

ECON 5101

Econometric Thinking for Business Strategists

Causal inference and business experimentation in the AI era

Instructor	Prof. Dean He, Ph.D. (deanhe@ust.hk)	TA	Emily Chen
Class Dates	Oct 31, Nov 14, Nov 28, and Dec 12, 2026 (four all-day sessions)	Office	LSK 5061B
Class Time	09:30–13:00 and 14:30–18:00	Credits	2, Letter Grade
Venue	Rm 6555, Academic Building (near Lifts 29–30) on Oct 31, due to School Information Day 2026; Rm 6602, Academic Building (near Lifts 31–32) on Nov 14, Nov 28, and Dec 12		

Course Overview

This course introduces the art and science of causal reasoning in business strategy. It equips business leaders and strategists to interpret, critique, and commission econometric analysis, and to design credible experiments of their own. The course is built around the methods that modern firms actually use: randomized controlled experiments, regression discontinuity, difference-in-differences, synthetic control, and the newer experimentation designs developed inside digital marketplaces and platform companies.

Econometrics has moved away from mathematical complexity and toward transparent tools that travel well into business practice. At the same time, the rise of AI has transformed both what firms experiment on and how analysis gets done. Companies such as Booking.com, Microsoft, Netflix, and ByteDance now run tens of thousands of controlled experiments each year, on everything from webpage layouts to recommendation algorithms and large language model features. Meanwhile, AI assistants have made sophisticated data analysis accessible to managers who do not code. This course embraces both shifts: students study how leading firms experiment in the AI era, and AI tools may be used as analysis aids in coursework, with the causal reasoning and the skepticism remaining the student's own.

No coding, spreadsheets, or advanced mathematics are required. The emphasis throughout is on strategic thinking: knowing when a causal claim deserves belief, when it does not, and how to design the test that settles the question.

Course Objectives and Learning Outcomes

By the end of the course, students will be able to:

- Approach claims of cause and effect in business decision making with disciplined skepticism, recognize flaws in causal claims, and communicate specific reasons to be doubtful, including for claims produced by AI systems.
- Explain why randomized experimentation is the benchmark for causal inference in strategy, why randomization is the key ingredient, and what precision, power, and balance mean for the credibility of a test.
- Recognize when natural experiments and quasi-experimental designs (regression discontinuity, difference-in-differences, synthetic control) are good substitutes for randomization, and articulate the assumptions each design rests on.

- Understand the experimentation designs of modern platform businesses, including switchback and geo experiments, interference and network effects, heterogeneous treatment effects and uplift targeting, and surrogate metrics for long-term outcomes.
- Design, pilot, and pitch an original experimental analysis for a real business strategy question, with primary focus on strategic thinking and experimental design rather than mathematics or coding.

Course Outline

The course outline and case readings are subject to change until the final version before the launch of the class.

Session	Topics	Required Readings and Cases
Section 1 · Experimental Design		
Class 1 Oct 31, AM	• Introduction • Average treatment effects • Selection bias • Why businesses (still) don't experiment • AI tool orientation	• HBS case: <i>Booking.com</i> (619-015) • Thomke, "Building a Culture of Experimentation," <i>HBR</i> • Thomke, "Why Businesses Don't Experiment"
Class 2 Oct 31, PM	• Precision and power • Randomization and balance • How platform firms industrialize testing	• Kohavi and Thomke, "The Surprising Power of Online Experiments," <i>HBR</i> • Discussion of group project ideas with the instructor • Problem Set 1 distributed
Section 2 · Econometric Analysis and Natural Experiments		
Class 3 Nov 14, AM	• Regression with covariates • Heterogeneous effects • Uplift targeting: who to treat under a limited budget	• HBS case: <i>Advertising Experiments at Restaurant Grades</i> • Instructor note on heterogeneous effects and uplift targeting in industry practice
Class 4 Nov 14, PM	• Non-compliance • Placebo and Hawthorne effects • Double machine learning at the level of intuition	• Bojinov, Saint-Jacques, and Tingley, "Avoid the Pitfalls of A/B Testing," <i>HBR</i> • "Virgin Atlantic Tested 3 Ways to Change Employee Behavior" • Instructor note on the predict, residualize, relate logic of double ML
Class 5 Nov 28, AM	• Regression discontinuity and its application: Uber • Synthetic control: building a synthetic twin • Geo experiments after the privacy shift	• Cohen et al., "Using Big Data to Estimate Consumer Surplus: The Case of Uber," pp. 6–10 • Instructor note on synthetic control and geo testing in marketing measurement • Problem Set 1 due before class
Class 6 Nov 28, PM	• Difference-in-differences and its application: bargaining at eBay • Modern DID: staggered adoption and event-study plots • Marketplace experiments: switchback designs, interference, cluster randomization	• Backus et al., sequential bargaining on eBay • INSEAD case: <i>Uber vs. Didi</i> (IN1306) • Instructor note on switchback experiments at ride-hailing and delivery platforms • Problem Set 2 distributed

Class 7 Dec 12, AM	• Sample attrition • Surrogate metrics and long-term effects • Experimentation in the AI era: LLM features, recommendation engines, algorithmic pricing	• Bloom et al., “Does Working From Home Work?” pp. 176–183 • HBS case: <i>TikTok: The Algorithm Will See You Now</i> (824-125) • Zhang et al., surrogate index study of 200 A/B tests at Netflix (excerpt) • Calvano et al., “Artificial Intelligence, Algorithmic Pricing, and Collusion” (discussion reading) • Group Project due; Problem Set 2 due
Class 8 Dec 12, PM	• Group presentations • Experiment pitches and oral defense	• Written pitch submitted with AI-use acknowledgment

Grading and Assessment

Component	Description	Weight	Additional Notes
Class Participation	In-class discussion is integral to the course. Students are expected to ask relevant questions, offer insight into the topic at hand, and behave professionally. Quality of contribution matters more than quantity. Contributions that catch flaws in causal claims, including AI-generated ones, are especially valued. Scores also account for attendance and lateness.	25%	Graded at the individual level. HKUST MS attendance policy applies: absence from 25% of the course (2 of 8 classes) incurs a one-third grade deduction (e.g., A– to B+); absence beyond 25% results in an F.
Problem Sets	Two group problem sets. Problem Set 1 covers experimental methods; Problem Set 2 covers analysis of natural experiments and includes a critique memo of a published causal claim. Each submission contains the interpreted results, a memo identifying threats to validity, and an acknowledgment of any AI use, per the course AI policy.	25%	Written up and submitted as groups; each group member shall contribute individually.
Group Project and Presentation	Teams of three develop a causal question tied to a real business strategy, design an experiment to test it, run a pilot study to inform the design where feasible, and pitch a proposed design and implementation strategy directed at the key decision maker in the company during Class 8.	50%	Experimental idea: 10% of project grade. Initial and revised design plan presented to the instructor: 10%. In-class pitch during Class 8: 20%, graded at the individual level. Written pitch: 60%. Students should be prepared to explain and defend their analysis, including any AI-assisted work, during the presentation. All other elements are graded at the group level.

Policy on the Use of Generative AI

Under HKUST's *Policy for GenAI Integration in Teaching and Learning*, faculty members set the course-level rules for generative AI and communicate them to students in writing. The policy for this course is as follows: generative AI tools (for example, the HKUST GenAI Platform, ChatGPT, or Claude) may be used as aids for analysis, brainstorming, drafting, and understanding course methods, subject to the conditions below.

- **Acknowledgment.** Consistent with University guidance on transparency, any substantive use of generative AI in a submitted deliverable must be acknowledged, stating the tool used and the purpose. Unacknowledged use of AI-generated content in assessed work is treated as a violation of the HKUST Academic Honor Code.
- **Ownership of reasoning.** Following CEI guidance, students should be prepared to explain their work, justify their sources, provide drafts, and discuss alternative perspectives, whether or not AI tools were used. Submitting AI-generated analysis that no group member can explain is not permitted.
- **Data care.** Per University guidelines, do not submit confidential, proprietary, or personal data to external AI tools. The HKUST GenAI Platform is the recommended option when working with course datasets.
- **Critical use.** AI tools make mistakes, including plausible but wrong causal claims. Verifying and challenging AI output is part of the discipline this course teaches.

Course Materials

(A) Required Reading

Huntington-Klein, N. (2025). *The effect: An introduction to research design and causality* (2nd ed.). CRC Press. <https://theeffectbook.net>

Thomke, S. H. (2020). *Experimentation works: The surprising power of business experiments*. Harvard Business Review Press.

(B) Optional Reading

Angrist, J. D., & Pischke, J.-S. (2014). *Mastering 'metrics: The path from cause to effect*. Princeton University Press.

Cunningham, S. (2026). *Causal inference: The remix*. Yale University Press.

Thomke, S. H. (2026). *The scientific mindset: How thinking scientifically drives better business decisions*. Harvard Business Review Press. Forthcoming in November 2026; an anticipated reference guide for this course.

(C) Cases

Advertising experiments at Restaurant Grades [Case]. Harvard Business School Publishing.

Booking.com (HBS Case No. 619-015). Harvard Business School Publishing.

TikTok: The algorithm will see you now (HBS Case No. 824-125). Harvard Business School Publishing.

Uber vs. Didi (INSEAD Case No. IN1306). INSEAD Publishing.

(D) Articles and Papers

Backus, M., Blake, T., Larsen, B., & Tadelis, S. (2020). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *The Quarterly Journal of Economics*, 135(3), 1319–1361.

Bloom, N., Liang, J., Roberts, J., & Ying, Z. J. (2015). Does working from home work? Evidence from a Chinese experiment. *The Quarterly Journal of Economics*, 130(1), 165–218.

Bojinov, I., Saint-Jacques, G., & Tingley, M. (2020, March–April). Avoid the pitfalls of A/B testing. *Harvard Business Review*.

- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297.
- Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). *Using big data to estimate consumer surplus: The case of Uber* (NBER Working Paper No. 22627). National Bureau of Economic Research.
- Kohavi, R., & Thomke, S. (2017, September–October). The surprising power of online experiments. *Harvard Business Review*.
- Thomke, S. (2020, March–April). Building a culture of experimentation. *Harvard Business Review*.
- Thomke, S. (n.d.). Why businesses don't experiment. *Harvard Business Review*.
- Zhang, V., Zhao, M., Le, A., & Kallus, N. (2023). *Evaluating the surrogate index as a decision-making tool using 200 A/B tests at Netflix* [Preprint]. arXiv. <https://arxiv.org/abs/2311.11922>
-

© Copyright 2026 HKUST Business School. All rights reserved. The course outline, cases, and readings are subject to change until the final version before the launch of the class.