

Introduction to Causal Inference

Course Code: IEDA 4000D

No. of Credits: 3

Prerequisites: IEDA 2520, IEDA 2540 or other equivalent courses.

Schedule:

Lectures: 13:30 – 14:50, Wednesdays and Fridays. Room 4579

Instructor:

Nian Si. Email: niansi@ust.hk, Office: Room 5559F.

TAs: TBD

Course webpage: <http://canvas.ust.hk>

Lecture notes, tutorials, assignments, solutions, and other announcements will be posted on this site. Please check this site regularly.

Course Description: This course provides an introduction to the fundamental concepts and statistical methods used in causal inference. Understanding causal relationships is essential in fields such as economics, digital marketing, healthcare, and policy evaluation. For example, students will explore how to assess the impact of a new product design on user experience, a price change on sales, or a welfare incentive on employee performance. The course equips students with essential tools for revealing causal effects and conducting hypothesis tests using both experimental and observational data. Students will also gain hands-on experience with real-world data sets. Special topics may include A/B testing on digital platforms and cutting-edge methods that connect causal inference with machine learning.

Course Learning Outcomes: This course is introduced to equip students with essential tools to understand and uncover cause-and-effect relationships, a crucial skill across fields like economics, marketing, and healthcare. In an era driven by data, the ability to move beyond correlation and make accurate causal inferences is highly valued. By learning methods like A/B testing, randomized controlled trials, and integrating causal inference with machine learning, students will gain practical, real-world skills that are foundational for data-driven decision-making and advanced studies in various disciplines. Upon the completion of this course, you will be able to:

- Understand key concepts in causal inference;
- Apply statistical methods for causal analysis;
- Use software tools for causal analysis.

Textbook: There is no required textbook for this course. All required class materials will be available on Canvas. However, if you want to pursue more advanced topics, here are three reference books that may be helpful:

- **Main:** Imbens, Guido W., and Donald B. Rubin. Causal inference in statistics, social, and biomedical sciences. Cambridge university press, 2015.
- **Optional:** Ding, Peng. A first course in causal inference. Chapman & Hall, 2024.
- **Optional:** Pearl, Judea. Causality. Cambridge university press, 2009.

Software and Coding:

- The class will use Python, a powerful and easy-to-use programming language.
- The class does assume prior programming/mathematical knowledge at the level of what's been introduced in IEDA 2520 and 2540. We will provide an effective tutorial so that those without such prior knowledge will be able to catch up by self-learning.
- The class is not a programming class. Programming is a tool to implement your idea and method. Like any programming language, Python is best learned through practice. Despite the steep learning curve, students can become proficient in a few weeks with reasonable effort.

Tentative Course Outline:

1. Randomized Controlled Trials
2. Observational Studies and Propensity Scores
3. Regression Discontinuity Designs
4. Panel data
5. Instrumental variables
6. Structural equation modeling
7. Causal inference in technology companies
8. Special Topics

Grading: project (30%), midterm (40%), homework (30%), participation (2%, bonus).

- The midterm exam will be held in class on 10th April, 2026 (tentatively). During the exam, you will not be allowed to use books or notes. However, you're allowed a one-page, double-sided cheat sheet.
- For the final course project, you are required to apply the methodologies covered in class to a real-world data analysis. The written report should consist text presentation plus supporting analysis, charts and graphs (up to 8 pages). You also need to submit your code for the project. Project presentation is scheduled during the last week. Precisely, on 6th and 8th May, 2026.
- You may discuss your homework with instructor, TA, fellow students, and others. However, you are expected to write up your solutions on your own.

Homework Submission Policy:

Due Date: All homework assignments must be submitted by 11:59 PM on the specified due date.

Late Submission Policy: • You are allocated a total of eight (5) late days for the entire term, which you can use across different homework assignments without incurring a penalty.

- A homework assignment is considered “ d days late” if it is submitted d days past the due date.
- We allow a maximum of 2 late days ($d = 2$) per homework assignment. Submissions will not be accepted if they are more than 2 days late.
- Once you have used all your allocated eight late days, any submissions past the due date will not be accepted, regardless of the circumstances.