

Instrumental Variables

Brenton Kenkel — PSCI 8357

March 31, 2016

A Brief Exposition

The derivations in these notes are taken largely from Angrist and Pischke (2009, chap. 4) and Angrist and Pischke (2015, chap. 4). I recommend buying both of these books if you are interested in instrumental variables or causal inference more generally. Hye Young is going to make you buy Angrist and Pischke (2009) next semester anyway, and Angrist and Pischke (2015) is a gentler introduction to the same material.

Assume we have a binary treatment T_i and an observed response $Y_i = Y_i(T_i)$. Like last week, suppose our goal is to estimate the average treatment effect,

$$\tau = E[Y_i(1) - Y_i(0)].$$

Assume there is an unobserved confounding variable X_i . Remember that confounding means X_i affects both treatment status and the potential outcomes. The ordinary difference of means estimator of the treatment effect is therefore biased.

The ideal solution to this problem would be to go out and collect data on X_i . If the costs of doing so are prohibitive, we have a second-best solution: *instrumental variables*. In this scenario, a variable Z_i is an *instrument* if it meets the following three conditions (Angrist and Pischke 2015, 106–7):

1. *First stage*: Z_i affects T_i . More generally, the instrument(s) affect the variable(s) of interest.
2. *Independence*: Z_i is independent from X_i . More generally, the instrument(s) are independent from any other confounding variables.
3. *Exclusion restriction*: Z_i only affects Y_i through its effect on T_i . More generally, the instrument(s) affect the response only through their effect on the variable(s) of interest.

Taken together, the latter two conditions imply that Z_i is independent of the potential outcomes:

$$Z_i \perp\!\!\!\perp (Y_i(0), Y_i(1)).$$

The exclusion restriction implies that

$$\text{effect of } Z_i \text{ on } Y_i = (\text{effect of } Z_i \text{ on } T_i) \times (\text{effect of } T_i \text{ on } Y_i)$$

We are most interested in the last term, the effect of the treatment. Rearranging the above expression, we see that the average treatment effect is a ratio of effects of the instrument:

$$\text{effect of } T_i \text{ on } Y_i = \frac{\text{effect of } Z_i \text{ on } Y_i}{\text{effect of } Z_i \text{ on } T_i}.$$

The first stage assumption is what allows us to make this rearrangement—if Z_i has no effect on T_i , then we are dividing by zero. The independence assumption allows us to individually estimate the two components of the ratio.

Assume linear models for both treatment and response:

$$\begin{aligned} T_i &= \gamma_0 + \gamma_1 Z_i + \gamma_2 X_i + \eta_i, \\ Y_i &= \beta_0 + \beta_1 T_i + \beta_2 X_i + \epsilon_i, \end{aligned}$$

where η_i and ϵ_i have mean zero and are independent of X_i , Z_i , and T_i . This model implies a constant treatment effect of β_1 for every observation, so the average treatment effect $\tau = \beta_1$. You will learn in Stat III how the interpretation of instrumental variable estimates changes when the treatment effect is allowed to vary across individuals.

Substituting the first equation into the second, we have

$$\begin{aligned} Y_i &= \beta_0 + \beta_1 T_i + \beta_2 X_i + \epsilon_i \\ &= \beta_0 + \beta_1 (\gamma_0 + \gamma_1 Z_i + \gamma_2 X_i + \eta_i) + \beta_2 X_i + \epsilon_i \\ &= \underbrace{\beta_0 + \beta_1 \gamma_0}_{\pi_0} + \underbrace{\beta_1 \gamma_1}_{\pi_1} Z_i + \underbrace{(\beta_1 \gamma_2 + \beta_2) X_i + \beta_1 \eta_i + \epsilon_i}_{\nu_i} \\ &= \pi_0 + \pi_1 Z_i + \nu_i. \end{aligned}$$

As in the verbal representation above, we can write the effect of the treatment as a ratio of effects of the instrument:

$$\beta_1 = \frac{\pi_1}{\gamma_1}.$$

The question, then, is whether we can consistently estimate π_1 and γ_1 . If so, then their ratio is a consistent estimate of the average treatment effect, β_1 . We will rely on the following fun fact about omitted variable bias (or the lack thereof):

If the included variables are uncorrelated with the error term and any omitted variables, then OLS gives an unbiased and consistent estimate of the coefficients of the included variables.

To estimate π_1 , we will run a regression of Y_i on Z_i (and a constant). Our independence assumption implies that

$$\text{Cov}(Z_i, v_i) = \text{Cov}(Z_i, (\beta_1\gamma_2 + \beta_2)X_i + \beta_1\eta_i + \epsilon_i) = 0,$$

so the coefficient estimate $\hat{\pi}_1$ is a consistent estimate of π_1 .

To estimate γ_1 , we will run a regression of T_i on Z_i . Our independence assumption implies that

$$\text{Cov}(Z_i, \gamma_2 X_i + \eta_i) = 0,$$

so $\hat{\gamma}_1$ is a consistent estimate of γ_1 .

Finally, then, the *instrumental variable estimator* of the average treatment effect is

$$\hat{\beta}_1 = \frac{\hat{\pi}_1}{\hat{\gamma}_1}.$$

Under the first stage, independence, and exclusion restriction assumptions, $\hat{\beta}_1$ is a consistent estimator of β_1 .¹ An instrumental variable lets us break a seemingly intractable problem—estimating a treatment effect when strong ignorability does not hold—into two tractable problems. The trick, of course, is finding a good instrument. These three conditions, particularly the independence assumption and the exclusion restriction, are hard to meet.

The first stage assumption is more plausible; in the social world, we tend to assume most things have some effect on most other things, though perhaps a small one. However, there is a cost to using a *weak instrument*—an instrument that only barely meets the first stage assumption. We divide by $\hat{\gamma}_1$ to form the IV estimator. If γ_1 is near zero, then $\hat{\gamma}_1$ will usually be near zero too.

¹The IV estimator is not unbiased in general, even though its two constituent parts are unbiased estimators of the respective parameters, because $E[A/B] \neq E[A]/E[B]$ in general.

When our denominator is close to zero, small changes in the numerator correspond to wild differences in the resulting ratio. Therefore, the IV estimator with a weak instrument is liable to have huge standard errors. If no better instrument is available, whether it is better to use IV with a weak instrument (consistent but inefficient) or to use OLS (inconsistent) depends on a number of factors, including the sample size and the overall strength of the unobserved confounding.

“Assignment”

There is no problem set this week, but I encourage you to undertake the following exercise:

- Read the three empirical papers assigned on the syllabus for March 24 (Angrist and Krueger 1991; Acemoglu, Johnson, and Robinson 2001; Miguel, Satyanath, and Sergenti 2004).
- Identify in each paper the instrument used, and explain why it meets (or doesn't!) each of the three conditions for an instrumental variable.

References

Acemoglu, Daron, Simon Johnson, and James A Robinson. 2001. “The Colonial Origins of Comparative Development: An Empirical Investigation.” *The American Economic Review* 91 (5): 1369–1401.

Angrist, Joshua D, and Alan B Krueger. 1991. “Does Compulsory School Attendance Affect Schooling and Earnings?” *The Quarterly Journal of Economics* 106 (4): 979–1014.

Angrist, Joshua D, and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. An Empiricist's Companion. Princeton University Press.

———. 2015. *Mastering 'Metrics*. The Path from Cause to Effect. Princeton University Press.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti. 2004. “Economic Shocks and Civil Conflict: An Instrumental Variables Approach.” *Journal of Political Economy* 112 (4): 725–53.