Columbia University School of International and Public Affairs
PhD in Sustainable Development

Causal Inference Workshop

Syllabus

(updated April 29, 2022)

Course Title: Causal Inference Workshop

Instructor: Suresh Naidu, sn2430@columbia.edu

Teaching Assistant: Claire Palandri, cp2913@columbia.edu

Semester: Spring 2022

Meeting Date/s Times: Fridays, 9:00-10:00am

Location: IAB 1101

Credits: 1.5 (Pass/Fail)

Course Description

This workshop designed for the students in the PhD program in Sustainable Development covers the fundamental theory and techniques of causal inference. Specifically tailored to students trained in econometrics and positioned to conduct interdisciplinary research, it ties back the econometrics approaches covered to the underlying statistical framework, and provides the students with the tools to conduct rigorous empirical analyses and to share and defend their approach in front of both economics and non-economics audiences. Lower-year students are presented the fundamental methods for observational studies; upper-year students can discuss how they employ them in their own current research. Participants are presented with the core methods in the field, their limitations and best practices, and less-common statistical methods relevant for causal inference.

Course Overview

The workshop will consist of a weekly class session, led by the teaching assistant. The 13 weeks are organized into four sections:

A. *[The aim of causal inference]* Fundamentals of inferential statistics are reviewed, followed by the theoretical framework of potential outcomes for causal inference, and how it's implemented with observational data through regression modeling.
B. *[Core of how it's done]* The most common identification strategies — special cases of regression adapted to particular forms of natural experiments — are reviewed. For each method, the canonical setup is presented in the first part of the session, in particular: the data generating process assumed; the identifying assumptions; the estimand of interest; the estimator used; best practices; strengths and weaknesses. The second part of the session then puts the theory in practice by discussing the analyzes of working or published papers: work in progress by a current PhD student and/or a published paper on sustainable development.

C. *[How to improve upon it]* How to obtain stronger causal inferences through steps at the analysis stage. Pre- and post-estimation best practices are presented, including how to support the assumptions on which the inferences rest, and the benefits of matching and prediction for causal inference.
D. Less common topics in causal inference are presented, such as randomization inference, the synthetic control method, and directed acyclic graphs.

The workshop does not follow a specific textbook, but the two references in which the participants will find most of the material covered — and that are highly recommended as complements of each other — are Angrist and Pischke (2008) and Gelman et al. (2020).

Grading The course is graded on a Pass/Fail basis. The course grade will be based mostly on attendance, and also on a home assignment to be turned in on the final week of class. It will consist of the replication of the analysis of a published paper, to supplement with statistical analyses covered during the course (e.g., diagnosis checks of underlying modeling assumptions, model evaluation, matching...).

Course Structure Week-by-week list of class topics:

A. Causal inference fundamentals

1. Overall presentation; Inferential statistics fundamentals

- The interpretation of slope parameter estimates and of their statistical significance relies on modeling assumptions. Regression models as conditional distributions; assumptions of the classical linear regression model, and the estimator properties depending on them.
- Dealing with some departures from the usual assumptions: sandwich estimators; limited y models.

2. The potential outcomes framework and identification

- The Neyman-Rubin causal model or potential outcomes framework. Identification from independence assumptions.
- The regression on the treatment recovers an average treatment effect (relation between observed and potential outcomes). Simplest case & extensions (limited Y, covariates X, continuous D).
- References Deaton and Cartwright (2018); Rubin (1974)

B. Design stage: Identification strategies

3. Instrumental Variables (IV)

- Theory: instruments and compliance behavior; two-stage least squares; local average treatment effect (LATE); computing average complier characteristics and getting more out of a LATE.
- Application: a working paper by a current PhD student and/or a published paper related to SDev.
- References Angrist and Pischke (2008, eq. 4.4.8); Abadie (2003); Andrews et al. (2019); Kowalski (2018)

4. Regression Discontinuity (RD)

- Theory: deterministic but discontinuous assignment; estimation with flexible functional forms; Sharp RD, Fuzzy RD (imperfect compliance).
- Application: a working paper by a current PhD student and/or a published paper related to SDev.
- References Almond and Doyle (2011)

5. Difference-in-Differences (DiD), ...-in-Differences (DiDiD), event-studies

- Theory: different models of the counterfactual; pre-trends; justifying a third difference; pitfalls of weighted sums of the average treatment effects with two-way fixed effect estimators.
- Application: a working paper by a current PhD student and/or a published paper related to SDev.
- References de Chaisemartin and D'Haultfœuille (2020, 2022); Goodman-Bacon (2021)

C. Analysis stage: Steps for stronger causal inferences

6. Limitations of identification strategies; Pre-estimation steps: restructuring the data

· Limitations of identification strategies; steps can be taken pre-/during-/post-estimation.
· Exploratory data analysis: scatterplot your raw data (and show some summary in your final paper).
· Restructuring the data by matching to improve overlap: *in place of* (Angrist and Pischke, 2008) or *on top of* (Gelman et al., 2020; Ho et al., 2007) regression, but never in place of design. Examples of distance metrics (propensity score, Mahalanobis distance). Algorithm for propensity score matching.
· References Almond et al. (2005)

7. Estimation steps: controls & TE heterogeneity

· Required/forbidden controls; optional good/bad controls; bias amplification.
· Treatment effect heterogeneity: interactions with pretreatment variables; pooling of group-level TEs and introduction to multilevel modeling.
· References Cinelli et al. (2021); Feller and Gelman (2015); Gelman (2006); Middleton et al. (2016)

8. Post-estimation steps #1: supporting assumptions

· Modeling assumptions; back to inference fundamentals: post-estimation model diagnostics. Fake data simulations: fit the model to simulated data where you know the true parameter values.
· Identifying assumptions; show a balance test table and do falsification tests. Examples of falsification tests for each identifying assumption of common identification strategies (IV, RDD, DiD).
· References Ex: Imbens and Wooldridge (2009, eq. 3); Imbens and Lemieux (2008); McCrary (2008)

9. Post-estimation steps #2: model selection; external validity

· Model selection; Regularization methods.
· Prediction isn't part of statistical inference, but can help 1. support your assumptions; 2. prove general interest of your results. Measures of performance: information criteria; cross-validation. Bayesian inference.

D. Other topics in causal inference

10. Another approach: the causal graph framework

· Elements of Directed Acyclic Graphs (DAGs). Two main identification strategies: 1. blocking back-door paths; 2. instruments.
· Strengths and weaknesses of the PO vs. DAG approaches.
· References Morgan and Winship (2015, ch. 1.5 & 3); Imbens (2020); Cunningham (2021, chap. 3)

11. Randomization inference

· Design-based vs sampling-based inference. Motivations: no true sampling variation to speak of; not having to rely on asymptotics; preserving unformalizable clustered data structures.
· Application: a working paper by a current PhD student and/or a published paper related to SDev.
· References Athey and Imbens (2017); Cooperman (2017)

12. Synthetic Control Method (SCM)

· Creating an optimized counterfactual: the "synthetic unit".
· Application: a working paper by a current PhD student and/or a published paper related to SDev.
· References Abadie (2021); Abadie et al. (2010, 2015); Arkhangelsky et al. (2021); Ferman et al. (2020)

E. Wrap-up & *'Replication+'* exercise

13. Wrap-up + Exercise

- Putting all the pieces together; How to present one's causal analysis.
- "Replication +" exercise: Using data from a published causal inference paper, we redo the analysis, conduct post-estimation checks and explore other methods.

Accommodations I am committed to teaching a class in which every student can learn and participate. If you have a disability or a circumstance that affects your learning in this course, please let me know as soon as possible so that we can discuss together the best way to meet your needs. Do not hesitate also to contact Disability Services to discuss options to removing barriers in the course, including accommodations.

References

Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics*, 113(2):231–263, ISSN: 0304-4076, DOI: 10.1016/S0304-4076(02)00201-4.

Abadie, A. (2021). Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects. *J. Econ. Lit.*, 59(2):391–425, ISSN: 0022-0515, DOI: 10.1257/jel.20191450.

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *J. Am. Stat. Assoc.*, 105(490):493–505, DOI: 10.1198/jasa.2009.ap08746.

Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *Am. J. Pol. Sci.*, 59(2):495–510, DOI: 10.1111/ajps.12116.

Almond, D., Chay, K. Y., and Lee, D. S. (2005). The Costs of Low Birth Weight. *Q. J. Econ.*, 120(3):1031–1083, ISSN: 0033-5533, 1531-4650.

Almond, D. and Doyle, J. J. (2011). After midnight: A regression discontinuity design in length of postpartum hospital stays. *American Economic Journal: Economic Policy*, 3(3):1–34, DOI: 10.1257/pol.3.3.1.

Andrews, I., Stock, J. H., and Sun, L. (2019). Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics*, 11(1):727–753, DOI: 10.1146/annurev-economics-080218-025643.

Angrist, J. and Pischke, J.-S. (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press, Princeton, NJ, ISBN: 9781400829828, DOI: 10.1515/9781400829828.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic Difference-in-Differences. *Am. Econ. Rev.*, 111(12):4088–4118, DOI: 10.1257/aer.20190159.

Athey, S. and Imbens, G. W. (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments*, volume 1, pages 73–140. Elsevier, ISSN: 2214-658X, DOI: 10.1016/bs.hefe.2016.10.003.

Cinelli, C., Forney, A., and Pearl, J. (2021). A crash course in good and bad controls, DOI: 10.2139/ssrn.3689437. Working paper.

Cooperman, A. D. (2017). Randomization Inference with Rainfall Data: Using Historical Weather Patterns for Variance Estimation. *Polit. Anal.*, 25(3):277–288, DOI: 10.1017/pan.2017.17.

Cunningham, S. (2021). *Causal Inference: The Mixtape*. Yale University Press, ISBN: 9780300251685, 9780300255881, https://mixtape.scunning.com.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *Am. Econ. Rev.*, 110(9):2964–2996, DOI: 10.1257/aer.20181169.

de Chaisemartin, C. and D'Haultfœuille, X. (2022). Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey, DOI: 10.3386/w29691, http://www.nber.org/papers/w29691. Working Paper 29691, National Bureau of Economic Research.

Deaton, A. and Cartwright, N. (2018). Understanding and misunderstanding randomized controlled trials. *Soc. Sci. Med.*, 210:2–21, ISSN: 0037-7856, DOI: 10.1016/j.socscimed.2017.12.005.

Feller, A. and Gelman, A. (2015). Hierarchical models for causal effects. In Scott, R. A., Kosslyn, S. M., and Buchmann, M. C., editors, *Emerging Trends in the Social and Behavioral Sciences*. John Wiley & Sons, ISBN: 9781118900772, DOI: 10.1002/9781118900772.etrds0160.

Ferman, B., Pinto, C., and Possebom, V. (2020). Cherry picking with synthetic controls. *J. Policy Anal. Manage.*, 39(2):510–532, ISSN: 0276-8739, 1520-6688, DOI: 10.1002/pam.22206.

Gelman, A. (2006). Multilevel (Hierarchical) Modeling: What It Can and Cannot Do. *Technometrics*, 48(3):432–435, DOI: 10.1198/004017005000000661.

Gelman, A., Hill, J., and Vehtari, A. (2020). *Regression and Other Stories*. Cambridge University Press, ISBN: 9781107023987.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *J. Econom.*, 225(2):254–277, ISSN: 0304-4076, DOI: 10.1016/j.jeconom.2021.03.014.

Ho, D. E., Imai, K., King, G., and Stuart, E. A. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15(3):199–236, DOI: 10.1093/pan/mpl013.

Imbens, G. W. (2020). Potential Outcome and Directed Acyclic Graph Approaches to Causality: Relevance for Empirical Practice in Economics. *J. Econ. Lit.*, 58(4):1129–1179, DOI: 10.1257/jel.20191597.

Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *J. Econom.*, 142(2):615–635, DOI: 10.1016/j.jeconom.2007.05.001.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent Developments in the Econometrics of Program Evaluation. *J. Econ. Lit.*, 47(1):5–86, ISSN: 0022-0515, DOI: 10.1257/jel.47.1.5.

Kowalski, A. E. (2018). Reconciling Seemingly Contradictory Results from the Oregon Health Insurance Experiment and the Massachusetts Health Reform. Working Paper 24647, National Bureau of Economic Research, DOI: 10.3386/w24647, http://www.nber.org/papers/w24647.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *J. Econom.*, 142(2):698–714, ISSN: 0304-4076, DOI: 10.1016/j.jeconom.2007.05.005.

Middleton, J. A., Scott, M. A., Diakow, R., and Hill, J. L. (2016). Bias Amplification and Bias Unmasking. *Polit. Anal.*, 24(3):307–323, DOI: 10.1093/pan/mpw015.

Morgan, S. L. and Winship, C. (2015). *Counterfactuals and Causal Inference*. Cambridge University Press, ISBN: 9781107065079.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.*, 66(5):688–701, ISSN: 0022-0663, 1939-2176, DOI: 10.1037/h0037350.