Section 5 Notes
Instrumental Variables and Regression Discontinuity

Stephen Pettigrew

November 17, 2015

1 Instrumental Variables

Setup: We have a variable that we want to know the treatment effect of, but it's not randomized. There's another variable that is randomized and it's correlated with the treatment we care about.

In Figure1, U and X_1 confound the relationship between the treatment and outcome (D and Y). X_1 is observed so you can control for it, but U is not so you can't get an unbiased estimate of the effect of D on Y. I is independent of everything, except through D, so we can use it as an instrument.

Figure 1:

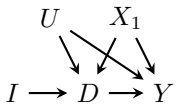

What you're estimating: If you're willing to assume that there are constant treatment effects across all units, then your IV estimate equals the ATE. If that assumption is too strong for you (which it probably is), then your IV estimate equals the ATE for compliers.
How to estimate it: Two stage least squares. In R, `ivreg()` in the `AER` package
Things to think about:

1. Strength of instrument: how strongly is the instrument correlated with the treatment?
    - If you have a strong instrument then you can use two stage least squares to estimate the treatment effect
    - If you have an extremely strong instrument (i.e. the instrument and treatment are almost perfectly correlated) and the instrument is randomized, then in effect the treatment itself is also randomized so 2SLS, linear regression, difference in means, or other techniques will all give similar results
    - If you have a very weak instrument (i.e. the instrument and treatment aren't correlated very much), then your estimates of the treatment effect will be unbelievably high variance. Asymptotically you'll still get an unbiased estimate of the treatment effect, but you can't possibly have enough data for those asymptotics to hold
2. Exclusion restriction: is the instrument uncorrelated with or independent of the outcome, except through the treatment?

- If you have an exclusion restriction violation and
  - you know the directionality of the relationship between I and Y, then you can say whether the true treatment effect is higher or lower than the one you estimated
  - you don't know the directionality, then you're really out of luck. It might even be that OLS gives you a better estimate of the true treatment effect than 2SLS, but there's absolutely no way of knowing.
- If you have an exclusion restriction violation, you should think about its size relative to the size of the treatment effect of D on Y. If there's a huge treatment effect of D on Y and only a tiny effect of I on Y, then you might only have a tiny amount of bias. If the effect of I on Y is larger, then the bias can be substantial and you can't draw good conclusions from the data
- The instrument could also be conditionally independent (Figure 2), but why should you be willing to believe that you can account for all confounders between the instrument and outcome if you're not willing believe that you can account for all confounders between treatment and outcome?

Figure 2:

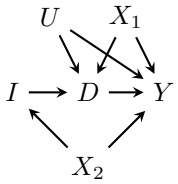

2 Regression Discontinuity

Setup: Treatment isn't randomized, but there's some process that deterministically (and exogenously) dictates whether a unit is treated or not. Using this seemingly arbitrary cutoff, we can compare the outcome for the treated units that are very close to one side of the cutoff to the outcome for the control units that are very close to the other side of the cutoff.

What you're estimating: You're estimating the effect of the treatment on the outcome, for any observations that lie exactly at the treatment/control cutoff of the forcing variable. If you're doing a fuzzy RD, then you're estimating the treatment effect, conditional on the forcing variable equaling the cutpoint *and* only for compliers. Whether that's an interesting/relevant quantity of interest is for you to decide/justify
How to estimate it: Linear regression with an interaction between the forcing variable and the treatment. If the forcing variable is adjusted so that its cutoff is zero then the CATE equals the lower-order coefficient on the treatment dummy variable.
Things to think about:

1. Bandwidth size: how close does an observation have to be to the cutoff to be included in the estimation?
   - Low bandwidths have less biased estimates of the treatment effect at the cutoff, but you're using less data so they're higher variance
   - High bandwidths have less variance because you have more data, but they might be biased
2. Functional form of forcing variable: In your regression, what functional form do you give your forcing variable
   - The typical convention is to include the forcing variable as a cubic term

- You should check a bunch of specifications and show that your results don't depend on the functional form you use

3. Broken-glass sharp or baby-duckling fuzzy?: If treatment and control are perfectly predicted based on whether an observation is above or below the cutpoint in the forcing variable, then you have a sharp RD. If the cutpoint is strongly correlated with your treatment group, but it's a not a perfectly predictive relationship, then it's a fuzzy RD

    - Sometimes this distinction comes down to what you're interested in calculating. Imagine the outcome of interest is crime or the number of police reports. A city has a policy that requires homeless shelters to open when the overnight temperature is expected to dip below 40 degrees, otherwise they're closed.
        - If you're interested in the effect of opening/closing homeless shelters on crime then you're in a sharp world
        - If you're interested in the effect of the number of people on the street on crime then you're in a fuzzy world because the treatment of being off the street isn't perfectly determined by whether it's above or below 40 degrees