mathbf{A}^{-}\mathbf{A} = \mathbf{A}$.

If part (1) of Assumption 3 holds, (74) equals zero. Suppose part (2) holds. As inference concerns the coefficients on \mathbf{X}_1 only, least squares estimation and inference is unaffected if we replace \mathbf{X}_1 in \mathbf{X} with the residual $\dot{\mathbf{X}}_1$. Then

$$\mathbf{X}'\mathbf{X} - \mathbf{X}'_g\mathbf{X}_g = \begin{bmatrix} \dot{\mathbf{X}}_1'\dot{\mathbf{X}}_1 - \dot{\mathbf{X}}_{1g}'\dot{\mathbf{X}}_{1g} & 0 \\ 0 & \mathbf{X}'_2\mathbf{X}_2 - \mathbf{X}'_{2g}\mathbf{X}_{2g} \end{bmatrix}.$$

Since $\dot{\mathbf{X}}_1'\dot{\mathbf{X}}_1 - \dot{\mathbf{X}}_{1g}'\dot{\mathbf{X}}_{1g}$ is invertible, (74) equals

$$R' \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{I} - (\mathbf{X}'_2\mathbf{X}_2 - \mathbf{X}'_{2g}\mathbf{X}_{2g})^{-} (\mathbf{X}'_2\mathbf{X}_2 - \mathbf{X}'_{2g}\mathbf{X}_{2g}) \end{bmatrix} = 0,$$

the final equality because R only loads on \mathbf{X}_1 . Hence under Assumption 3, (73) holds, as claimed.

It will be useful to observe that $\text{rank}(\mathbf{B}) \leq \max[G, n-k]$. This holds because, firstly, $\text{rank}(\mathbf{B}) \leq G$ as \mathbf{B} is $n \times G$, and secondly, because $\text{rank}(\mathbf{B}) \leq n-k$ as (73) shows that $G \times n$ \mathbf{B}' has null space \mathbf{X} , which is of dimension k .

We next write the coefficient estimator and squared standard error \hat{v}_5^2 as explicit linear/quadratic functions of the regression error vector \mathbf{e} . First,

$$\hat{\theta} - \theta = R'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{e} = \mathbf{Z}'\mathbf{e} \quad (75)$$

where $\mathbf{Z} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}R$ is the vector of stacked \mathbf{Z}_g . Second, let $\hat{\mathbf{e}}_-$ denote the stacked prediction errors $\hat{\mathbf{e}}_{-g}$, which satisfies $\hat{\mathbf{e}}_- = \widetilde{\mathbf{M}}\mathbf{Y}$. Using (14), the definitions of \mathbf{Z}_g and $\bar{\mathbf{Z}}$, the equation

$\hat{\mathbf{e}}_- = \widetilde{\mathbf{M}}\mathbf{Y}$, (72), $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$, and finally (73), we can write

$$\begin{aligned}\hat{v}_5^2 &= \mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{g=1}^G \mathbf{X}'_g \hat{\mathbf{e}}_{-g} \hat{\mathbf{e}}'_{-g} \mathbf{X}_g \right) (\mathbf{X}'\mathbf{X})^{-1} \mathbf{R} \\ &= \sum_{g=1}^G \mathbf{Z}'_g \hat{\mathbf{e}}_{-g} \hat{\mathbf{e}}'_{-g} \mathbf{Z}_g\end{aligned}\quad (76)$$

$$\begin{aligned}&= \hat{\mathbf{e}}'_- \bar{\mathbf{Z}} \bar{\mathbf{Z}}' \hat{\mathbf{e}}_- \\ &= \mathbf{Y}' \widetilde{\mathbf{M}}' \bar{\mathbf{Z}} \bar{\mathbf{Z}}' \widetilde{\mathbf{M}} \mathbf{Y} \\ &= (\mathbf{X}\boldsymbol{\beta} + \mathbf{e})' \mathbf{B}\mathbf{B}' (\mathbf{X}\boldsymbol{\beta} + \mathbf{e}) \\ &= \mathbf{e}' \mathbf{B}\mathbf{B}' \mathbf{e}.\end{aligned}\quad (77)$$

It will be useful to observe that Theorem 1 and (77) imply that

$$v^2 \leq \mathbb{E}[\hat{v}_5^2] = \mathbb{E}[\mathbf{e}' \mathbf{B}\mathbf{B}' \mathbf{e}] = \mathbb{E}[\text{tr}(\mathbf{B}\mathbf{B}' \mathbf{e}\mathbf{e}')] = \text{tr}(\mathbf{B}\mathbf{B}' \boldsymbol{\Sigma}). \quad (78)$$

Since $\mathbf{e} \sim \text{N}(\mathbf{0}, \boldsymbol{\Sigma})$ we can write $\mathbf{e} = \boldsymbol{\Sigma}^{1/2} \boldsymbol{\psi}$ with $\boldsymbol{\psi} \sim \text{N}(\mathbf{0}, \mathbf{I}_n)$. Using $\widehat{C}_5(c) = \widehat{\theta} \pm \widehat{v}_5 c$, (75), (77), and $\mathbf{e} = \boldsymbol{\Sigma}^{1/2} \boldsymbol{\psi}$, we find

$$\begin{aligned}\mathbb{P}[\theta \in \widehat{C}_5(c)] &= \mathbb{P}\left[(\widehat{\theta} - \theta)^2 \leq c^2 \widehat{v}_5^2\right] \\ &= \mathbb{P}[\mathbf{e}' \mathbf{Z}\mathbf{Z}' \mathbf{e} \leq c^2 \mathbf{e}' \mathbf{B}\mathbf{B}' \mathbf{e}] \\ &= \mathbb{P}[0 \leq \mathbf{e}' (c^2 \mathbf{B}\mathbf{B}' - \mathbf{Z}\mathbf{Z}') \mathbf{e}] \\ &= \mathbb{P}[0 \leq \boldsymbol{\psi}' \mathbf{C} \boldsymbol{\psi}]\end{aligned}\quad (79)$$

where

$$\mathbf{C} = \boldsymbol{\Sigma}^{1/2} (c^2 \mathbf{B}\mathbf{B}' - \mathbf{Z}\mathbf{Z}') \boldsymbol{\Sigma}^{1/2}. \quad (80)$$

By the spectral decomposition, $\mathbf{C} = \mathbf{H}\boldsymbol{\Lambda}\mathbf{H}'$ where $\boldsymbol{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_n\}$, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ are the eigenvalues of \mathbf{C} , and $\mathbf{H}'\mathbf{H} = \mathbf{I}_n$. Define $\boldsymbol{\xi} = \mathbf{H}'\boldsymbol{\psi} \sim \text{N}(\mathbf{0}, \mathbf{I}_n)$ and make the partition $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)'$. Set $\lambda_0 = -\lambda_n$ and $\xi_0^2 = \xi_n^2$. We find that (79) equals

$$\mathbb{P}[0 \leq \boldsymbol{\psi}' \mathbf{H}\boldsymbol{\Lambda}\mathbf{H}' \boldsymbol{\psi}] = \mathbb{P}[0 \leq \boldsymbol{\xi}' \boldsymbol{\Lambda} \boldsymbol{\xi}] = \mathbb{P}\left[\lambda_0 \xi_0^2 \leq \sum_{j=1}^{n-1} \lambda_j \xi_j^2\right]. \quad (81)$$

Next, we establish the following properties of the eigenvalues:

$$\lambda_j \geq 0 \text{ for } j < n \quad (82)$$

$$\lambda_0 \geq 0 \quad (83)$$

$$\lambda_0 \leq v^2 \quad (84)$$

$$\bar{\lambda} \stackrel{\text{def}}{=} \sum_{j=1}^{n-1} \lambda_j \geq c^2 \lambda_0 \quad (85)$$

with (85) holding for $c \geq 1$.

Let $\lambda_j(\mathbf{D})$ denote the j th largest eigenvalue of a Hermitian matrix \mathbf{D} . By a corollary of the Weyl eigenvalue inequality for Hermitian matrices (Corollary 4.3.15 of Horn and Johnson (2013)), for any Hermitian matrices \mathbf{A} and \mathbf{B} , and any j ,

$$\lambda_j(\mathbf{A}) + \lambda_{\min}(\mathbf{B}) \leq \lambda_j(\mathbf{A} + \mathbf{B}) \leq \lambda_{\max}(\mathbf{A}) + \lambda_j(\mathbf{B}). \quad (86)$$

Using the definition (80) and lower bound in (86), for each $1 \leq j < n$,

$$\begin{aligned} \lambda_j &\geq \lambda_{\min}(c^2 \boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) + \lambda_j(-\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) \\ &= c^2 \lambda_{\min}(\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) - \lambda_{n-j+1}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) = 0, \end{aligned}$$

the final equality since $\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}$ and $\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}$ are positive semi-definite with deficient rank. (Indeed, $\text{rank}(\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) \leq \text{rank}(\mathbf{B}) \leq \max[G, n - k] < n$ and $\text{rank}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) = 1$.) This is (82).

Using the lower bound in (86) and similar reasoning,

$$\begin{aligned} \lambda_0 &= \lambda_{\max}(-\mathbf{C}) \\ &= \lambda_{\max}(-c^2 \boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2} + \boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) \\ &\geq c^2 \lambda_{\max}(-\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) + \lambda_{\min}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) \\ &= -c^2 \lambda_{\min}(\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) + \lambda_{\min}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) = 0. \end{aligned}$$

This is (83).

Using the upper bound in (86),

$$\begin{aligned}
\lambda_0 &= \lambda_{\max}(-c^2 \boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2} + \boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) \\
&\leq \lambda_{\max}(-c^2 \boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) + \lambda_{\max}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) \\
&= -c^2 \lambda_{\min}(\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) + \mathbf{Z}' \boldsymbol{\Sigma} \mathbf{Z} \\
&= v^2,
\end{aligned}$$

since $\mathbf{Z}' \boldsymbol{\Sigma} \mathbf{Z} = v^2$. This is (84).

Using the fact that the trace of a matrix equals the sum of its eigenvalues, $\lambda_n = -\lambda_0$, (80), (78), and $\mathbf{Z}' \boldsymbol{\Sigma} \mathbf{Z} = v^2$,

$$\begin{aligned}
\bar{\lambda} &= \text{tr}(\mathbf{C}) + \lambda_0 \\
&= c^2 \text{tr}(\boldsymbol{\Sigma}^{1/2} \mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}^{1/2}) - \text{tr}(\boldsymbol{\Sigma}^{1/2} \mathbf{Z} \mathbf{Z}' \boldsymbol{\Sigma}^{1/2}) + \lambda_0 \\
&= c^2 \text{tr}(\mathbf{B} \mathbf{B}' \boldsymbol{\Sigma}) - \mathbf{Z}' \boldsymbol{\Sigma} \mathbf{Z} + \lambda_0 \\
&\geq c^2 v^2 - v^2 + \lambda_0 \\
&= c^2 \lambda_0 + (v^2 - \lambda_0)(c^2 - 1) \geq c^2 \lambda_0,
\end{aligned}$$

The final inequality holds by (84) and $c^2 \geq 1$. This is (85). We have established (82)-(85) as stated.

From (79) and (81) we have

$$\mathbb{P}[\theta \in \widehat{C}_5(c)] = \mathbb{P}\left[\lambda_0 \xi_0^2 \leq \sum_{j=1}^{n-1} \lambda_j \xi_j^2\right]. \quad (87)$$

From (82)-(83), $\lambda_j \geq 0$ for $j = 0, \dots, n-1$. Suppose $\lambda_0 = 0$. Then (87) equals 1, which satisfies (24). Now suppose $\lambda_0 > 0$. For $j \geq 1$ define $w_j = \lambda_j / \bar{\lambda}$, which satisfy $\sum_{j=1}^{n-1} w_j = 1$ and $w_j \geq 0$. Then (87) equals

$$\mathbb{P}\left[\frac{\xi_0^2}{\sum_{j=1}^{n-1} w_j \xi_j^2} \leq \frac{\bar{\lambda}}{\lambda_0}\right] \geq \mathbb{P}\left[\zeta^2 \leq \frac{\bar{\lambda}}{\lambda_0}\right] \geq \mathbb{P}[\zeta^2 \leq c^2] = \mathbb{P}[|\zeta| \leq c].$$

The first inequality is the right-side inequality of Corollary 2 of Makshanov and Shalaevski (1986). The second inequality is (85). This satisfies (24), and we have shown the latter holds whether $\lambda_0 = 0$ or $\lambda_0 > 0$, completing the proof. \blacksquare

The proofs of Theorems 4-7 are in the Online Appendix.

Supplementary Data

Supplementary data are available at the *Review of Economic Studies* online.

Data Availability

Full replication code and data is available on Zenodo at <https://doi.org/10.5281/zenodo.17712099>.

References

- [1] ANDREWS, D. W. K. (1991), "Asymptotic Normality of Series Estimators for Nonparametric and Semiparametric Regression Models," *Econometrica*, 59, 307-345.
- [2] ARELLANO, M. (1987), "Computing Robust Standard Errors for Within Groups Estimators," *Oxford Bulletin of Economics and Statistics* 49, 431-434.
- [3] BAHADUR, R. R. and SAVAGE, L. J. (1956), "The Nonexistence of Certain Statistical Procedures in Nonparametric Problems," *Annals of Mathematical Statistics*, 27, 1115-1122.
- [4] BELL, R. M., and McCAFFREY, D. F. (2002), "Bias Reduction in Standard Errors for Linear Regression with Multi-Stage Samples," *Survey Methodology*, 28, 169-181.
- [5] BERA, A. K., SUPRAYITNO, T., and PREMARATNE, G. (2002), "On Some Heteroskedasticity-Robust Estimators of Variance-Covariance Matrix of the Least-Squares Estimators," *Journal of Statistical Planning and Inference*, 108, 121-136.
- [6] BERTANHA, M. and MOREIRA, M. J. (2020), "Impossible Inference in Econometrics: Theory and Applications," *Journal of Econometrics*, 218, 247-270.
- [7] BODENHAM, D. A. and ADAMS, N. M. (2016), "A Comparison of Efficient Approximations for a Weighted Sum of Chi-Squared Random Variables," *Statistics and Computing*, 26, 917-928.
- [8] CAMERON, A. C., GELBACH, J. B., and MILLER, D. L. (2008), "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90, 414-427.
- [9] CANAY, I. A., SANTOS, A., and SHAIKH, A. M. (2021), "The Wild Bootstrap with a Small Number of Large Clusters," *Review of Economics and Statistics*, 103, 346-363.
- [10] CATTANEO, M. D., JANSSON, M., and NEWWEY, W. K. (2018), "Inference in Linear Regression Models with Many Covariates and Heteroskedasticity," *Journal of the American Statistical Association*, 113, 1350-1361.
- [11] CHESHER, A. D. (1989), "Hájek Inequalities, Measures of Leverage, and the Size of Heteroskedasticity Robust Wald Tests," *Econometrica*, 57, 971-977.
- [12] CHESHER, A. D. and AUSTIN, G. (1991), "The Finite-Sample Distributions of Heteroskedasticity Robust Wald Statistics," *Journal of Econometrics*, 47, 153-173.
- [13] CHESHER, A. D. and JEWITT, I. D. (1987), "The Bias of the Heteroskedasticity Consistent Covariance Matrix Estimator," *Econometrica*, 55, 1217-1272.

- [14] COCHRAN, W. G. (1977), *Sampling Techniques*, 3rd Edition, (New York: Wiley).
- [15] CONLEY, T. G. and TABER, C. R. (2011), "Inference with 'Difference in Differences' with a Small Number of Policy Changes," *Review of Economics and Statistics*, 93, 113-125.
- [16] COUDIN, E. and DUFOUR, J-M (2009), "Finite-Sample Distribution-Free Inference in Linear Median Regressions Under Heteroskedasticity and Nonlinear Dependence of Unknown Form," *Econometrics Journal*, 12, S19-S49.
- [17] DAVIDSON, R., and MACKINNON, J. G. (1993), *Estimation and Inference in Econometrics*, Oxford University Press.
- [18] DJOGBENOU, A. A., MACKINNON, J. G. and Nielsen, M. Ø. (2019), "Asymptotic Theory and Wild Bootstrap Inference with Clustered Errors," *Journal of Econometrics*, 212, 393-412.
- [19] DUFOUR, J-M (1997), "Some Impossibility Theorems in Econometrics, with Applications to Structural and Dynamic Models," *Econometrica*, 65, 1365-1389.
- [20] DUFOUR, J-M (2003), "Identification, Weak Instruments and Statistical Inference in Econometrics," *Canadian Journal of Economics*, 36, 767-808.
- [21] EFRON, B. (1982), *The Jackknife, the Bootstrap, and Other Resampling Plans*, Society for Industrial and Applied Mathematics.
- [22] EFRON, B. and STEIN, C. (1981), "The Jackknife Estimate of Variance," *The Annals of Statistics*, 9, 586-596.
- [23] EICKER, F. (1963), "Asymptotic Normality and Consistency of the Least Squares Estimators for Families of Linear Regressions," *Annals of Mathematical Statistics*, 34, 447-456.
- [24] FERMAN, B. and PINTO, C. (2019), "Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity," *Review of Economics and Statistics*, 101, 452-467.
- [25] HAGEMANN, A. (2019), "Placebo Inference on Treatment Effects When the Number of Clusters is Small," *Journal of Econometrics*, 213, 190-209.
- [26] HAGEMANN, A. (2025), "Inference With a Single Treated Cluster," *Review of Economic Studies*, 92, 3968-3994.
- [27] HANSEN, B. E. (2025), "Standard Errors for Difference-in-Difference Regression," *Journal of Applied Econometrics*, 40, 291-309.
- [28] HANSEN, B. E. and LEE, S. (2019), "Asymptotic Theory for Clustered Samples," *Journal of Econometrics*, 210, 268-290.
- [29] HINKLEY, D. V. (1977), "Jackknifing in Unbalanced Situations," *Technometrics*, 19, 285-292.
- [30] HIRANO, K. and PORTER, J. R. (2012), "Impossibility Results for Nondifferentiable Functionals," *Econometrica*, 80, 1769-1790.

- [31] HORN, R. A. and JOHNSON, C. R. (2013), *Matrix Analysis*, Second Edition, Cambridge University Press.
- [32] HUBER, P. J. (1967), "The Behavior of Maximum Likelihood Estimates Under Nonstandard Conditions," *Proceedings of the Berkeley Symposium on Mathematical Statistics and Probability*, L. M. Le Cam and Neyman, J. editors, 1, 221-223.
- [33] IBRAGIMOV, R. and MÜLLER, U. K. (2010), "*t*-Statistic Based Correlation and Heterogeneity Robust Inference," *Journal of Business and Economic Statistics*, 28, 453-468.
- [34] IBRAGIMOV, R. and MÜLLER, U. K. (2016), "Inference with a Few Heterogeneous Clusters," *Review of Economics and Statistics*, 98, 83-96.
- [35] IMBENS, G. W. and KOLESÁR, M. (2016), "Robust Standard Errors in Small Samples: Some Practical Advice," *Review of Economics and Statistics*, 98, 701-712.
- [36] KHURI, A. I. (1995), "A Measure to Evaluate the Closeness of Satterthwaite's Approximation," *Biometrical Journal*, 37, 547-563.
- [37] KLINE, P, SAGGIO, R. and SØLVSTEN, M. (2020), "Leave-Out Estimation of Variance Components," *Econometrica*, 88, 1859-1898.
- [38] KOLESÁR, M. (2023), "Robust Standard Errors in Small Samples," (unpublished, Princeton University).
- [39] KRANZ, S. (2024), "From Replications to Revelations: Heteroskedasticity-Robust Inference," (unpublished, Ulm University).
- [40] LIANG, K-Y, and ZEGER, S. L. (1986), "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 73, 13-22.
- [41] LONG, J. S., and ERVIN, L. H. (2000), "Using Heteroscedasticity Consistent Standard Errors in the Linear Regression Model," *The American Statistician*, 54, 217-224.
- [42] MACKINNON, J. G., and WEBB, M. D. (2020), "Randomization Inference for Difference-in-Differences with Few Treated Clusters," *Journal of Econometrics*, 218, 435-450.
- [43] MACKINNON, J. G. and WHITE, H. (1985), "Some Heteroskedasticity-Consistent Covariance Matrix Estimators with Improved Finite Sample Properties," *Journal of Econometrics*, 29, 305-325.
- [44] MACKINNON, J. G., NIELSEN, M. Ø., and WEBB, M. D. (2023a), "Cluster-Robust Inference: A Guide to Empirical Practice," *Journal of Econometrics*, 232, 272-299.
- [45] MACKINNON, J. G., NIELSEN, M. Ø., and WEBB, M. D. (2023b), "Fast and Reliable Jackknife and Bootstrap Methods for Cluster-Robust Inference," *Journal of Applied Econometrics*, 38, 671-694.

- [46] MACKINNON, J. G., NIELSEN, M. Ø., and WEBB, M. D. (2023c), “Leverage, Influence, and the Jackknife in Clustered Regression Models: Reliable Inference Using `summclust`,” *Stata Journal*, 4, 942-982.
- [47] MAKSHANOV, A. V. and SHALAEVSKI, O. V. (1986), “Some Problems of Asymptotic Approximation of Distributions,” *Journal of Mathematical Sciences*, 34, 1433-1445.
- [48] MENG, X., QIAN, N., and YARED, P. (2015), “The Institutional Causes of China’s Great Famine, 1959-1961,” *Review of Economic Studies*, 82, 1568-1611.
- [49] NICCODEMI, G. and WANSBEEK, T. (2022), “A New Estimator for Standard Errors with a Few Unbalanced Clusters,” *Econometrics*, 10, 6.
- [50] PÖTSCHER, B. M. and PREINERSTORFER, D. (2025), “Valid Heteroskedasticity Robust Testing,” *Econometric Theory*, 41, 249-301.
- [51] PREINERSTORFER, D. and PÖTSCHER, B. M. (2016), “On Size and Power of Heteroskedasticity and Autocorrelation Robust Tests,” *Econometric Theory*, 32, 261-358.
- [52] ROMANO, J. P. (2004), “On Non-Parametric Testing, the Uniform Behavior of the t -test, and Related Problems,” *Scandinavian Journal of Statistics*, 31, 567-584.
- [53] RUST, K. F. and RAO, J. N. K. (1996), “Variance Estimation for Complex Surveys Using Replication Techniques,” *Statistical Methods in Medical Research*, 5, 283-310.
- [54] SATTERTHWAITTE, F. E. (1946), “An Approximate Distribution of Estimates of Variance Components,” *Biometrics Bulletin*, 2, 110-114.
- [55] SHAO, J. and TU, D. (1995), *The Jackknife and Bootstrap*, Springer.
- [56] TUKEY, J. (1958), “Bias and Confidence in Not Quite Large Samples,” *Annals of Mathematical Statistics*, 29, 614.
- [57] WHITE, H. (1980), “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity,” *Econometrica*, 48, 817-838.
- [58] YOUNG, A. (2016), “Improved, Nearly Exact, Statistical Inference With Robust and Clustered Covariance Matrices Using Effective Degrees of Freedom Corrections,” (unpublished, London School of Economics).
- [59] YOUNG, A. (2019), “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *Quarterly Journal of Economics*, 134, 557-598.