Abstract

Abadie et al. (QJE 2023) propose two new variance estimators, one analytical and one bootstrap, for the important use case of a binary treatment where the number of clusters is finite (e.g., U.S. states), there are many treated and control units in each cluster, and only some clusters are observed. They study a setting where treatment is randomly assigned within cluster but the treatment effect and probability of assignment to treatment vary across clusters. Their simulations show that when all clusters are observed, conventional clustered standard errors (c.s.e.'s) are severely conservative, but they find close-to-correct s.e.'s and coverage for their proposed two-stage cluster bootstrap (TSCB), with both OLS and cluster fixed effects, and good performance for an analytical variance estimator with cluster fixed effects. We confirm that c.s.e.'s are conservative when all clusters (out of 52) are observed, but show that their results are a corner case. When less than all clusters are observed (even slightly less), the TSCB produces neither correct s.e.'s nor correct coverage and can be inferior to c.s.e.'s or the wild cluster bootstrap. The analytical variance estimator is also unreliable when only some of 52 clusters are observed, but performs well for a universe of 500 clusters. Which variance estimator comes closer to correct coverage depends on the sample, the treatment, the outcome, the total number of clusters, and the fraction of observed clusters. Moreover, a conservative s.e. does not necessarily imply conservative coverage. Including versus omitting even a single influential cluster (California, in their dataset) can greatly affect s.e. and coverage estimates.