BOLD THINKERS DRIVING REAL-WORLD IMPACT

Unbiased Instrumental Variables (IV) in Stata

Austin Nichols

@austnnchols

Magic Bullets

- Instrumental Variables (IV) methods are the only way to estimate causal effects in a variety of settings, including experiments (randomized control trials or RCTs) with imperfect compliance

❖ IV methods often exhibit poor performance

  – Bias & size distortion with many weak instruments
  – No finite moments when exactly identified

- Andrews and Armstrong (2017) offer a solution

Causal Diagram

- Conditioning on confounders does not in general solve the problem of endogenous participation in a treatment of interest

- The receipt of a treatment (R=1) whose effect β we want to measure may be randomly assigned (Z=1), but we still need IV to estimate impact

We want to measure the effect β of receiving treatment R on the outcome Y

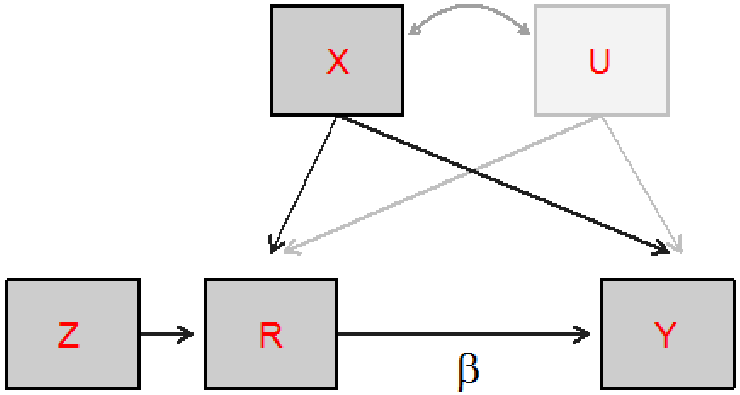

Conditioning on confounders X and U can eliminate bias, but conditioning on X alone can either reduce or amplify bias. Instrumenting for R with Z asymptotically eliminates that bias.

Sign restriction allows unbiased IV

- IV has one fewer moments than overid restrictions, so exactly identified IV has no moments
  - Hirano and Porter (2015) show that mean, median, and quantile unbiased estimation are all impossible in the linear IV model with an unrestricted parameter space for the first stage
- This result no longer holds when the sign of the first stage is known (e.g. no defiers, some compliers):
  - In models with a single instrumental variable, Andrews and Armstrong (2017) show that there is a unique unbiased estimator based on the reduced form and first-stage regression estimates
  - This estimator is substantially less dispersed than the usual 2SLS estimator in finite samples
- In an RCT, we are very confident the first stage is positive

Model and Estimator

$Y=Z\pi\beta+u$ ← reduced form coef $\xi_1=(Z'Z)^{-1}(Z'Y)$

$R=Z\pi+v$ ← first stage coef $\xi_2=(Z'Z)^{-1}(Z'R)$

IV estimator constructs Wald ratio ξ_1 / ξ_2

Assume u,v normal so (ξ_1 , ξ_2)~$N(\mu,\Sigma)$ w/variance $\Sigma=(\sigma_1^2 , \sigma_{12} \backslash \sigma_{12} , \sigma_2^2)$

Let $d=(\xi_1 - \xi_2 \sigma_{12} /\sigma_2^2)$. $E[d]=\pi\beta-\pi\sigma_{12} /\sigma_2^2$

Voinov and Nikulin (1993) show that unbiased estimation of $1/\pi$ is possible if its sign is known:

Let $t=\Phi(-\xi_2 /\sigma_2)/\phi(\xi_2 /\sigma_2)\sigma_2$ then $E[t]= 1/\pi$ and $E[dt]= E[d]E[t]= \beta-\sigma_{12} /\sigma_2^2$

Estimator $b_U=dt+s_{12}/v_2$

Further considerations

- b_U is asymptotically equivalent to 2SLS when instruments are strong and thus b_U can be used together with conventional 2SLS standard errors

- Optimal *estimation* and optimal *testing* are distinct questions in the context of weak instruments
  - b_U is uniformly minimum risk unbiased for convex loss, but it follows from the results of Moreira (2009) that the Anderson–Rubin test is the uniformly most powerful unbiased two-sided test in the just-identified context (not a conditional t-test based on b_U)
  - more research needed on *tests* based on this unbiased IV estimator…

Small-Sample Properties

- Note this applies to bivariate normal errors with known variance, not the focal case of random assignment Z={0,1} and endogenous receipt of treatment R={0,1}
  - Appendix B (Nonnormal errors and unknown reduced-form variance) “derives asymptotic results for the case with non-normal errors and an estimated reduced-form covariance matrix. Appendix B.1 shows asymptotic unbiasedness in the weak-instrument case. Appendix B.2 shows asymptotic equivalence with 2SLS in the strong-instrument case”
  - How does this approach perform in finite samples?

Stata command

- Estimator implemented as `aaniv` on SSC
- Download using `ssc install aaniv`
- So far, just one endogenous treatment and one excluded instrument (as of today), as is ideal for an RCT, but the command will be updated in future releases to a larger set of use cases

Small-Sample Properties

- Even with binary R and Z, so non-normal errors by design, standard linear regression rejects the truth all the time, and unbiased IV outperforms standard IV/2SLS

  (this simulation has a high correlation between a normal variate that predicts R and the unobserved error that predicts the outcome Y)

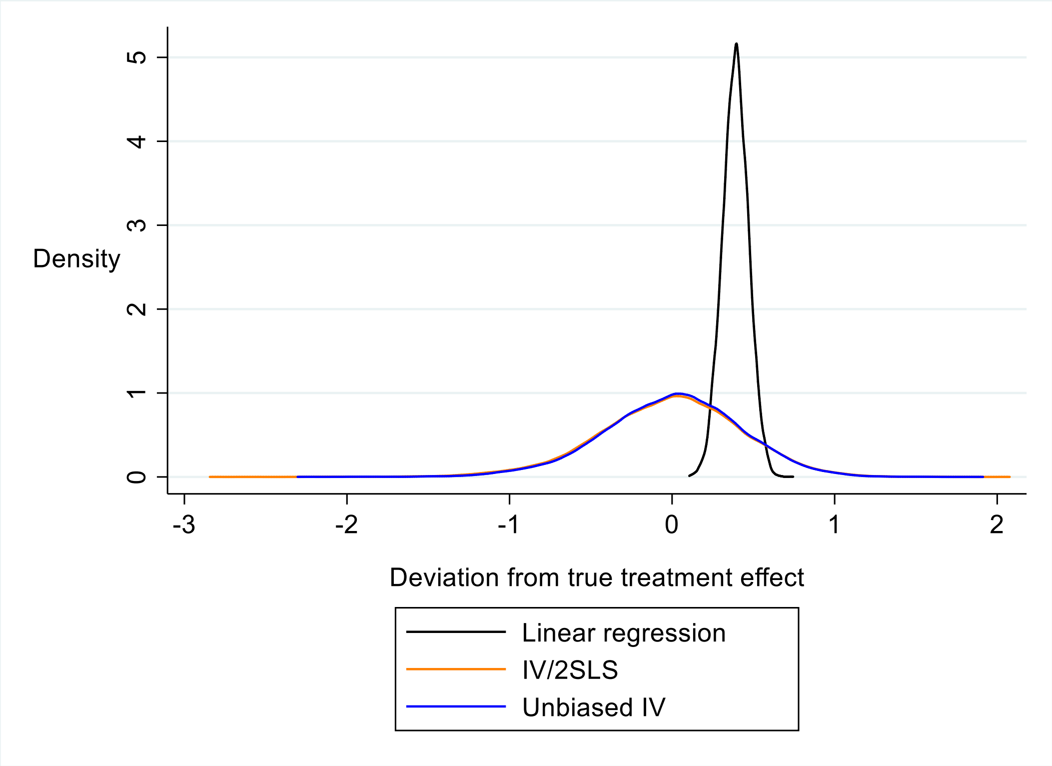

Distributions of Estimators by Sample Size and Correlation

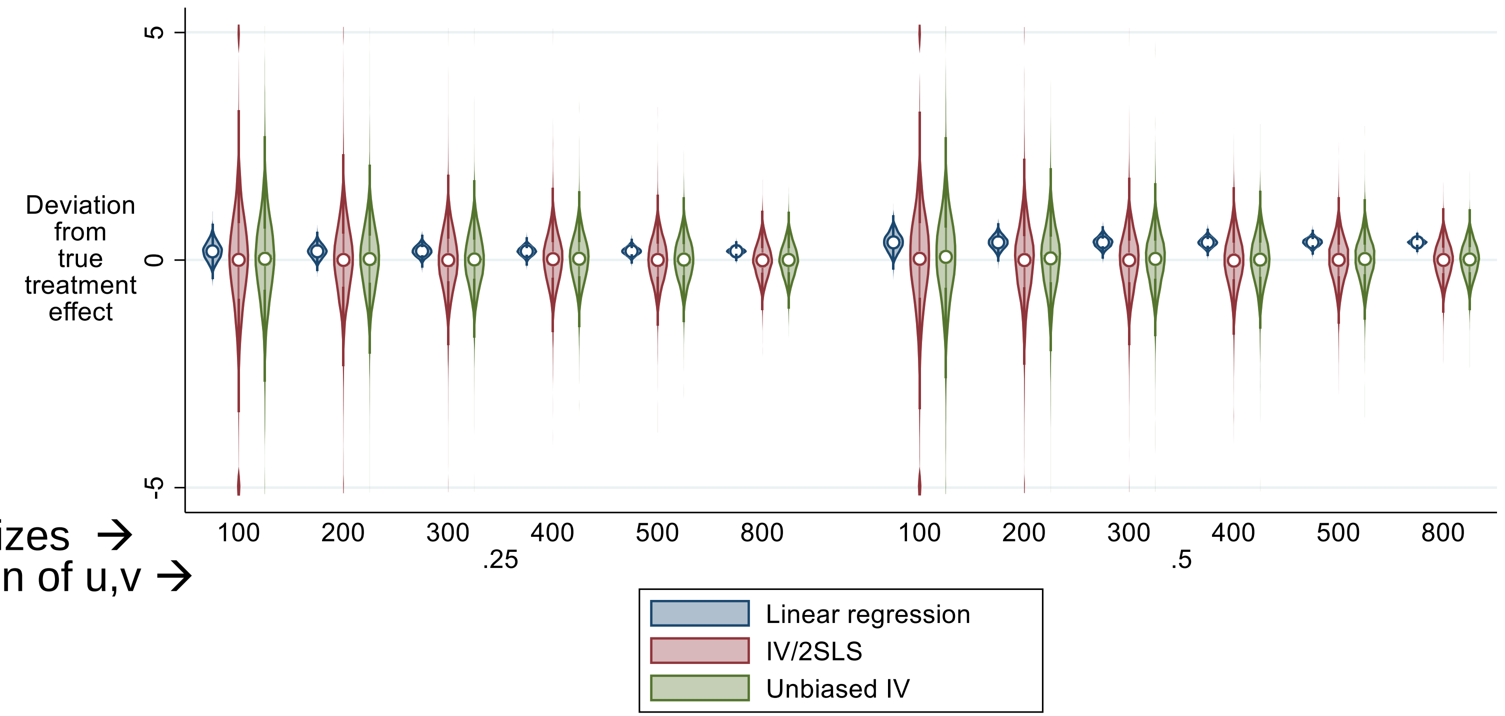

Sample sizes →
Correlation of u,v →

Rejection rates about right for IV models, in large samples

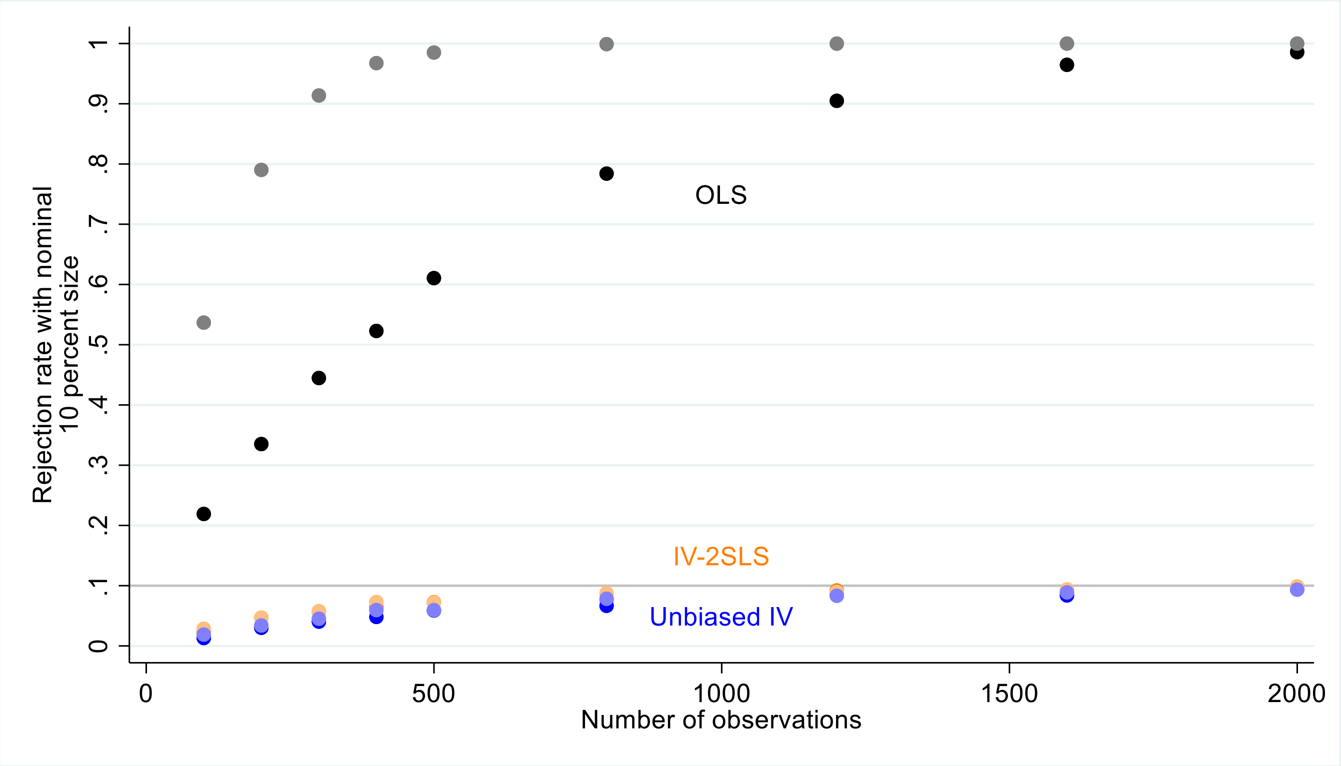

Conclusion

- Unbiased IV performs as well as IV-2SLS in a setting that it is not designed for, with no bias and lower evident dispersion (but neither has a finite variance)
  - Report unbiased IV for an experiment, if only to enable meta-analysis; use `aaniv` (ssc install aaniv) in Stata
- Rejection rates for both Unbiased IV and IV 2SLS approximately at the nominal rate when sample size is over a thousand
  - At smaller sample sizes, there is some under-rejection of a true null—needs further study

Contact

Austin Nichols

Principal Scientist

austinnichols@gmail.com

BOLD
THINKERS
DRIVING
REAL-WORLD
IMPACT

abtassociates.com