Econometrics II: Econometric Modelling

Jürgen Meinecke

Research School of Economics, Australian National University

28 September, 2018

Assignment 1

If you want to collect your assignment you will need to see me (just come by or send me an e-mail for an appointment)

If you want to be able to get a remark on an uncollected assignment, please collect it from me before 5 October

Assignments collected later than that cannot get a remark!

Roadmap

Some new terminology is unavoidable:

- A person i is given a treatment $X_i \in \{0, 1\}$
- Treatment will have an effect on the person's outcome Y_i
- The outcome therefore is a function of the treatment
- Only two outcomes possible per person: $Y_i(0)$ and $Y_i(1)$
- The individual treatment effect (ITE) is: $Y_i(1) - Y_i(0)$
- This is the effect of the treatment on person i
- The individual treatment effect is never observed because we can only know $Y_i(0)$ (the outcome if the person did not receive the treatment) or $Y_i(1)$ (outcome if the person did receive the treatment) but never both

Example: blood pressure medication

- Let's say I am person i
- A medical researcher randomly assigns me to one of two groups: control group or treatment group
- Subjects in the control group are made to believe that they do get the actual medication but in fact do NOT get the medication (they get an ineffective fake pill)
- Subjects in the treatment group are made to believe that they do get the actual medication and in fact they do
- Let's say the medical researcher assigns me to the treatment group

- At the end of the trial, the researcher observes $Y_i(1)$ (my outcome given that I received the actual medication)
- The researcher cannot know $Y_i(0)$ because I did receive the treatment
- It therefore is impossible to calculate $Y_i(1) - Y_i(0)$
- The outcome $Y_i(1)$ is called my *factual outcome*
- The outcome $Y_i(0)$ is called my *counterfactual outcome*
- (Had I been assigned to the control group instead, then $Y_i(0)$ would have been observed, and my factual outcome would have been $Y_i(0)$ while my counterfactual outcome would have been $Y_i(1)$)
- Counterfactual outcomes are unknown

How can we solve the problem of the missing counterfactual?

One idea would be to find an otherwise identical person $j \neq i$ who did not receive the treatment

For that person, the observed factual would be $Y_j(0)$

The individual treatment effect would be $Y_i(1) - Y_j(0)$

Problem: no two persons are ever "otherwise identical"

It is practically impossible to find two people who are the same

What else can be done?

We need to say good-bye to the idea of learning about the *individual* treatment effect

Instead, a more realistic goal is to learn about the *average* treatment effect

Do you remember its definition?

The ATE is given by $\mathrm{E}[Y_i(1) - Y_i(0)]$

It is the effect on the average person in the population

How can we estimate the ATE using a regression model?

Roadmap

In the data you observe (X_i, Y_i) where

- treatment X_i is
  - randomly assigned to person i
  - equal to 0 or 1
- observed outcome is given by

$$\begin{aligned} Y_i &:= Y_i(1) \cdot X_i + Y_i(0) \cdot (1 - X_i) \\ &= [Y_i(1) - Y_i(0)]\, X_i + Y_i(0) \\ &= \mathrm{E}[Y_i(0)] + [Y_i(1) - Y_i(0)]\, X_i + (Y_i(0) - \mathrm{E}[Y_i(0)]) \end{aligned}$$

The last line has a familiar appearance…

$$Y_i = \underbrace{\mathrm{E}[Y_i(0)]}_{\beta_0} + \underbrace{[Y_i(1) - Y_i(0)]}_{\beta_{1i}} X_i + \underbrace{(Y_i(0) - \mathrm{E}[Y_i(0)])}_{u_i}$$

$$= \beta_0 + \beta_{1i} X_i + u_i$$

Careful: this is NOT our standard OLS equation!

The slope coefficient is *indidvidual specific* (it has an i-subscript)

In this equation here, the slope coefficient is allowed to differ by individuals

The coefficient β_{1i} is the individual treatment effect

It is impossible to estimate n slope coefficients from the sample data

Let's turns this last equation into a model which allows us to estimate the ATE instead of the ITE

$$
\begin{aligned}
Y_i &= \mathrm{E}[Y_i(0)] + [Y_i(1) - Y_i(0)]\, X_i + (Y_i(0) - \mathrm{E}[Y_i(0)]) \\
&= \mathrm{E}[Y_i(0)] + [Y_i(1) - Y_i(0)]\, X_i + (Y_i(0) - \mathrm{E}[Y_i(0)]) \\
&\quad + \mathrm{E}[Y_i(1) - Y_i(0)] \cdot X_i - \mathrm{E}[Y_i(1) - Y_i(0)] \cdot X_i \\
&= \mathrm{E}[Y_i(0)] + \mathrm{E}[Y_i(1) - Y_i(0)] \cdot X_i \\
&\quad + \Big((Y_i(0) - \mathrm{E}[Y_i(0)]) \\
&\quad + \quad [Y_i(1) - Y_i(0)]\, X_i - \mathrm{E}[Y_i(1) - Y_i(0)] \cdot X_i \Big) \\
&= \beta_0 + \beta_1 \cdot X_i + u_i,
\end{aligned}
$$

where $\beta_0 := \mathrm{E}[Y_i(0)]$ and $\beta_1 := \mathrm{E}[Y_i(1) - Y_i(0)]$ and everything in big parentheses is u_i

Notice that $\beta_1 := \mathrm{E}[Y_i(1) - Y_i(0)]$ is equal to the ATE

Putting things together, we get:

Theorem (OLS in Randomized Controlled Trial)

Suppose you have available data (X_i, Y_i) from a randomized controlled trial. In particular, X_i is a randomly assigned treatment dummy variable. Then the OLS estimator of β_1 in the model $Y_i = \beta_0 + \beta_1 X_i + u_i$ is an estimator of the average treatment effect $E[Y_i(1) - Y_i(0)]$.

For this theorem to be valid, however, we still need to establish that $\mathrm{E}[u_i|X_i] = 0$

Let's quickly do this...

$$
\begin{aligned}
\mathrm{E}[u_i|X_i] &= \mathrm{E}\Big[(Y_i(0) - \mathrm{E}[Y_i(0)]) \\
&\qquad + [Y_i(1) - Y_i(0)]\, X_i - \mathrm{E}[Y_i(1) - Y_i(0)] \cdot X_i | X_i\Big] \\
&= \mathrm{E}[Y_i(0)|X_i] - \mathrm{E}[\mathrm{E}[Y_i(0)]|X_i] \\
&\qquad + \mathrm{E}[Y_i(1)|X_i] \cdot X_i - \mathrm{E}[Y_i(0)|X_i] \cdot X_i \\
&\qquad - \mathrm{E}[\mathrm{E}[Y_i(1)]|X_i] \cdot X_i + \mathrm{E}[\mathrm{E}[Y_i(0)]|X_i] \cdot X_i \\
&= \mathrm{E}[Y_i(0)|X_i] - \mathrm{E}[Y_i(0)] \\
&\qquad + \mathrm{E}[Y_i(1)|X_i] \cdot X_i - \mathrm{E}[Y_i(0)|X_i] \cdot X_i \\
&\qquad - \mathrm{E}[Y_i(1)] \cdot X_i + \mathrm{E}[Y_i(0)] \cdot X_i \\
&= \mathrm{E}[Y_i(0)] - \mathrm{E}[Y_i(0)] \\
&\qquad + \mathrm{E}[Y_i(1)] \cdot X_i - \mathrm{E}[Y_i(0)] \cdot X_i \\
&\qquad - \mathrm{E}[Y_i(1)] \cdot X_i + \mathrm{E}[Y_i(0)] \cdot X_i \\
&= 0,
\end{aligned}
$$

where

- the first equality follows by definition
- the second equality follows by breaking up all individual terms (using the fact that the expected value of the sum is equal to the sum of the expected values)
- the third equality follows because the expected value of an expected value is equal to the (inner) expected value
- the last equality follows because X_i is assigned randomly

What if you have additional regressors available in the data?

Let X_i = treatment variable and W_i = control variable

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i$$

You don't actually have to include W_i in the regression

Could simply look at it as part of the error term

If X_i is randomly assigned, then it cannot be correlated with W_i, so OLS Assumption 1 still satisfied

But you should include W_i nevertheless...

Two reasons to include W in a regression analysis of the effect of a randomly assigned treatment:

1. reduces the error variance and can narrow standard errors
2. if the probability of treatment assignment depends on W_i, so that X_i is randomly assigned given W_i, then omitting W_i can lead to ovb, but including it eliminates that OV bias

Example

- men ($W_i = 0$) and women ($W_i = 1$) are randomly assigned to a course on table manners (X_i)
- but women are assigned with a higher probability than men
- Suppose women have, on average, better table manners than men prior to the course
- Then even if the course has no effect, the treatment group will have better post-course table manners than the control group because the treatment group has a higher fraction of women than the control group
- That is, the OLS estimator of β_1 in the regression of Y_i on X_i will have ovb

$$Y_i = \beta_0 + \beta_1 X_i + \beta_2 W_i + u_i$$

- In this example, X_i is randomly assigned, given W_i, so $E(u_i|X_i, W_i) = E(u_i|W_i)$.
- In words, among women, treatment is randomly assigned, so among women, the error term is independent of X_i so, among women, its mean doesn't depend on X_i
- Same is true among men
- Thus if randomization is based on covariates, conditional mean independence holds
- so that once W_i is included in the regression the OLS estimator is unbiased

Roadmap

If you can't effectively randomize treatment, an alternative is to randomize *eligibility* for treatment instead

Last week we learned that in this case you should use TSLS estimation

But does TSLS estimate the ATE?

Like before, let's study a model in which treatment effect is heterogeneous

$$Y_i = \beta_{0i} + \beta_{1i} X_i + u_i \qquad \text{(equation of interest)}$$

$$X_i = \pi_{0i} + \pi_{1i} Z_i + v_i \qquad \text{(first stage of TSLS)}$$

Using some math, it can be shown that

$$\hat{\beta}_1^{TSLS} \xrightarrow{p} \frac{\mathrm{E}[\beta_{1i} \cdot \pi_{1i}]}{\mathrm{E}[\pi_{1i}]}$$
$$\neq \mathrm{E}[\beta_{1i}] =: \mathrm{ATE}$$

Two results here:

- TSLS estimator does not converge to the ATE (bad news)
- Instead it converges to $\mathrm{E}[\beta_{1i} \cdot \pi_{1i}] / \mathrm{E}[\pi_{1i}]$ (looks complicated)

Here is a useful way to think about $\frac{\mathrm{E}[\beta_{1i} \cdot \pi_{1i}]}{\mathrm{E}[\pi_{1i}]}$:

- interpret $\frac{\pi_{1i}}{\mathrm{E}[\pi_{1i}]}$ as weights
- then the rhs is equal to the expected value of β_{1i} adjusted for these weights
- in other words: the rhs is the weighted average of β_{1i}
- ideally, we would not want any weights in there (because we are after the ATE, which is the simple average)
- some intuition for the weights: when π_{1i} is large relative to $\mathrm{E}[\pi_{1i}]$ then the weight is large; therefore people with large π_{1i} influence the TSLS estimator more (their Z_i have a strong impact on X_i)

Putting things together: TSLS estimates the causal effect for those individuals for whom Z_i is most influential
(those with large π_{1i})

The probability limit $\frac{\mathrm{E}[\beta_{1i} \cdot \pi_{1i}]}{\mathrm{E}[\pi_{1i}]}$ is called the
local average treatment effect (LATE)

The LATE can be understood as the ATE for the subpopulation whose treatment X_i is most heavily influenced by the instrument Z_i

LATE is an ATE only for this peculiar ("local") subpopulation; it is not equal to the ATE in the population

How does LATE relate to ATE?

-

$$
\begin{aligned}
\text{LATE} &:= \frac{\mathrm{E}[\beta_{1i} \cdot \pi_{1i}]}{\mathrm{E}[\pi_{1i}]} \\
&= \frac{\mathrm{E}[\beta_{1i}]\mathrm{E}[\pi_{1i}] + \mathrm{Cov}(\beta_{1i}, \pi_{1i})}{\mathrm{E}[\pi_{1i}]} \\
&= \mathrm{E}[\beta_{1i}] + \frac{\mathrm{Cov}(\beta_{1i}, \pi_{1i})}{\mathrm{E}[\pi_{1i}]} \\
&= \text{ATE} + \frac{\mathrm{Cov}(\beta_{1i}, \pi_{1i})}{\mathrm{E}[\pi_{1i}]}
\end{aligned}
$$

(Second equality holds b/c $\mathrm{E}[X \cdot Y] = \mathrm{E}[X] \cdot \mathrm{E}[Y] + \mathrm{Cov}(X, Y)$ as you surely remember from STAT1008)

In words: LATE equals ATE plus "some stuff"

But what exactly is "some stuff"?

It is the covariance between the two individual-specific parameters β_{1i} and π_{1i}

If the treatment effect β_{1i} tends to be large for individuals for whom the effect of the instrument π_{1i} is also large, then $\text{Cov}(\beta_{1i}, \pi_{1i}) > 0$ and therefore LATE $>$ ATE
(supposing $\text{E}[\pi_{1i}] > 0$)

On the other hand, if the treatment effect β_{1i} tends to be small for individuals for whom the effect of the instrument π_{1i} is also large, then $\text{Cov}(\beta_{1i}, \pi_{1i}) < 0$ and therefore LATE $<$ ATE

When does TSLS estimate the ATE?

- If $\beta_{1i} = \beta_1$ (no heterogeneity in equation of interest)
- If $\pi_{1i} = \pi_1$ (no heterogeneity in first stage equation)
- If β_{1i} and π_{1i} vary but are independently distributed

But these three are unrealistic

In general, $\hat{\beta}_1^{TSLS}$ does not estimate ATE

Whether this is important depends on the application ...

Example: Cardiac catheterization

Recall the example from lecture 6

TSLS setup:

$$\text{Main equation:} \quad SurvivalDays_i = \beta_0 + \beta_{1i} CardCath_i + u_i$$

$$\text{First stage:} \quad CardCath_i = \pi_0 + \pi_{1i} Distance_i + v_i$$

The coefficient β_1 captures the effect of cardiac catheterization on survial outcomes

Recall: TSLS estimate of β_1 small and insignificant

That seemed puzzling: why would CC not be beneficial?

Problem: the TSLS estimator is not estimating the ATE, instead it is estimating LATE

LATE is the ATE for the subpopulation for which the instrument Z_i is particularly influential for the treatment X_i

Applied to this example: TSLS estimates the causal effect for those whose value of X_i is most heavily influenced by Z_i

Who are they?
Patients whose proximity to CC hospital affects their chances of receiving CC

Patients whose proximity to CC hospital affects their chances of receiving CC

- presume that this applies to patients who are on average healthier
- if they live far away from a CC hospital, they will not receive CC but some alternative treatment
- if they happen to live nearby a CC hospital, they will receive CC
- for them, we would not expect to find a positive CC effect because they would do equally well if they did not receive CC

That story is consistent with the actual research finding of a small and insignificant TSLS estimate

Problem: it's all speculative, but that's the trouble with TSLS

Roadmap

RCT are generally considered the gold standard of internal validity

A well designed and executed RCT is the best way of learning about a causal effect

However, even here there are pitfalls

Let's use an example to illustrate the potential problems with RCT:

- You are interested if a particular job training program helps unemployed people in Canberra to get back into work
- You randomly choose 100 unemployed people
- You randomly choose a subset of 50 and make the eligible for the job training program
- After 3 years you compare the outcomes between the two

1. Failure to randomize
   (or imperfect randomization)
   - Example: We make people eligible on a first-come, first-serve basis (the first 50 people qualify for the job training program); late-comers will be controls
2. Failure to follow treatment protocol
   (or "partial compliance")
   - some people in the control group receive job training from a different institution
   - some people in the treatment group do not actually attend the job training sessions (and we fail to notice)

3. Attrition
   (some subjects drop out)
   - Example: Some subjects move to other cities for unexpected family reasons (Would that be a problem?)
   - But what if some successful trainees move to Sydney to find better jobs (and you fail to follow up with them)
   - Similar to sample selection bias
4. Experimental effects
   - experimenter bias (conscious or subconscious): treatment X is associated with "extra effort" or "extra care,"
   - subject behavior might be affected by being in an experiment ("Hawthorne effect")

Remember: threats to internal validity always mean that OLS Assumption 1 is in danger

If OLS Assumption 1 fails, then we are not estimating the correct causal effect

What are we estimating instead? We have no idea!

We always strive for internal validity, or else everything we do is useless

Roadmap

1. Nonrepresentative sample
    - We use our Canberra job training RCT to make inferences on a nation-wide basis
    - Canberra's unemployed are not representative of the nation's unemployed and Canberra's economy also is not
2. Nonrepresentative "treatment" (that is, program or policy)
    - Suppose we have learned that the job training program is effective (we estimated a large positive causal effect)
    - Canberra politicians may be interested in rolling out a job training program on a larger basis (as their signature labor market policy, say)
    - When rolling out the program, it needs to be implemented in exactly the same way the RCT program was done; otherwise we cannot strictly say that the lessons from the RCT are valid for a different program

3. General equilibrium effects

- If we were to scale up the program, then it is conceivable that it will "crowd out" other similar programs
- Again, if Canberra policy makers decide to introduce a job training program just like the one conducted in the RCT, then employers may decide not to offer their own training programs

Roadmap

RCT can be regarded as planned or "unnatural" experiments; they are intentionally created by researchers

A *natural experiment* or *quasi experiment* in most cases is not the result of deliberate planning for the purpose of studying treatment effects

Natural experiments are not planned; they come about by "nature"

The split between treatment and control group happens (almost) unintentionally and are not the result of an explicit RCT

Nevertheless, natural experiments create treatment and control groups and that division is regarded "as if" it was randomly assigned

Natural experiments are the closest we can get to an RCT without actually doing an RCT

Two types of natural experiments

1. A variable Z_i which influences receipt of treatment X_i is "as if" randomly assigned
2. Treatment X_i is "as if" randomly assigned

Examples: Z_i "as if" randomly assigned

Effect on longevity of cardiac catheterization

- Y_i = longevity
- X_i = received CC surgery
- Z_i = distance to nearest CC hospital
- In lecture 6 we said that Z_i can be viewed as random

Effect of Studying on grades

- Y_i = student grades
- X_i = amount of studying (hours)
- Z_i = dummy equal one if roommate brought video game
- In lecture 4 we said that Z_i can be viewed as random

Effect of economic growth on civil conflict

- Y_i = dummy equal one if civil conflict
- X_i = change in real GDP
- Z_i = rainfall
- In lecture 4 we said that Z_i can be viewed as random

The last example, in particular, shows why we speak of "natural experiments": rainfall can indeed be thought of as forced by nature

Example: X_i "as if" randomly assigned

Jha, N. (2015),
"*Late Start with Extra Schooling The Effect of Increase in School Entry Age and Preschool Provision*", Economic Record

Research question:
Does school entry age affect student outcomes?

- Y_i = NAPLAN outcomes
- X_i = primary school entry age

Exploit Queensland law change in 2008:

- Before 2008, children needed to be 6 years of age by 30 December of the year they enroll in grade 1
- Since 2008, the cutoff was moved to 30 June

As result, the compulsory school starting age increased by up to 6 months

This effectively creates four groups:

- children in Queensland admitted before 2008
- children in Queensland admitted since 2008 (treatment)
- children elsewhere admitted before 2008
- children elsewhere admitted since 2008

Whether or not a child ends up in one of these groups can be regarded “as if” randomly assigned

Other states have not changed their laws during that time

Estimate following regression:

$$\Delta Y_i = \beta_0 + \beta_1 Q_i + u_i,$$

where

- ΔY_i: difference in average test scores at school i between 2006 and 2010
- Q_i is a dummy equal to one if child from Queensland

The coefficient β_1 captures the average treatment effect of increased school entry age on testscores

This is an example of a *difference-in-differences* estimation:

- Compare the testscores of students from Queensland before and after 2008
- Compare the testscores of students from elsewhere before and after 2008
- Then compare these two differences to each other