

Causal Inference

Mondays 11:00 – 13:00 & Fridays 13:00 – 15:00

Location: Seminar Room 2

Administrative contact: [Marco Pistolesi](#)

Instructor: [Sascha Riaz](#)

Email: sascha.riaz@eui.eu

Teaching Assistants:

[Martín Alberdi](#) (martin.alberdi@eui.eu)

[Ludwig Schulze](#) (ludwig.schulze@eui.eu)

Course Description: This course provides an introduction to causal inference in the social sciences, equipping students with the skills to critique methods used in contemporary academic work and apply these methods in their research. It begins with an overview of causality in the social sciences, covering the potential outcomes framework and randomized experiments. The second part of the course delves into the most widely used approaches for drawing causal inferences from observational data, including “selection on observables” (subclassification, matching), instrumental variables, difference-in-differences, and the regression discontinuity design.

Class meetings: The course meets twice per week over a period of five weeks. The first class meeting is on January 5, and the last will be on February 6. The course is divided into a lecture component and a lab component. Please bring your own laptop to the labs and ensure that you have R and RStudio installed and running. Instructions [here](#).

Officially, the EUI allows you to miss two class meetings and still receive credit for the course, but I urge you to attend *all* sessions except in cases of severe illness.

Readings: The main textbook for the course is “[Causal Inference: The Mixtape](#)” by Scott Cunningham. The book is available for free online. If you want to read something more advanced on the same topics, see Angrist and Pischke’s “[Mostly Harmless Econometrics](#)” or “[Causal Inference for Statistics, Social, and Biomedical Sciences](#)” by Guido Imbens and Donald Rubin. Slides and lab materials will be provided on Brightspace.

Course Prerequisites: Students are expected to have a basic understanding of probability and statistics. If in doubt, review [chapter 2](#) of the textbook.

The software used throughout the course and in the lab sessions is R/RStudio. You should be familiar with using R Markdown to generate HTML/PDF documents. If you don't, you can find a tutorial [here](#) or elsewhere online.

Getting help: While you are welcome to reach out to the teaching assistants with any coding questions, you are expected to make every effort to resolve coding issues on your own before seeking assistance. Please utilize all available resources and attempt to troubleshoot independently. Only if absolutely necessary should you contact the teaching assistants via email. You may also reach out to the TA to schedule office hour appointments if needed.

Requirements: The course will be graded pass/fail – no letter or numerical grade will be given. Beyond regular attendance (see above), the following is required to pass the course:

1. **Critical review:** You will write a critique (1,000 words max) on a paper that (i) aims to establish a causal effect and (ii) was published in one of the top political science or sociology journals. Your review should focus on critically assessing the identification strategy employed by the authors. Clearly describe the empirical methods they use to establish causality, identify the identification assumptions required for their strategy to be valid, and discuss potential violations of these assumptions. Additionally, evaluate any evidence the authors provide to support their assumptions and discuss whether you find it convincing. If you believe further tests are necessary, you may attempt to implement them using the authors' replication files that should be available through the journal websites. While critical reviews can take many forms, your goal for this assignment is to engage with and critically assess the authors' causal identification strategy, not the substantive merits of their theoretical argument or contribution to the literature. The critical review is due on February 15 and must be submitted on Brightspace. Email submissions will be ignored. Late submissions will not be accepted.
2. **Final Exam:** The final take-home exam will be posted on Brightspace on Feb. 16 and will include a mix of conceptual and coding questions that require you to implement the methods we have covered throughout the course. The exam is due on March 15 and must be submitted on Brightspace; email submissions will not be accepted. Your submission must be in PDF or HTML format, generated using RMarkdown. For a tutorial, see [here](#). Late submissions will not be accepted, and collaboration on the take-home exam is not permitted.

Topics and Readings

Jan. 5, 2026: Introduction, Causality, Potential Outcomes Framework

- *Mixtape:* Chapters 1 and 4

Jan. 9, 2026: Randomized Experiments

- Angrist, J.D. and Pischke, J.S., 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Chapter 2.

Jan. 12, 2026: Lab Session I

Jan. 16, 2026: Selection on Observables: Matching and Subclassification

- *Mixtape:* Chapter 5

Jan. 19, 2026: Instrumental Variables

- *Mixtape:* Chapter 7

Jan. 23, 2026: Lab Session II

Jan. 26, 2026: Difference-in-Differences

- *Mixtape:* Chapter 9

Jan. 30, 2026: Lab Session III

Feb. 2, 2026: Regression Discontinuity Designs

- *Mixtape:* Chapter 6

Feb. 6, 2026: Lab Session IV