



Stockholm
University

Pär Nyman

Methods in Political Science
25 May 2026

Course review

Introduction

- The exam
- Central topics during the course
- Other questions

Your suggestions were spread widely, so I will not be able to address them all.

The exam

Practical info

- It will be a digital on-campus exam using Inspera.
- You are allowed to bring snacks (in transparent bags) and drinks, as well as a calculator (not required).
- Always bring a physical ID.
- Room information will be available two days before the exam, at the latest.

The exam

Format

- The exam has a maximum score of 30 points.
- One block of 6 multiple choice questions.
- Three hypothetical studies with four sub-questions each.
- Most questions are about (not specific to this exam):
 - Describing concepts.
 - Interpreting results.
 - Discussing methodological choices.
 - Evaluating assumptions.
 - Understanding implications.

The exam

Content: preparing for what is on the exam

- There are typically 1–2 questions from each lecture.
- My ambition: test a general ability to identify key problems.
 - I think we have been practicing this for two months.
 - If I succeed, cramming will be harder and less useful.
 - Yet exams always test more specific knowledge about concepts and assumptions.
- I suspect that the most time-effective ways to study are:
 - Answer the three previous exams on Athena.
 - Repeat central concepts.
 - Go over the lecture slides once more.

The exam

Content: what is not on the exam

- No questions about Stata.
- No questions about the literature
 - Many answers are found in the literature.
 - But not exclusively in the literature.
- We do not require explicit use of equations.
 - We may use related language (e.g., x , y and z).
 - Understanding an equation is often the best way to understand a methodological concept.
 - For me, equations and regressions are not really separable.

Central topics

Main themes in the course (non-exhaustive)

- Sampling errors.
- Exogeneity.
- Design-based: experiments, RD, IV, DiD.
- Multiple regression.
- Other topics.

Questions welcome, but we can't go too much into details (office hours!). I can only repeat and place into context.

Central topics

Today's example

- What is the economic return to education?
- We could regress earnings on the years of education and get an estimate of how income is affected by one additional schooling year
- That estimate would have sampling error and bias.

Sampling errors

- Comparing the income of two individuals, one with high and one with low education
- Chances are that our sample would not be representative – the difference between them would not be the same as the population difference.
- This would be a sampling error.

Sampling errors

- The larger our sample, the smaller this error would be (in expectation), but it would never be zero.
- Small standard errors indicate good precision.
- Precision does not say anything about bias.

Sampling errors

- The standard error estimates the expected size of the sampling errors.
- 95 percent of samples would provide a confidence interval that includes the true income difference (effect if unbiased).
- The p-value would tell us the probability of obtaining an income difference at least this large if income is unrelated to education length.
- Small p-value equals statistically significant – we reject the null hypothesis that earnings are unrelated to education.

Sampling errors

Words of warning

- Standard errors are estimates. Cluster them to adjust for observations being correlated within specified groups.
- Consider the number of tests. About 1 in 20 p-values will be smaller than 0.05 when the null hypothesis is true. Is the significant effect we are looking at one of those?

Exogeneity assumption

- Confidence intervals and p-values handle sampling errors. They say nothing about possible bias.
- Violating the exogeneity assumption is the most common source of serious bias:
 - Omitted variables (confounders).
 - Reverse causation.
- What if income affects education? Or if both are affected by class, grit or cognitive ability?

Exogeneity assumption

Many techniques for meeting the exogeneity assumption:

- Randomization: income or grit does not affect the dice roll
- As-if random: x is arguably unaffected by income or grit
- RD: just passed vs just failed entry test
- IV: something else is exogenous
- DiD: the parallel trends assumption provides a counterfactual
- Control variables: control for class and cognitive ability

Experiments

- Randomization ensures that potential outcomes are the same in treatment and control group (in expectation).
- They still differ because of sampling errors.
- A mean comparison or bivariate regression is enough for unbiased estimates.

Balance and placebo tests

- The idea of experiments, RD, IV and DiD is to identify a (plausibly) exogenous variation.
- That argument is stronger if we can show that the treatment and control group (for example) are similar on observable characteristics (balance test) or that the model does not find any effects on outcomes that should not be affected (placebo tests).

Experiments

Assumptions (condensed definitions)

- Random assignment into treatment and control group.
- Compliance: treatment corresponds to assignment.
- Excludability: assignment only affects through treatment.
- Non-interference: treatment does not affect others.

Experiments

Adjusting for non-compliers

- The intent-to-treat (ITT) effect is the effect of assignment.
- With one-sided non-compliance it is in expectation equal to the share of compliers times the treatment effect.
- We can therefore estimate the treatment effect on the compliers (LATE/CACE) by dividing the estimated ITT with the share of compliers.



Controlled and natural experiments

- In a controlled experiment, the researcher assigns treatment through randomization.
- In a natural experiment, assignment is as if random.
 - Some are analyzed as if there were randomization.
 - Others require specific techniques such as RD or IV.
 - DiD designs are sometimes seen as natural experiments.

Regression Discontinuity

- Bleemer & Mehta 2022 analyze those close to a GPA threshold for majoring in economics.
- Those just below and above will have similar grit and ability, but the latter are more likely to major in economics.
- Is there a jump at the threshold? Evaluate visually as well as model specification (reasonable bw and polynomial?).
- Key assumption: no sorting around the threshold.
- This is a fuzzy RD: the probability of treatment does not go from 0 to 1.

Regression Discontinuity



Instrumental variables

When and why

- We are interested in the effect of x on y but the exogeneity assumption is violated (very common).
- And there is an instrument that meets this assumption and that affects our x (rare).

If true, we have identified some exogenous variation in x .

Instrumental variables

- Compulsory time in school differs depending on birth month (start at different ages, allowed to drop out at the same).
- Relevance? Look at the first stage regression's F-test.
- Exogeneity: Is birth month exogenous to class, grit and cognitive ability?
- Exclusion: Is years of schooling the only channel through which birth month affects earnings?

Difference-in-Differences

- Some event has affected one group but not another.
- The variation is not random (so comparing means is not sufficient), but the two groups are similar enough to develop the same way over time.
- Key assumption: parallel trends in potential outcomes – treated units would have changed in the same way as controls in the absence of treatment.
- Estimation: Change among treated minus change among controls (i.e., a difference in differences).

Diff-in-Diff: Income and education

Made-up example about mandatory schooling age and earnings

In State A, a new law is passed in 2010 raising the school-leaving age from 16 to 18. State B, a similar state, does not change its education policy.

Group	Income (age 25)
State A, pre-policy	30,000 SEK
State A, post-policy	33,000 SEK
State B, pre-policy	30,000 SEK
State B, post-policy	31,000 SEK

$$(33,000 - 30,000) - (31,000 - 30,000) = 3,000 - 1,000 = 2,000$$

Multiple regression

- **Confounders** affect both x and y . Unless we include them as **control variables** we get **omitted variable bias**.
- **Mediators** are affected by x and affect y . They are **bad controls**, but include them if you do a mediation analysis.
- Moderators condition the effect of x . We can analyze conditional effects through **split samples** and **interactions**.

We assume causal direction based on theory or common sense.

Multiple regression

- When estimating slope coefficients, we only use the unique variation in each independent variable.
- In this case we would use the variation in education that is unrelated to our measures of class, grit and cognitive ability.
- We cannot (easily) control for reverse causation or things we cannot measure.
- What happens if we control for occupation?
 - Mediator (include it to see if effect goes through occupation)
 - Bad control (we remove a causal effect without realizing)

Other topics

Regressing monthly income on years of education

	(1)	(2)	(3)	(4)
Years of education	1651.32*** (91.27)	1184.73*** (104.58)	421.86*** (78.44)	1196.54*** (102.31)
Parents' highest education		617.49*** (68.92)	208.35*** (54.17)	610.22*** (67.84)
Parents' average income (centered)		0.08*** (0.01)	0.03*** (0.01)	0.10*** (0.01)
Years of education $\times$ parents' income				-0.04*** (0.01)
Average income in occupation			0.74*** (0.02)	
Constant	9817.64*** (1148.39)	6234.18*** (1261.72)	3412.57*** (897.45)	6218.39*** (1248.66)
Observations	2500	2500	2500	2500
R^2	0.18	0.29	0.62	0.31
Standard error	8423.71	7812.46	5594.83	7698.41

Other topics

Column (1) estimates the following model:

$$\widehat{\text{Monthly income}} = 9817.64 + 1651.32 \times \text{Years of education}$$

- Each additional year of education is associated with about 1,651 SEK higher monthly income.
- The standard error for that coefficient is 91.27, suggesting that the estimate is quite precise.
- Years of education explains about 18 percent of the variation in monthly income.
- A typical prediction error from this simple model is about 8,424 SEK per month.

Other topics

Keep in mind when interpreting slope coefficients

- How are the variables coded? Crucial both for informative interpretations and fair comparisons.
- What is the causal direction between x and the control? Occupational income would be a mediator/bad control.
- What happened to the coefficient's standard error and the number of observations? Not relevant here.

Other topics

Interpreting interaction models

- If we suspect that the effect of x is conditioned on z we could estimate this equation: $\hat{y} = a_1 + b_1x + b_2xz + b_3z$
- b_1 is the effect of x when $z = 0$.
- b_2 is how much the effect of x increases if z increases by one.
- b_3 is the effect of z when $x = 0$.

The negative interaction coefficient indicated that the effect decreases with parental income – education is compensatory.

Other topics

Working with panel data

- Correlated residuals bias standard errors. Cluster!
- Unit fixed effects use *within-unit* variation:
 - Controls for all time-invariant unit characteristics.
 - Removes a lot of variation and often hurts precision.
- Time fixed effects control for shocks common to all units.
- Not all effects are immediate. Think about dynamics:
 - Lag the dependent variable to model persistence.
 - Lag the independent variable to model a delayed effect.

Other topics

Why not choose the most flexible specification?

Adding controls, allowing non-linearities, sample restrictions, etc. may reduce bias and relax assumptions. But often a trade-off:

- Statistical precision and overfitting.
- Interpretability and presentation.
- Conventions can be convincing.

Conclusion: data and purpose should determine model choice.

Concluding

Remaining questions?

Concluding

- Office hours are Wednesday, 10:00–11:45, in F271.
- Your questions tend to focus on previous exams.
- Thanks to everyone. It has been fun!