

ISOM 7000M: IS Doctoral Seminar

Causal Inference and Large Language Models in Digitization Research

2024 Spring

Ohchan Kwon (ohchankw@ust.hk) and Zhitao Yin (zhitaoyin@ust.hk)

- Meets Thursday 6-9pm; Rm 1026, LSK Bldg.
- 3-credit letter-graded, based on 20% class participation, 40% assignments, 40% empirical project

Overview

This course aims to equip you with the skills to critically read and write empirical papers leveraging observational data for causal analysis—identifying "what causes what and why." As consumers, you will learn to discern when regression results can be interpreted as causal evidence. As future producers, you will discover how to spot and develop research projects that establish compelling causal relationships. In addition, we discuss how recent large language models can provide useful tools in empirical research.

There are already good econometrics (e.g., SOSC 5090 Quantitative Methods for Social Science Research, ECON 5300 Econometrics) and causal inference courses (e.g., SOSC 5340 Econometric Approaches to Social Science Research II). Students planning to do empirical research should certainly take these courses. However, it has been my experience that these methods courses still leave some important skillsets uncovered. Typical econometrics courses focus more on estimation and inference issues, that is, how to establish relationships among variables and less on causal identification issues. Causal inference courses discuss identification issues, but because many "classical" examples are from papers published in disciplinary journals a long time ago, students often do not learn some implicit rules and patterns to write a publishable paper in their respective fields as of now.

To fill the gap, this course uses recent papers in the economics of IT/digitization area. We will learn about causal inference techniques by carefully examining and replicating examples from recent papers. Compared to other typical quantitative methods courses, this course will focus on a few specific techniques (e.g., difference-in-differences) and discuss those techniques in much greater depth. Moreover, this course tries to cover more practical tools and techniques side of social science research, often described as the "hidden curriculum" in doctoral education.

Anyone interested in empirical research and/or digitization is welcome to take this course. We assume your familiarity with undergraduate-level statistics and econometrics. We also assume some familiarity with at least one statistical programming language (e.g., R, Stata, Python). Students without such background can still take the course as long as they are willing to spend considerable time to catch up.

Class Structure

Each week, our class will start by reviewing the main concepts from the previous session. We will do this by asking each student to introduce a paper that utilizes the ideas learned and evaluate its strengths and weaknesses. After that, our instructors will introduce new concepts and tools. Our primary focus will be on understanding and practical application, rather than theoretical proofs. We will also examine selected papers and their code to see these methods in action. Furthermore, we will cover the best practices for effectively communicating methods and results to an audience, ensuring clarity without overstating our findings.

Recommended Readings

- Angrist, Joshua D., and Jorn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.
- Scott Cunningham, *Causal Inference: The Mixtape*, <https://mixtape.scunning.com/>
- Goldfarb, A. and Tucker, C.E., 2014. *Conducting research with quasi-experiments: A guide for marketers*. Rotman School of Management Working Paper, (2420920).

Class Requirements

1) Class Participation (20%). Students must thoroughly read each assigned paper and actively engage in class discussions. This course is designed as a seminar rather than a traditional lecture series. Active participation and hands-on learning are key; therefore, it is crucial that everyone arrives well-prepared. Regular attendance is strongly encouraged.

2) Assignment (40%). There will be two assignments (20% x 2). The first assignment involves applying the difference-in-differences method to a provided dataset. The second assignment focuses on utilizing Large Language Models (LLMs). These assignments will be covered in detail during sessions in late March and early April, with specific guidelines and requirements to be announced closer to these dates.

3) Empirical Project (40%). For your final project, you are required to write a short empirical paper. You can choose any question of your interest, provided that you are asking a clear causal question that you can answer using the tools covered in the course. It is recommended that you use your ongoing research projects, for example, with your advisors. You are allowed to write a replication paper, but it should extend the original analysis by applying techniques covered in the course. The final paper should be a minimum of 10 pages long and should resemble an empirical journal article, except for the literature review and theory development. You will be expected to present your project in the final class.

Class Schedules and Readings

Week 1 (Feb 1): Introduction

Framing:

- Gelman, Andrew, and Guido Imbens. Why ask why? Forward causal inference and reverse causal questions. No. w19614. National Bureau of Economic Research, 2013.
- Little, Andrew T. "Three Templates for Introductions to Political Science." (2016). http://andrewlittle.com/papers/little_intros.pdf
- Cachon, G.P., 2012. What is interesting in operations management?. *Manufacturing & Service Operations Management*, 14(2), pp.166-169.

Elaborate theory:

How to write a social science paper:

- Imai, Kosuke. "How to Write an Empirical Social Science Paper." <https://imai.fas.harvard.edu/teaching/files/HowToPaper.pdf>
- Goldfarb and Tucker, Conducting Research with Quasi-Experiments https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2420920
- Simcoe, Timothy. Empirical Etiquette <http://people.bu.edu/tsimcoe/etiquette.html>

Readings for Economics of Innovation and Digitization (as a reference)

- Goldfarb, Avi, and Catherine Tucker. "Digital economics." *Journal of Economic Literature* 57.1 (2019): 3-43.

Week 2 (Feb 8/Feb 15): Causality, Experiments, and Matching

Background readings:

- (*) Angrist and Pischke: Chapters 1 and 2.
- Ho, D. E., K. Imai, G. King, E. A. Stuart. Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference. *Political Analysis* 15(3) 199–236.

Hands-on exercise:

- (*) Bloom, N., J. Liang, J. Roberts, Z. J. Ying. 2015. Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics* 130(1) 165–218.
- (*) Edelman, B., M. Luca, D. Svirsky. 2017. Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics* 9(2) 1–22.

Recent examples:

- Cui, R., Li, J. and Zhang, D.J., 2020. Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb. *Management Science*, 66(3), pp.1071-1094.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial

intelligence. *Science*, 381(6654), 187–192.

- Gunarathne, P., H. Rui, A. Seidmann. 2021. Racial Bias in Customer Service: Evidence from Twitter. *Information Systems Research*.

Week 4 (Feb 22): Difference-in-Differences (1) – Natural Experiment

Background readings:

- Angrist and Pischke: Chapters 5.
- Card, D., A. B. Krueger. 1994. Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania. *American Economic Review* 84(4) 772–793.
- Kahn-Lang, A., & Lang, K. (2020). The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications. *Journal of Business & Economic Statistics*, 38(3), 613–620.

Hands-on exercise:

- Impact of Amazon certifications

Recent examples:

- Zhang, L. 2018. Intellectual Property Strategy and the Long Tail: Evidence from the Recorded Music Industry. *Management Science* 64(1) 24–42.
- Sen, A., C. Tucker. 2020. Product Quality and Performance in the Internet Age: Evidence from Creationist-Friendly Curriculum. *Journal of Marketing Research*.
- Zhang, S., Lee, D., Singh, P. and Mukhopadhyay, T., 2022. Demand interactions in sharing economies: Evidence from a natural experiment involving airbnb and uber/lyft. *Journal of Marketing Research*, 59(2), pp.374-391.

Week 5 (Feb 29): Difference-in-Differences (2) – Staggered Adoption Design

Background readings:

- (*) Cunningham: Difference-in-Differences
- Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), pp.254-277.
- Roth, J., Sant’Anna, P.H., Bilinski, A. and Poe, J., 2023. What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*.
- Baker, A.C., Larcker, D.F. and Wang, C.C., 2022. How much should we trust staggered difference-in-differences estimates?. *Journal of Financial Economics*, 144(2), pp.370-395.

- Athey, S. C., G. W. Imbens. 2022. Design-based analysis in Difference-In-Differences settings with staggered adoption. *Journal of Econometrics* 226(1) 62–79.
- Cengiz, D., A. Dube, A. Lindner, B. Zipperer. 2019. The Effect of Minimum Wages on Low-Wage Jobs. *The Quarterly Journal of Economics* 134(3) 1405–1454. – Appendix D for “Stacked Design”.

Hands-on exercise:

- Impact of Amazon entry on online marketplace

Recent examples:

- Wang, J., G. Li, K.-L. Hui. 2021. Monetary Incentives and Knowledge Spillover: Evidence from a Natural Experiment. *Management Science*.
- Pattabhiramaiah, A., E. M. Overby, L. Xu. 2021. Spillovers from Online Engagement: How a Newspaper Subscriber’s Activation of Digital Paywall Access Affects Her Retention and Subscription Revenue. *Management Science*.
- Aral, S., P. S. Dhillon. 2020. Digital Paywall Design: Implications for Content Demand and Subscriptions. *Management Science*.
- Xu, Y., Ghose, A. and Xiao, B., 2023. Mobile payment adoption: An empirical investigation of Alipay. *Information Systems Research*.

Week 6 (Mar 7): Large Language Model for Scientific Research [Zhitao]

Reading:

- de Kok T (2023) Generative LLMs and textual analysis in accounting: (Chat) GPT as research assistant?. Available at SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4429658
- Susarla A, Gopal R, Thatcher JB, Sarker S (2023) The Janus Effect of generative AI: Charting the path for responsible conduct of scholarly activities in Information Systems. *Inform. Systems Res.* 34(2): 399-408.
- Li P, Castelo N, Katona Z, Sarvary M (2024) Frontiers: Determining the validity of Large Language Models for automated perceptual analysis. *Marketing Sci.* Articles in advance.

Week 7 (Mar 14): Large Language Model for Simulating Economic Agents [Zhitao]

Reading:

- Horton JJ (2023) Large language models as simulated economic agents: What can we learn from homo silicus? (No. w31122). National Bureau of Economic Research.
- Chen, Yiting, Tracy Xiao Liu, You Shan, and Songfa Zhong. "The Emergence of Economic Rationality of GPT." *PNAS*.12763 (2023).

Hands-on exercise:

- <https://github.com/johnjosephhorton/homosilicus>

Week 8 (Mar 21): Hackathon Demo Day of LLM for Research [Zhitao]

Choose one of the following topics:

- Research Assistant Tool: Develop a Python application that leverages LLMs to automate parts of the research process, such as generating literature review summaries, identifying research gaps, or suggesting potential methodologies.
- Economic Agent Simulator: Create a Python simulation application where LLMs act as consumers, firms, or policymakers, making decisions based on simulated economic scenarios. This task could explore market dynamics, negotiation strategies, or the impact of policy changes.
- Perceptual Analysis Application: Utilizing LLMs for automated perceptual analysis, you could attempt to validate the AI's interpretations against human judgments in various contexts, such as consumer product perceptions, brand sentiment analysis, or social media content interpretation.

Note:

- Present a demo of your application and discuss what you find and so what [30 mins]
- This is an individual task.

Week 9 (Apr 11): Difference-in-Differences (3) – Checklist and Other Topics

Background readings:

- Abadie, A., 2021. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2), pp.391-425.

Recent example:

- Lu, S., Rajavi, K. and Dinner, I., 2021. The effect of over-the-top media services on piracy search: Evidence from a natural experiment. *Marketing Science*, 40(3), pp.548-568.

Week 10 (Apr 18): Instrumental Variables

Background readings:

- Angrist and Pischke: Chapters 5.
- Cunningham: Instrumental Variables
- Borusyak, K., Hull, P. and Jaravel, X., 2022. Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1), pp.181-213.
- Goldsmith-Pinkham, P., Sorkin, I. and Swift, H., 2020. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8), pp.2586-2624.

Recent examples:

- Sun, M., F. Zhu. 2013. Ad Revenue and Content Commercialization: Evidence from Blogs. *Management Science*. (*)
- Seiler, S., S. Yao, W. Wang. 2017. Does Online Word of Mouth Increase Demand? (And How?) Evidence from a Natural Experiment. *Marketing Science* 36(6) 838–861. (*)
- Todri, V. 2021. Frontiers: The Impact of Ad-Blockers on Online Consumer Behavior. *Marketing Science*.

Week 11 (Apr 25): Robustness Checks

Background readings:

- Eggers, A.C., Tuñón, G. and Dafoe, A., 2023. Placebo tests for causal inference. *American Journal of Political Science*.
- Oster, E. 2017. Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics* 40(2) 1–18.
- Chen, J. and Roth, J., 2023. Logs with zeros? Some problems and solutions. *The Quarterly Journal of Economics*, p.qjad054.

Week 12 (May 2): Mechanisms

Background readings:

- Goldfarb, A., Tucker, C. and Wang, Y., 2022. Conducting research in marketing with quasi-experiments. *Journal of Marketing*, 86(3), pp.1-20.
- Aneshensel, Carol S. *Theory-based data analysis for the social sciences*. Sage Publications, 2012. Chapter 1, Introduction to Theory-Based Data Analysis.
- Rosenbaum, Paul R. *Observation and experiment*. Harvard University Press, 2017. Chapter 7, Elaborate Theories

Week 13 (May 9): Final Presentation
