POLITICAL SCIENCE 406, SPRING 2025

FACULTY: JAYE SEAWRIGHT
PHONE#: 847-467-1148

Email: j-seawright@northwestern.edu

1. Basic Course Information

Prerequisites: Political Science 403 and 405 or equivalent.

Course Objectives: This course offers an introduction to quantitative approaches to causal inference in the social sciences.

The goals of the course involve starting a lifetime of engagement with the rapidly evolving literature behind applied quantitative causal inference. While causal inference is difficult and far from straightforward, even in most experiments, scholars and practitioners have developed and continue to produce clever and insightful ideas that help us design studies and analyze results in ways that are more coherent, insightful, and reliable. But because these ideas are both exciting and important, new approaches are constantly emerging — and that state of affairs is likely to continue! We need to become good not just at a set of techniques but also at picking up new approaches.

Thus, through this seminar, students will practice a set of skills that prepare them for the future, as well as for knowledge of the current state of causal inference. By the end of this seminar, students will be able to:

- Translate between mathematical and verbal descriptions of causal inference estimators
- Use a published article and its affiliated R package to implement and correctly interpret a causal inference estimator with data
- Evaluate a collection of causal inference estimators related to a single research design in order to either select a best estimator for the research context of interest or to conclude that the estimators are interchangeable given current knowledge
- Correctly describe, and when feasible, test the assumptions involved with each family of causal inference research designs
- Communicate about quantitative causal inference at a professional level in a way appropriate for workshop, conference, and other relevant conversational settings
- Produce a written "grant proposal" that features a research design that shows mastery of at least one cutting-edge quantitative causal inference estimator

Note: Any student requesting accommodations related to a disability or other condition is required to register with AccessibleNU (accessiblenu@northwestern.edu; 847-467-5530) and provide professors with an accommodation notification from AccessibleNU, preferably within the first two weeks of class. All information will remain confidential.

Books:

Causal Inference: The Mixtape by Scott Cunningham.

The Effect: An Introduction to Research Design and Causality by Nick Huntington-Klein, available online in an extremely useful Markdown version at `https://theeffectbook.net/index.html`.

Statistical Software: Students in this course are invited but not required to learn, use, and enjoy the statistical package, **R**. Most (all?) projects in this course can also be completed using Stata or Python, but there may be advantages to R, including:

- R is free. Anyone can download the newest version of R from the internet (see, for example, `http://cran.r-project.org/`). To get R for Windows from that address, click "Windows." Then click "base." Finally, click "Download R ¡latest version¿ for Windows" and walk through the installation process.
- R is one of the most extensive statistical packages in the world because it is open-source, although Python is clearly a close competitor at this point. As a result, new econometric and statistical techniques are often implemented and available online within days of their invention.
- Many people feel that R has better support for statistical graphics.
- I am a lot more familiar with R than with Stata. I've used Stata a few times, and use Python when relevant, but I do much of my research in R. So classroom examples will be done in R, and I'm a lot more likely to be able to help you solve problems in R than in Stata, Python, or other statistical environments. That said, the TA and/or I can usually point you toward relevant resources for doing the course labs in Stata, Python, or something else if that's your preference.

Class Time: Thursdays, 2-4:50.

Class Room: Scott Hall 212

Office Hours: Tuesdays, 3:30-5:30

2. Assessment

There are three major categories of assessments for this seminar.

The first involves presenting an estimator in class. Presenting an estimator will involve:

(1) Explaining the problem the estimator is intended to address

(2) Discussing the equation or equations that instantiate the estimator, interpreting them to the audience
(3) Describing the assumptions needed for the estimator, and ideally relating those to the equation(s)
(4) Spelling out the strengths and weakness of the estimator in terms of statistical properties like bias, consistency, variance, mean squared error, etc.
(5) Showing an applied example of the estimator, either one from published work or an original application

An in-class presentation should be ten to twelve minutes of prepared content, and the presenter should be ready for a period of eight to ten minutes of audience questions as the end. The presentation should have professional slides that help illustrate the key ideas — it is very challenging to discuss statistical estimators, coding, and results without visuals! The goal with all of this is to simulate the experience of a professional conference presentation discussing an idea in methodology, but using a fully attributed discussion of someone else's published ideas as a classroom analogue.

Please meet with Jaye Seawright as soon as possible (in office hours, by email, etc.) to choose an estimator to present and to plan your presentation. Please feel free to talk through drafts of your presentation and even to rehearse a version before the class arrives.

The presentation is graded as follows. An "outstanding" presentation includes all of the key elements listed above, is clear, professional, and helpful to the audience, has few if any mistakes (saying "I don't know" is not a mistake!), and has at least one aspect that shows unusually deep research — whether that is a careful original empirical example, more work than usual on strengths and weaknesses, or something else. A "pass" presentation includes all of the key elements listed above, is clear, professional, and helpful to the audience, and may have some mistakes or misunderstandings but not at a level that signals a lack of preparation. A "marginal" presentation is completely missing one or more key elements, is substantially confusing or unprofessional in spots, or has such deep mistakes that it raises questions about the degree of preparation. Finally, a "failed" presentation is missing multiple key elements, likely has no or only a very brief slide show, is casual and unprepared in an unprofessional way, and generally does not contribute to the educational progress of the class.

The second assessment for the seminar involves weekly lab exercises, which involve applying causal inference ideas from class to real data using partially guided code scripts. Assignments involve real data and real R packages for causal inference; they will give some steps in full code, some steps in incomplete hints, and some steps will be left for you to complete. Finally, there will be questions about what we are trying to accomplish, what certain results mean, etc., that ask you to talk about the methods we're learning in your own words. Lab assignments for the entire quarter are currently available on the course github site (`https://github.com/jnseawright/PS406`).

Labs are given one of three scores: outstanding, pass, or revisit. An "outstanding" lab has either no errors or only a tiny scattering of very minor glitches and also shows insight and mastery at a very high professional level. A grade of "outstanding" means that the student is likely ready, right now, to use the techniques in question in professional research. A "pass" lab shows clear understanding of the core ideas at work and gets results, but who still has one or more important misunderstandings (whether in terms of coding or concepts) that have been pointed out in the feedback. This grade goes to a student who is well on the road to mastery. A grade of "revisit" goes to a lab in which the student has more fundamental misunderstandings and would benefit from retrying the lab. Students who get this grade have two weeks to resubmit the lab if they choose to do so, at which time it will be regraded.

The third and final assessment is a mock grant proposal that features a research design that shows mastery of at least one cutting-edge quantitative causal inference estimator from this class. The proposal will be evaluated based on the criteria listed for the Northwestern Graduate School's Graduate Research Grant (`https://www.tgs.northwestern.edu/funding/fellowships-and-grants/internal-fellowships-grants/graduate-research-grant.html`), and the format must meet the rules for the "Description of the project" section of a proposal for that grant — five pages, double spaced, up to three pages of references/endnotes/figures — with the exception that it does not need to already have IRB approval.

Note that a successful proposal will not only score well on the criteria for grant review — which is good practice for your professional future! — but will also show mastery of at least one cutting edge causal inference estimator. I am not going to list out here which estimators are cutting edge and which are not, but you are probably best off if you use an estimator that you did not know about before this quarter.

To get an A in the course, a student must:

- Get an overall score of 1.5 out of 5 or higher on the grant proposal and also a score of "pass" on the presentation; or an overall score of 2 out of 5 or higher on the grant proposal and also a score of "outstanding" on the presentation.
- Complete labs with one of the following combinations of scores: four "outstanding;" three "outstanding" and two "pass;" two "outstanding" and four "pass;" one "outstanding" and six "pass;" or simply complete at least seven labs with at least a pass and no unresolved "revisits."

To get an A- in the course, a student must:

- Get an overall score of 2.5 out of 5 or higher on the grant proposal and also a score of "pass" on the presentation; or an overall score of 3 out of 5 or higher on the grant proposal and also a score of "outstanding" on the presentation.
- Complete labs with one of the following combinations of scores: three "outstanding;" two "outstanding" and two "pass;" one "outstanding" and four "pass;" or simply complete at least six labs with at least a pass and no unresolved "revisits."

To get a B+ in the course, a student must:

- Get an overall score of 3.5 out of 5 or higher on the grant proposal and also a score of "pass" on the presentation; or an overall score of 4 out of 5 or higher on the grant proposal and also a score of "outstanding" on the presentation.
- Complete labs with one of the following combinations of scores: two "outstanding;" one outstanding" and two "pass;" or simply complete at least five labs with at least a pass and no unresolved "revisits."

2.1. **Academic Honesty.** Group work is encouraged for labs. Your in-class presentation and grant proposal must, of course, reflect your original work. Any quotations from other people's work must be fully cited and documented. The same is true for paraphrases or for statistics or facts that are not general knowledge. Please do not hesitate to ask for additional details if you are confused about this assignment. The WCAS policy on academic integrity reads:

> In a scholarly community like Northwestern, academic integrity is of the utmost importance. If you are guilty of dishonesty in academic work, you may receive a failing grade in the course and be suspended or permanently excluded from the University. The brochure "Academic Integrity at Northwestern: A Basic Guide" details the types of offenses that constitute academic dishonesty and contains a thorough discussion of the proper citation of sources. You can get this brochure at the Office of Undergraduate Studies and Advising. A document on how instances of alleged academic dishonesty are handled is available online. The Undergraduate Catalog contains a non-exhaustive list of behaviors that violate standards of academic integrity. These include: cheating, plagiarism, fabrication, obtaining an unfair advantage, aiding and abetting dishonesty, falsification of records and official documents, and unauthorized access to computerized academic or administrative records or systems. Each of these is described in more detail in the catalog. One important type of academic dishonesty is plagiarism. Plagiarism includes more than just copying someone else's work. Northwestern's "Principles Regarding Academic Integrity" defines plagiarism as "submitting material that in part or whole is not entirely one's own work without

> attributing those same portions to their correct source." A Northwestern web page provides links to additional information on academic integrity, including information on relevant policies and on how to recognize and avoid violations of academic integrity in your own work. More tips on avoiding plagiarism are available from Northwestern's Writing Place. Sometimes students think that another student has acted in a way that is academically dishonest. In this situation you should consult with the Weinberg College Adviser.

This course's projects will be submitted electronically, via the course's Canvas page, rather than in printed form. As per university policy, all student work may by analyzed electronically for violations of the university's academic integrity policy and may also be included in a database for the purpose of testing for plagiarized content.

3. Course Schedule and Readings

This schedule is subject to changes (minor or major) depending on how long each topic actually takes us to cover, as well as on the needs of the class. Slides and code for in-class discussion and examples are available on the course github site (`https://github.com/jnseawright/PS406`).

Week 1: *Experiments.*

- Cunningham, Chapters 1, 3-4.
- Huntington-Klein, Chapters 1-11.

Examples:

- Bush, Sarah Sunn and Lauren Prather. 2021. "Islam, gender segregation, and political engagement: Evidence from an experiment in Tunisia." *Political Science Research and Methods* 9(4): 728-744. `https://doi.org/10.1017/psrm.2020.37`
- Krishnarajan, Suthan. 2023. "Rationalizing Democracy: The Perceptual Bias and (Un)Democratic Behavior." American Political Science Review 117(2): 474–96. `https://doi.org/10.1017/S0003055422000806`
- Cheema, Ali, Sarah khan, Asad Liaqat, and Shandana Khan Mohmand. 2023. "Canvassing the Gatekeepers: A Field Experiment to Increase Women Voters' Turnout in Pakistan." American Political Science Review 117(1): 1–21. doi: 10.1017/S0003055422000375. `https://doi.org/10.1017/S0003055422000375`

Week 2: *Regression.*

- Cunningham, Chapter 2.
- Huntington-Klein, Chapters 12-13.

- Keele, Luke, Randolph T. Stevenson, and Felix Elwert. 2020. "The Causal Interpretation of Estimated Associations in Regression Models." *Political Science Research and Methods* 8(1): 1–13. doi: 10.1017/psrm.2019.31.

Examples:

- Noam Lupu and Leonid Peisakhin, 2017, "The Legacy of Political Violence across Generations." *American Journal of Political Science* 61 (Oct.): 836-51.
- Sergi Pardos-Prado, 2015, "How Can Mainstream Parties Prevent Niche Party Success? Center-Right Parties and the Immigration Issue." *Journal of Politics* 77 (Feb.): 352-67.
- Rambotti, S., and Breiger, R. L. 2020. "Extreme and Inconsistent: A Case-Oriented Regression Analysis of Health, Inequality, and Poverty." *Socius*. `https://doi.org/10.1177/2378023120906064`

Week 3: *Matching.*

- Cunningham, Chapter 5.
- Huntington-Klein, Chapter 14.
- King, Gary, and Richard Nielsen. 2019. "Why Propensity Scores Should Not Be Used for Matching." *Political Analysis* 27(4): 435–54. doi: 10.1017/pan.2019.11.

Examples:

- Andrea Ruggeri, Han Dorussen, and Theodora-Ismene Gizelis, 2017, "Winning the Peace Locally: UN Peacekeeping and Local Conflict." *International Organization* 71 (Winter): 163-85.
- Lilliana Mason, 2015, "'I Disrespectfully Agree': The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59 (Jan.): 128-45.
- Mark Hill. 2020. "Traditional and Non-traditional Forms of Political Participation in Multigenerational Households." Working paper. `https://osf.io/2nsbh/download`

Week 4: *Natural Experiments I.*

- Dunning, *Natural Experiments in the Social Sciences*, Chapters 5, 6, 8-10.

Examples:

- Allison Carnegie and Nikolay Marinov, 2017, "Foreign Aid, Human Rights, and Democracy Promotion: Evidence from a Natural Experiment." *American Journal of Political Science* 61 (July): 671-83.
- Silva, Bruno Castanho, and Sven-Oliver Proksch. 2021. "Fake It 'Til You Make It: A Natural Experiment to Identify European Politicians' Benefit from Twitter

Bots." American Political Science Review 115(1): 316–22. `https://doi.org/10.1017/S0003055420000817`
- Umair Khalil, Sulagna Mookerjee, and Ryan Tierney. 2019. "Social interactions in voting behavior: Evidence from India." *Journal of Economic Behavior & Organization* 163: 158-171.

Week 5: *Natural Experiments II.*

- Cunningham, Chapter 5.
- Huntington-Klein, Chapters 19-20.

Examples:

- Jeremy Ferwerda and Nicholas L. Miller, 2014, "Political Devolution and Resistance to Foreign Rule: A Natural Experiment." *American Political Science Review* 108 (Aug.): 642-60.
- Matthew A. Kocher and Nuno P. Monteiro, 2016, "Lines of Demarcation: Causation, Design-Based Inference, and Historical Research." *Perspectives on Politics* 14 (Dec.): 952-75.
- Fernanda Brollo and Tommaso Nannicini, 2012, "Tying Your Enemy's Hands in Close Races: The Politics of Federal Transfers in Brazil." *American Political Science Review* 106 (Nov.): 742-61.

Week 6: *Difference-in-Differences.*

- Cunningham, Chapters 8-9.
- Huntington-Klein, Chapters 16-18.

Examples:

- Kevin B. Smith, John R. Alford, John R. Hibbing, Nicholas G. Martin, and Peter K. Hatemi, 2017, "Intuitive Ethics and Political Orientations: Testing Moral Foundations as a Theory of Political Ideology." *American Journal of Political Science* 61 (April): 424-37.
- Andy Baker, Barry Ames, Anand E. Sokhey, Lucio R. Renno, 2015, "The Dynamics of Partisan Identification When Party Brands Change: The Case of the Workers Party in Brazil." *Journal of Politics* 78 (Oct.): 197-213.
- Edmund J. Malesky, Cuong Viet Nguyen, and Anh Tran, 2014, "The Impact of Recentralization on Public Services: A Difference-in-Differences Analysis of the Abolition of Elected Councils in Vietnam." *American Political Science Review* 108 (Feb.): 144-68.

Week 7: *Synthetic Controls.*

- Cunningham, Chapter 10.

- Huntington-Klein, Chapter 15.

Examples:

- Luke Keele and William Minozzi, 2013, "How Much Is Minnesota Like Wisconsin? Assumptions and Counterfactuals in Causal Inference with Observational Data." *Political Analysis* 21 (Spring): 193-216.
- Mourtgos, SM, Adams, IT, Nix, J. (2022). "Elevated police turnover following the summer of George Floyd protests: A synthetic control study." *Criminology and Public Policy* 21: 9-33. `https://doi.org/10.1111/1745-9133.12556`

Week 8: *Machine Learning for Causal Inference.*

- Mario Molina and Filiz Garip. 2019. "Machine Learning for Sociology." *Annual Review of Sociology* 45 (1): 27-45.
- Stefan Wager and Susan Athey, 2017 "Estimation and Inference of Heterogeneous Treatment Effects using Random Forests." Journal of the American Statistical Association `https://www.tandfonline.com/doi/full/10.1080/01621459.2017.1319839`
- Naoki Egami, Christian J. Fong, Justin Grimmer, Margaret E. Roberts, and Brandon M. Stewart, 2018. "How to Make Causal Inferences Using Texts." `https://arxiv.org/abs/1802.02163`
- Huntington-Klein, Chapter 19.

Examples:

- Robert A Blair, Christopher Blattman, and Alexandra Hartman, 2017, "Predicting Local Violence: Evidence from a Panel Survey in Liberia." *Journal of Peace Research* 54 (2): 298-312.
- In Song Kim, 2017, "Political Cleavages within Industry: Firm-level Lobbying for Trade Liberalization." *American Political Science Review* 111 (1): 1-20.
- ANASTASOPOULOS, L. JASON, and ANTHONY M. BERTELLI. 2020. "Understanding Delegation Through Machine Learning: A Method and Application to the European Union." *American Political Science Review* 114(1): 291–301. doi: 10.1017/S0003055419000522.

Week 9: *Selection and Missing Data.*

- Morgan and Winship, Chapter 8.
- Christopher Winship and Robert D. Mare, 1992, "Models for Sample Selection Bias," *Annual Review of Sociology.*
- Arel-Bundock, Vincent, and Krzysztof J. Pelc. 2018. "When Can Multiple Imputation Improve Regression Estimates?" *Political Analysis* 26 (2): 240–45. doi: 10.1017/pan.2017.43.

- Pepinsky, Thomas B. 2018. "A Note on Listwise Deletion versus Multiple Imputation." *Political Analysis* 26(4): 480–88. doi: 10.1017/pan.2018.18.

Examples:

- Vladimir Gimpelson and Daniel Treisman, 2018, "Misperceiving Inequality." *Economics & Politics* 30 (March): 27-54.
- Ranjit Lall, 2017, "The Missing Dimension of the Political Resource Curse Debate." *Comparative Political Studies* 50 (10): 1291-1324.
- Zimran, Ariell. 2019. "Sample-Selection Bias and Height Trends in the Nineteenth-Century United States." *The Journal of Economic History* 79(1): 99–138. doi: 10.1017/S0022050718000694.

Department of Political Science, Northwestern University

Email address: `j-seawright@northwestern.edu`