HARVARD T.H. CHAN SCHOOL OF PUBLIC HEALTH

Biostatistics 258 Spring 2025 Syllabus
Causal Inference: Theory and Practice

credit hours: 4.0 (via FAS) or 5.0 (via HSPH)
lecture meets TuTh 02:00–03:30PM, FXB G03
laboratory meets F 02:00–03:30PM, Kresge LL6

Faculty Instructor: Nima Hejazi, PhD
E-mail: nhejazi@hsph.harvard.edu
Webpage: https://nimahejazi.org
Office Location: HSPH 1-411
Office Hours: Th 01:00–02:00PM

Teaching Fellow: Alexander Mercier
E-mail: amercier@g.harvard.edu
Office Location: HSPH 2-434
Office Hours: Tu 03:30–4:45PM

The instructional staff reserve the right to make changes to this syllabus at any time.

Course Description: Randomized experimentation is the gold standard for the measurement of the causal effect of an intervention (i.e., treatment, exposure) in the public health and biomedical sciences; however, randomization is often impossible, impractical, or unethical, leading to real-world scenarios in which causal inferences are drawn from observational studies. This course will review the foundations of causal inference in (bio)statistics, outlining causal-analytic methods that help to extract as much evidence as imperfect observational studies carry about causal effects commonly of interest in applied health science settings. This doctoral-level survey of statistical causal inference will introduce a structured analytic roadmap for formulating causal (i.e., counterfactual) effect measures tied to clearly defined and well-articulated scientific queries. Such a roadmap begins with a formal model encoding the temporal ordering of variables in a system under study and the possibly *a priori*-known causal relationships between these variables. This is then followed by giving a clear mathematical formulation of the causal effect measure necessary to answer the scientific question of interest. Finally, state-of-the-art mathematical and statistical techniques must be applied to derive best-in-class estimators of the causal effect measures of interest, accompanied by valid statistical inference. Emphasis will be placed on understanding why a formal theory of causation is necessary and how intuition alone or the rote application of traditional statistical modeling frameworks can lead to logical mistakes that invalidate a data analysis and undermine its scientific conclusions.

This tour of statistical causal inference begins with foundational concepts: counterfactual random variables, the potential outcomes and graphical modeling frameworks (e.g., directed acyclic graphs), and necessary assumptions and common strategies for identification of the causal effects of static interventions. Building on these foundations, we will discuss elements of semi-parametric efficiency theory necessary to formulate asymptotically efficient (e.g., augmented inverse probability weighted, targeted maximum likelihood) estimators of causal effect estimands. Theory for studying the causal effects of time-varying, dynamic interventions (e.g., marginal structural models) will be touched upon as well. Further topics to be addressed include causal mediation analysis, heterogeneous treatment effects and optimal dynamic treatment regimes, the causal dose-response curve, and modified treatment policies. Time permitting, additional topics motivating current research will be introduced too. Wherever possible, this course will emphasize the application of non-parametric regression and machine learning tools in the estimation of causal effects.

Prerequisites and Recommendations:

This doctoral-level course is designed for students who are already equipped with a foundational understanding of probability theory and mathematical statistics, working fluency with scientific programming languages (e.g., `R`, numerical `Python`, `Julia`), and working proficiency with tools and best practices for scientific programming, including version control with `git`.

- *Pre-requisites:* BST 231 (Statistical Inference I), BST 232 (Methods I), and BST 233 (Methods II), or equivalent; experience in statistical computing/programming; and mathematical maturity.
- *Recommended:* EPI 207 (Advanced Epidemiologic Methods); BST 241 (Statistical Inference II).
- *Related/Similar:* STAT 186 (Introduction to Causal Inference); STAT 286 (Causal Inference with Applications); STAT 234 (Sequential Decision Making).

Acknowledging diversity of academic backgrounds, the instructional staff suggest that all students review key concepts and software tools in the initial weeks of instruction in a self-directed manner. While self-motivation and the continuous pursuit of learning are highly valued, we recognize the importance of support and collaboration. Each student is encouraged to actively engage in their own learning process, balancing independent growth with the readiness to seek *and offer* help.

Requirements and Materials:

Lecture notes, covering the topics discussed in class, will be distributed each week as the course progresses. A list of articles from the primary and secondary literature will also be provided. The course will not rely upon a single text, but several have been written on this topic, addressing it from a variety of angles. Below is a non-exhaustive list of well-written monographs, given in reverse chronological order of their publication. Students are *not required* to purchase any single text.

- Hernán and Robins (2024), *Causal Inference: What If*
- van der Laan and Rose (2018), *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*
- Peters et al. (2017), *Elements of Causal Inference: Foundations and Learning Algorithms*
- Pearl et al. (2016), *Causal Inference in Statistics: A Primer*
- VanderWeele (2015), *Explanation in Causal Inference: Methods for Mediation and Interaction*
- van der Laan and Rose (2011), *Targeted Learning: Causal Inference for Observational and Experimental Data*
- Pearl (2009), *Causality: Models, Reasoning, and Inference*
- Freedman (2009), *Statistical Models: Theory and Practice*
- Angrist and Pischke (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*
- van der Laan and Robins (2003), *Unified Methods for Censored Longitudinal Data and Causality*

In the first half of the course, you may also find it interesting to read Pearl and Mackenzie (2018)'s *The Book of Why*, a recently published book on the emerging science of causal inference.

Some elements of semi-parametric efficiency theory, especially the theory of influence functions, have become major tools for the formulation of asymptotically efficient estimators of causal effect estimands, and these will be introduced and reviewed with an eye towards their application in causal inference. For those inclined to dive deeper into this area, some relevant texts include

- Kosorok (2008), *Introduction to Empirical Processes and Semiparametric Inference*
- Tsiatis (2007), *Semiparametric Theory and Missing Data*
- van der Vaart (1998), *Asymptotic Statistics*

Pedagogic note: This is a doctoral-level course intended to equip students with the necessary background to begin pursuing research in causal inference. A critical aspect of what makes research challenging is that one does not know in advance the answer to a problem being worked on, or, much more frustratingly, that any satisfying answer even exists. A key skill is to learn to be at peace when working under uncertainty, and to be comfortable in making guesses and pursuing these to wherever they may lead. To build this skill, some of the problems assigned will pose conjectures without indicating whether one should pursue an answer or seek a counterexample. When stuck, try reversing your chosen strategy. While you probably will not find the answer to every question—and that's alright!—you will develop a sense for what it means to conduct research in an area unfamiliar to you. This will prove to be a valuable, and highly transferable, skill. Educational research has shown that encountering challenges and feeling stuck is often when the deepest learning occurs, so embrace such moments as invaluable opportunities to grow and to sharpen understanding.

Laboratory section: In the laboratory section, students will actively apply the causal inference methodologies introduced in the lecture meetings. The aim is to go beyond abstract discussions, allowing students hands-on, practical opportunities to implement causal inference techniques *in a collaborative setting* to answer scientific questions while emphasizing transparency and reproducibility, as needed in real-world projects. The laboratory exercises are designed not only to supplement lecture meetings but to facilitate opportunities to explore select topics in depth and to facilitate guided practice in the application of the formal methodological framework of causal inference to scientific problems. The instructional staff strongly believe that proficiency in the use of open-source software tools for version control (e.g., `git`, GitHub) and literate computing (e.g., Quarto, Jupyter notebooks) is crucial for the modern applied statistician—that is, these tools constitute the applied statistician's "workbench" and their proper setup a form of *mise en place*, helping to organize work in a manner that improves efficiency and enforces the "laboratory hygiene" necessary to ensure clear, reproducible results are obtained consistently. Laboratory exercises should be completed using such tools. Since competency in the use of these tools is expected from the course's start, additional time and effort may be required to develop the expected skill set; we will provide references to tutorials and provide guidance to support this learning process. Hands-on experience in the laboratory section aims to enhance and reinforce understanding of causal inference concepts and to equip students with skills that are vital for the responsible practice of applied statistics.

Learning Outcomes and Objectives:

At the completion of this course, students will have learned. . .

- **I** To translate scientific questions of interest into causal inference questions (i.e., written in the language of counterfactuals) and to mathematically express counterfactual questions, including via causal diagrams expressing *a priori* subject matter knowledge and assumptions.
- **II** To assess when it is possible to learn a causal effect from experimental or observational studies and to state and evaluate common identification assumptions in appropriate scientific context.
- **III** To understand the inferential obstacles posed by intermediate and time-varying confounders and appreciate how formal frameworks, including for principal stratification, causal mediation analysis, and sequential adjustment or re-weighting, can help to address these issues.
- **IV** To apply concepts from semi-parametric efficiency theory to construct and evaluate state-of-the-art, asymptotically efficient estimators of target causal estimands (e.g., average treatment effect, average treatment effect on the treated, natural direct and indirect effects).
- **V** To implement both simple (e.g., inverse probability weighted) and robust and efficient (e.g., augmented inverse probability weighted) estimation approaches to compute a given target causal estimand (e.g., average treatment effect) with the data at hand.

VI To interpret the results of implementing causal analytic approaches in statistical data analysis and to understand when to give up on pursuing causal interpretations of study findings.

Grading Scheme:

- *Course Project* ***(35%):*** An in-depth exploration (or "deep dive") into a single or closely related set of topics in statistical causal inference. Discoveries will be shared as a 10-page report and a 40-minute oral group presentation. We envision this as a collaborative endeavor in which small teams (of 2-4 learners) invest 3-4 weeks of dedicated effort into focused exploration. Guidance on expectations and a range of suggested topics will be shared at the semester's midpoint. Please note that the presentation of previously conducted research—whether past or in-progress—falls strictly outside of the scope of this assignment.
- *Assignments* ***(28%):*** Seven problem sets designed to highlight key concepts by including a range of conceptual, mathematical, and computational exercises will be distributed across the term. Any late submissions will be subjected to a 20% overall deduction per day late.
- *Presentation* ***(17%):*** In lieu of a mid-term examination, students will, in small groups (of 2–3 learners), have the opportunity to delve into a published research manuscript and to lead an in-depth discussion of it during a lecture meeting. Presenters should carefully and closely read the content and prepare presentation materials that summarize key ideas, stimulate discussion, and engage their peers in critical reflections.
- *Participation* ***(10%):*** Students' active participation in class is highly valued. This encompasses regular attendance and substantive contributions to in-class discussions and activities. Engagement and understanding will be assessed occasionally via brief, open-note "concept checks" distributed at random during the lecture and/or laboratory meetings. Successful completion of the mid-term course evaluation is also expected.
- *Journal Club* ***(10%):*** Most weeks, students will read a manuscript assigned from the primary literature and distill their insights via brief write-ups that will guide in-class discussions. Specific expectations will be made available as the term progresses.

Course Policies and Expectations

Accommodations: Please make sure to speak with the Student Support Services staff (studentsupport@hsph.harvard.edu) as soon as possible if you may require any particular accommodations, and they will work out any necessary arrangements as best as possible. We are committed to making feasible adjustments to support educational success for all.

Scheduling Conflicts: Please notify the instructional staff by the second week of the term about any known or potential conflicts, such as religious observances or job interviews. If you foresee the possibility of missing the equivalent of two or more weeks of class, we recommend postponing enrollment in the course to a future term when your full engagement with the course material will be possible. Regular attendance and participation are vital for a complete learning experience.

Collaboration and Independence: Learning is a collective journey, so we encourage you to work together on homework assignments. That said, all homework assignment submissions should clearly list collaborators and references and should be written up independently. Submitted assignments will not be considered for credit if they are a replicate of another submission; further, such forms of academic misrepresentation may trigger disciplinary action. Use of `ChatGPT` and the like is strongly discouraged—*caveat emptor*—and we encourage you to practice forging your own insights and testing your own understanding. These guidelines are designed to support a constructive and inclusive learning experience for all course participants.

Key Dates

Unless otherwise stated, all deliverables are due by 5:00pm EST on the specified date.

Problem set 1	Thursday, 20th February
Problem set 2	Thursday, 6th March
Problem set 3	Thursday, 13th March
Midterm Feedback Survey	Thursday, 13th March
Problem set 4	Thursday, 27th March
Problem set 5	Thursday, 10th April
Problem set 6	Thursday, 24th April
Problem set 7	Thursday, 8th May
Final Project Report	Monday, 12th May
Final Project Presentations	Week of 12th May

Class Session Structure

The primary class sessions focus on lecture-oriented *synchronous learning*. Each of the 90-minute learning sessions will be divided unevenly into an interactive, review component (Part I) and a lecture component (Part II), as elaborated upon below.

- Part I (20 minutes) begins with a brief 3-5 question, open-note "concept check" intended to review previously covered material and cover any assigned reading material. A random selection of students will present their answers, which will prompt a brief in-class discussion.
- Part II (70 minutes) is dedicated to a traditional lecture delivered by the faculty instructor, another member of the instructional staff (only sparingly) or, if appropriate, an invited guest lecturer. This will include time for questions from and discussion with the audience.

The laboratory sessions are a crucial part of the learning experience, providing an interactive environment to apply concepts introduced in the lectures through hands-on activities and exercises.

Course Outline

The course is divided into a series of modules. The weekly coverage of topics is subject to change, as it will depend on the progress of the class. Anticipated time for the completion of each module is indicated below. As the semester runs 16 weeks (n.b., Spring 2024 runs 22 January–10 May with Spring Break 11–15 March), we will cover only a selection of the modules listed. To strive for a reasonably comprehensive introduction to some core topics in statistical causal inference, part of the menu—the underlined modules—is served *prix fixe*. These will be supplemented by other modules to be selected based on the audiences' indicated interests and at the instructor's discretion.

1. **Introduction to and overview of causal inference.** Randomized controlled experiments, observational studies, and the pitfalls of trying to draw causal inferences by rote application of traditional statistical methods (i.e., regression); a roadmap for causal inference.
    - Time anticipated: 1 week (2 lectures)
    - Learning objectives: **I**, **II**, **VI**
    - Topics: Measures of association (risk difference, risk ratio, odds ratio, difference-in-means, regression coefficients); randomization as control (i.e., probability by fiat) and the design-based perspective; observational studies (i.e., probability "from nature") and their shortcomings; the Yule-Simpson paradox and some real-life examples; perspectives on causality, statistics, and ontological commitments in applied statistical science.

- Books and tutorials: Hernán and Robins (2024, Ch. 1–3), Freedman (2009, Ch. 1–2, 4), Pearl (2009, Ch. 1), Pearl et al. (2016, Ch. 1), Angrist and Pischke (2009, Ch. 1–2, Sec. 3.1), Pearl and Mackenzie (2018, Ch. 1–3), Starmans (2018)
- Primary literature: Rubin (1974), Bickel et al. (1975), Holland (1986), Freedman (1999), Hernán and Taubman (2008), Petersen and van der Laan (2014), Vansteelandt (2021)

2. **Potential outcomes, graphical models, and identification strategies.** Links between missing data and causal inference; structural causal models and graphical models; assumptions for identification of causal effects; g-computation and inverse probability weighting.
   - Time anticipated: 3-4 weeks (6-8 lectures)
   - Learning objectives: **I**, **II**, **V**, **VI**
   - Topics: Week 1: Intervention and outcome variables; counterfactual random variables; the Neyman-Rubin potential outcomes framework and the stable unit treatment value assumption; overview of graphical modeling frameworks via single-world intervention graphs. Week 2: The positivity and overlap assumptions; the ignorability assumption and no unmeasured confounding; the propensity score; the g-computation algorithm. Week 3: Estimation of causal effects using inverse probability weighting and outcome regression modeling; perils of model misspecification and strategies to circumvent.
   - Books and tutorials: Hernán and Robins (2024, Ch. 6–8, 10, 12, 13, 15), van der Laan and Rose (2011, Ch. 2), Pearl (2009, Ch. 3, 6), Pearl et al. (2016, Ch. 2, Sec. 3.1–3.6, 4.1–4.4), Angrist and Pischke (2009, Sec. 3.2), Pearl and Mackenzie (2018, Ch. 4–8)
   - Primary literature: Rosenbaum and Rubin (1983), Greenland and Robins (1986), Pearl (1995), Greenland et al. (1999a), Greenland et al. (1999b), Rubin (2005), Pearl (2010), Hubbard and van der Laan (2008), Robins et al. (2007), Petersen et al. (2012)

3. **Semiparametric efficiency theory in causal inference and causal machine learning.** Infinite-dimensional or large statistical models; semi-parametric local efficiency; the efficient influence function and its key properties (e.g., double robustness); estimation strategies based on the efficient influence function; considering the role of machine learning in causal inference.
   - Time anticipated: 3 weeks (6 lectures)
   - Learning objectives: **IV**, **V**
   - Topics: Week 1: Nonparametric statistical models incorporating real-world knowledge; influence functions and asymptotic linearity (revisited); the efficient influence function and semi-parametric efficiency. Week 2: Alleviating model misspecification via machine learning (cross-validation, loss-based estimation, and the super learner algorithm); cross-fitting and the role of regularity conditions. Week 3: Asymptotic curse of dimensionality and asymptotic bias-correction; asymptotically efficient estimation based on the one-step correction, augmented inverse probability weighting, and targeted maximum likelihood.
   - Books and tutorials: Kennedy (2016), Fisher and Kennedy (2020), Kennedy (2024), Hines et al. (2022), van der Laan and Rose (2011, Ch. 4–6, Appx. A), van der Vaart (1998, Ch. 7–9, 19, 20), Tsiatis (2007, Ch. 2–3)
   - Primary literature: Bang and Robins (2005), van der Laan and Rubin (2006), Rubin and van der Laan (2007), Tsiatis et al. (2008), Rubin and van der Laan (2008), Moore and van der Laan (2009), Rubin and van der Laan (2011), Wang et al. (2019), Vansteelandt et al. (2010), Gruber and van der Laan (2015), Schnitzer et al. (2016), Hahn (1998), Hirano et al. (2003), Hejazi and van der Laan (2023), Zheng and van der Laan (2011), Chernozhukov et al. (2018), Ju et al. (2019), van der Laan et al. (2004), Dudoit and van der Laan (2005), van der Laan et al. (2007), Wyss et al. (2018), Phillips et al. (2023)

4. **The same in a relative way: Time-varying interventions and confounding feedback.** Confounding due to dependencies between time-varying covariates and treatment schedules across time; causal effects of dynamic and time-varying intervention schemes; identification using the longitudinal g-computation algorithm and by inverse probability weighting; overview of marginal structural models and their properties.
   - Time anticipated: 2 weeks (4 lectures)
   - Learning objectives: **II**, **III**, **IV**, **V**, **VI**
   - Topics: Week 1: Treatment-confounder feedback and the shortfall of "typical" methods; defining causal effects of time-varying interventions; extensions of g-computation and inverse probability weighting, and estimation based on these identification strategies; overview of the construction of efficient estimators. Week 2: Formulation of marginal structural models (MSMs) and longitudinal MSMs, MSMs as working projections, and efficient estimation of the parameters of an MSM.
   - Books and tutorials: Hernán and Robins (2024, Ch. 19–21), VanderWeele (2015, Ch. 6)
   - Primary literature: Robins (1986), Tsiatis et al. (2011), Rotnitzky et al. (2012), van der Laan and Gruber (2012), Luedtke et al. (2017), Rotnitzky et al. (2017), Robins et al. (2000), Hernán et al. (2000) Neugebauer and van der Laan (2007), Cole and Hernán (2008), Rosenblum and van der Laan (2010), Schnitzer et al. (2014a), Schnitzer et al. (2014b)

5. **Caught in the middle: Principal stratification and mediation analysis.** Issues that arise due to intermediate variables; principal strata; mediation as effect decomposition; well-known direct and indirect effect definitions and estimands; the interventionist perspective.
   - Time anticipated: 3 weeks (6 lectures)
   - Learning objectives: **II**, **III**, **V**, **VI**
   - Topics: Week 1: Post-treatment confounding due to intermediate variables; principal stratification and principal stratum causal effects; brief history of mediation analysis in statistics (e.g., path analysis). Week 2: Nonparametric identification of mediation effects (i.e., sequential ignorability, nested potential outcomes, "cross-world" counterfactuals); the controlled direct effect; the natural direct and indirect effects. Week 3: Estimation of direct and indirect effects; separability criteria and interventionist mediation analysis; intermediate confounding and the interventional direct and indirect effects.
   - Books and tutorials: Hernán and Robins (2024, Ch. 4–5, 23), Freedman (2009, Ch. 6), VanderWeele (2015, Ch. 1–2, 5–6, 8), Pearl (2009, Sec. 4.5), Pearl et al. (2016, Sec. 3.7, 4.5), Pearl and Mackenzie (2018, Ch. 9)
   - Primary literature: Frangakis and Rubin (2002), Rubin (2006), Hudgens and Halloran (2006), Jemiai et al. (2007), VanderWeele (2011), Pearl (2011), Gilbert et al. (2011), Dawid and Didelez (2012), Tchetgen Tchetgen (2014), VanderWeele (2008), Robins and Greenland (1992), Pearl (2001), Avin et al. (2005), Didelez et al. (2006), Petersen et al. (2006), van der Laan and Petersen (2008), VanderWeele and Vansteelandt (2009), Zheng and van der Laan (2012), VanderWeele and Vansteelandt (2014), Andrews and Didelez (2020), Díaz and Hejazi (2020), Robins et al. (2022), Miles (2023), Díaz (2023)

6. **Born this way: Treatment effect heterogeneity and personalized interventions.** Causal effect heterogeneity and its connections to both effect modification and statistical interaction; the conditional average treatment effect and optimal dynamic treatment regimes.
   - Time anticipated: 2 weeks (4 lectures)
   - Learning objectives: **II**, **IV**, **V**

- Topics: Week 1: Measures of heterogeneity (i.e., interaction, effect modification) and differences between the associational and causal perspectives; the conditional average treatment effect (CATE)—challenges for identification and inference. Week 2: Optimal dynamic treatment regimes and tailoring rules based on expected benefit or harm (i.e., assignment based on the CATE); challenges posed by individualized treatment rules.
- Books and tutorials: VanderWeele (2015, Ch. 9–10)
- Primary literature: Bland and Altman (2011), Moodie et al. (2007), Murphy (2003), Chakraborty et al. (2010), Zhang et al. (2012), Laber et al. (2014), Luedtke and van der Laan (2016), Luedtke and van der Laan (2017), Wager and Athey (2018), VanderWeele et al. (2019b), Nie and Wager (2021), van der Laan et al. (2024), Boileau et al. (2025), van der Laan and Petersen (2007), Bembom and van der Laan (2007), Qian and Murphy (2011), Zhao et al. (2012), Qiu et al. (2022)

7. **Dose makes the poison: The dose-response curve and modified treatment policies.** Defining causal effects suited to ordinal and continuous treatment variables; the causal dose-response curve; effects of interventions that would modify the treatment actually received.
   - Time anticipated: 2 weeks (4 lectures)
   - Learning objectives: **II**, **IV**, **V**
   - Topics: Week 1: Defining and interpreting contrasts for continuous treatment variables; inferential perils of dichotomizing continuous treatment variables; on the definition and identification of the causal dose-response curve (CDRC) and some key challenges (e.g., structural positivity violations, pathwise differentiability). Week 2: Inference for the CDRC via (1) marginal structural models (as "working" projections) and (2) efficient estimation leveraging semi-parametric theory; definition and identification of the causal effects of modified treatment policies (MTPs); inference for the causal effects of MTPs.
   - Books and tutorials: Hernán and Robins (2024, Sec. 12.4)
   - Primary literature: Altman and Royston (2006), Imbens (2000), Hirano and Imbens (2004), Imai and Van Dyk (2004), Díaz and van der Laan (2012), Haneuse and Rotnitzky (2013), Young et al. (2014), Díaz et al. (2021), Kennedy et al. (2017), Westling et al. (2020), Westling (2022), van der Laan et al. (2023)

8. **Buying a stairway to heaven: Instrumental variables and non-identifiability.** Using an instrument for point- and set-identification; estimation of bounds in instrumental variables models; Mendelian randomization; negative controls and proximal causal inference.
   - Time anticipated: 2 weeks (4 lectures)
   - Learning objectives: **II**, **VI**
   - Topics: Week 1: Brief history of instrumental variables; using instruments to address issues of measurement error and imperfect compliance; identification assumptions in the use of instruments (e.g., monotonicity); connections with principal stratification; the local average treatment effect curve and its nonparametric identification. Week 2: Biological phenomena as instrumental variables (i.e., Mendelian randomization)—some promises and pitfalls; negative controls as instrumental variables when some confounders may be unmeasured, and connections to the framework of proximal causal inference.
   - Books and tutorials: Hernán and Robins (2024, Ch. 16), Freedman (2009, Ch. 9), Pearl (2009, Ch. 8), Angrist and Pischke (2009, Ch. 4)
   - Primary literature: Imbens and Angrist (1994), Angrist et al. (1996), Balke and Pearl (1997), Baiocchi et al. (2014), Ogburn et al. (2015), Small et al. (2017), Kennedy et al.

(2019), Davey Smith and Ebrahim (2003), Didelez and Sheehan (2007), VanderWeele et al. (2014), Lipsitch et al. (2010), Miao et al. (2018)

9. **Where are your friends tonight? Interference, contagion, and spillover.** Violations of the stable unit treatment value assumption; between-unit dependence structures and their complications; identification of causal effects under common dependence structures.
   - Time anticipated: 2 weeks (4 lectures)
   - Learning objectives: **II**, **VI**
   - Topics: Week 1: The no-interference part of the stable unit treatment value assumption and its plausibility; extensions of study designs and graphical models to accommodate interference; contagion versus interference; clustered treatment assignment and spillover causal effects. Week 2: Networks and network interference; relaxing the no-interference assumption in the identification of causal effects; overview of complications for inference on causal effect estimands when the no-interference assumption is weakened.
   - Books and tutorials: VanderWeele (2015, Ch. 14–15)
   - Primary literature: Rosenbaum (2007), Hudgens and Halloran (2008), Tchetgen Tchetgen and VanderWeele (2012), Ogburn and VanderWeele (2014), van der Laan (2014), Carnegie et al. (2016), Ogburn and VanderWeele (2017), Forastiere et al. (2021), Sävje et al. (2021), Zivich et al. (2022), Lee et al. (2022), Ogburn et al. (2024)

10. **Yeah, you wanted the truth: Sensibility and sensitivity analysis.** How randomization validates RCTs; epistemological uncertainty in drawing causal inferences from observational studies; sensitivity analysis in statistics and epidemiology; some assumption-lean approaches.
    - Time anticipated: 1 week (2 lectures)
    - Learning objectives: **II**, **VI**
    - Topics: R.A. Fisher's smoking–cancer controversy; when association really is causation; the Cornfield conditions in epidemiology; the "causal gap" and sensitivity analysis via its quantification; model-agnostic approaches to sensitivity analysis and the E-value.
    - Books and tutorials: N/A
    - Primary literature: Fisher (1957), Cornfield et al. (1959), Gilbert et al. (2003), Brumback et al. (2004), Shepherd et al. (2006), Díaz and van der Laan (2013), Ding and VanderWeele (2016), VanderWeele and Ding (2017), VanderWeele et al. (2019a), Haneuse et al. (2019), Zhao et al. (2019), Díaz et al. (2023)

Harvard Chan School Policies and Expectations

Inclusivity Statement

Diversity and inclusiveness are fundamental to public health education and practice. It is a requirement that you have an open mind and respect differences of all kinds. We share responsibility with you for creating a learning climate that is hospitable to all perspectives and cultures; please contact us if you have any concerns or suggestions.

Bias Related Incident Reporting

The Harvard Chan School believes all members of our community should be able to study and work in an environment where they feel safe and respected. As a mechanism to promote an inclusive community, we have created an anonymous bias-related incident reporting system. If you have experienced bias, please submit a report here so that the administration can track and address concerns as they arise and to better support members of the Harvard Chan School community.

Title IX

For information on Harvard University policies and procedures and Title IX Resource Coordinators at the Harvard Chan School, please see:

- Harvard University Interim Title IX Sexual Harassment and Interim Other Sexual Misconduct policies and procedures: https://titleix.harvard.edu/policies-procedures
- Title IX Resource Coordinators: https://titleix.harvard.edu/coordinators
- Title IX Sexual Harassment and Other Sexual Misconduct resource guide: https://titleix.harvard.edu/resource-guide

Academic Integrity

You are expected to abide by the Harvard University and the Harvard T.H. Chan School of Public Health School's standards of Academic Integrity in conjunction with the expectations outlined in the Course Structure and Assessment of Learning section of this syllabus. All work submitted to meet course requirements is expected to be your own work. In the preparation of work submitted to meet course requirements, you should always take great care to distinguish your own ideas and knowledge from information derived from sources.

You must assume that collaboration in the completion of assignments is prohibited unless explicitly specified. You must acknowledge any collaboration and its extent in all submitted work. This requirement applies to collaboration on editing as well as collaboration on substance.

Should academic misconduct occur, you may be subject to disciplinary action as outlined in the Student Handbook. See the Student Handbook for additional policies related to academic integrity and disciplinary actions.

Accommodations for Students with Disabilities

Harvard University provides academic accommodations to students with disabilities. Any requests for academic accommodations should ideally be made before the first week of the semester, except for unusual circumstances, so arrangements can be made. You must register with a Local Disability Coordinator in the Office for Student Affairs to verify their eligibility for appropriate accommodations. Contact studentsupport@hsph.harvard.edu in all cases, including temporary disabilities.

Absence Due to Religious Holidays

According to Chapter 151c, Section 2B, of the General Laws of Massachusetts, any student in an educational or vocational training institution, other than a religious or denominational training institution, who is unable, because of his or her religious beliefs, to attend classes or to participate in any examination, study, or work requirement on a particular day shall be excused from any such examination or requirement which he or she may have missed because of such absence on any particular day, provided that such makeup examination or work shall not create an unreasonable burden upon the School. See the Student Handbook for more information.

Course Evaluation

Constructive feedback from students is a valuable resource for improving the teaching and learning experience. The feedback should be specific, focused, and respectful. It should address aspects of the course and teaching that are positive, as well as those which need improvement.

For registered students, submission of course evaluations is considered to be a School requirement because of their importance. The course evaluation system will open during the last week of the term and remain open for a three week period. You will gain access to your grades for the term after you have completed your course evaluations, and the course evaluation system has closed.

References

Douglas G Altman and Patrick Royston. The cost of dichotomising continuous variables. *The BMJ*, 332(7549):1080, 2006. doi:10.1136/bmj.332.7549.1080.

Ryan M Andrews and Vanessa Didelez. Insights into the cross-world independence assumption of causal mediation analysis. *Epidemiology*, 32(2):209–219, 2020. doi:10.1097/EDE.0000000000001313.

Joshua D Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009. doi:10.1515/9781400829828.

Joshua D Angrist, Guido W Imbens, and Donald B Rubin. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):444–455, 1996.

Chen Avin, Ilya Shpitser, and Judea Pearl. Identifiability of path-specific effects. In *IJCAI International Joint Conference on Artificial Intelligence*, pages 357–363, 2005.

Michael Baiocchi, Jing Cheng, and Dylan S Small. Instrumental variable methods for causal inference. *Statistics in Medicine*, 33(13):2297–2340, 2014. doi:10.1002/sim.6128.

Alexander Balke and Judea Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997. doi:10.1080/01621459.1997.10474074.

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005. doi:10.1111/j.1541-0420.2005.00377.x.

Oliver Bembom and Mark J van der Laan. A practical illustration of the importance of realistic individualized treatment rules in causal inference. *Electronic Journal of Statistics*, 1:574–596, 2007. doi:10.1214/07-EJS105.

Peter J Bickel, Eugene A Hammel, and J William O'Connell. Sex bias in graduate admissions: Data from Berkeley. *Science*, 187(4175):398–404, 1975. doi:10.1126/science.187.4175.398.

J Martin Bland and Douglas G Altman. Comparisons within randomised groups can be very misleading. *The BMJ*, 342, 2011. doi:10.1136/bmj.d561.

Philippe Boileau, Ning Leng, Nima S Hejazi, Mark J van der Laan, and Sandrine Dudoit. A nonparametric framework for treatment effect modifier discovery in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 87(1):157–185, 2025. doi:10.1093/jrsssb/qkae084.

Babette A Brumback, Miguel A Hernán, Sebastien J P A Haneuse, and James M Robins. Sensitivity analyses for unmeasured confounding assuming a marginal structural model for repeated measures. *Statistics in Medicine*, 23(5):749–767, 2004. doi:10.1002/sim.1657.

Nicole Bohme Carnegie, Rui Wang, and Victor De Gruttola. Estimation of the overall treatment effect in the presence of interference in cluster-randomized trials of infectious disease prevention. *Epidemiologic Methods*, 5(1):57–68, 2016. doi:10.1515/em-2015-0016.

Bibhas Chakraborty, Susan Murphy, and Victor Strecher. Inference for non-regular parameters in optimal dynamic treatment regimes. *Statistical Methods in Medical Research*, 19(3):317–343, 2010. doi:10.1177/0962280209105013.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi:10.1111/ectj.12097.

Stephen R Cole and Miguel A Hernán. Constructing inverse probability weights for marginal structural models. *American Journal of Epidemiology*, 168(6):656–664, 2008. doi:10.1093/aje/kwn164.

Jerome Cornfield, William Haenszel, E Cuyler Hammond, Abraham M Lilienfeld, Michael B Shimkin, and Ernst L Wynder. Smoking and lung cancer: Recent evidence and a discussion of some questions. *Journal of the National Cancer institute*, 22(1):173–203, 1959. doi:10.1093/jnci/22.1.173.

George Davey Smith and Shah Ebrahim. 'Mendelian randomization': Can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, 32(1):1–22, 2003. doi:10.1093/ije/dyg070.

A Philip Dawid and Vanessa Didelez. "Imagine a Can Opener"–the magic of principal stratum analysis. *The International Journal of Biostatistics*, 8(1), 2012. doi:10.1515/1557-4679.1391.

Iván Díaz. Non-agency interventions for causal mediation in the presence of intermediate confounding. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, page qkad130, 2023. doi:10.1093/jrsssb/qkad130.

Iván Díaz and Nima S Hejazi. Causal mediation analysis for stochastic interventions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(3):661–683, 2020. doi:10.1111/rssb.12362.

Iván Díaz and Mark J van der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, 68(2):541–549, 2012. doi:10.1111/j.1541-0420.2011.01685.x.

Iván Díaz and Mark J van der Laan. Sensitivity analysis for causal inference under unmeasured confounding and measurement error problems. *The International Journal of Biostatistics*, 9(2): 149–160, 2013. doi:10.1515/ijb-2013-0004.

Iván Díaz, Nicholas Williams, Katherine L Hoffman, and Edward J Schenck. Nonparametric causal effects based on longitudinal modified treatment policies. *Journal of the American Statistical Association*, 118(542):846–857, 2021. doi:10.1080/01621459.2021.1955691.

Iván Díaz, Hana Lee, Emre Kıcıman, Edward J Schenck, Mouna Akacha, Dean Follman, and Debashis Ghosh. Sensitivity analysis for causality in observational studies for regulatory science. *Journal of Clinical and Translational Science*, 7(1):e267, 2023. doi:10.1017/cts.2023.688.

Vanessa Didelez and Nuala Sheehan. Mendelian randomization as an instrumental variable approach to causal inference. *Statistical Methods in Medical Research*, 16(4):309–330, 2007. doi:10.1177/0962280206077743.

Vanessa Didelez, Philip Dawid, and Sara Geneletti. Direct and indirect effects of sequential treatments. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*, pages 138–146, 2006.

Peng Ding and Tyler J VanderWeele. Sensitivity analysis without assumptions. *Epidemiology*, 27 (3):368–377, 2016. doi:10.1097/EDE.0000000000000457.

Sandrine Dudoit and Mark J van der Laan. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology*, 2(2):131–154, 2005. doi:10.1016/j.stamet.2005.02.003.

Aaron Fisher and Edward H Kennedy. Visually communicating and teaching intuition for influence functions. *The American Statistician*, 75(2):162–172, 2020. doi:10.1080/00031305.2020.1717620.

Ronald A Fisher. Dangers of cigarette-smoking. *British Medical Journal*, 2(5039):297–298, 1957.

Laura Forastiere, Edoardo M Airoldi, and Fabrizia Mealli. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534):901–918, 2021. doi:10.1080/01621459.2020.1768100.

Constantine E Frangakis and Donald B Rubin. Principal stratification in causal inference. *Biometrics*, 58(1):21–29, 2002. doi:10.1111/j.0006-341X.2002.00021.x.

David Freedman. From association to causation: some remarks on the history of statistics. *Journal de la Société Française de Statistique*, 140(3):5–32, 1999.

David A Freedman. *Statistical Models: Theory and Practice*. Cambridge University Press, 2009. doi:10.1017/CBO9780511815867.

Peter B Gilbert, Ronald J Bosch, and Michael G Hudgens. Sensitivity analysis for the assessment of causal vaccine effects on viral load in HIV vaccine trials. *Biometrics*, 59(3):531–541, 2003. doi:10.1111/1541-0420.00063.

Peter B Gilbert, Michael G Hudgens, and Julian Wolfson. Commentary on "Principal stratification — a goal or a tool?" by Judea Pearl. *The International Journal of Biostatistics*, 7(1):1341, 2011. doi:10.2202/1557-4679.1341.

Sander Greenland and James M Robins. Identifiability, exchangeability, and epidemiological confounding. *International Journal of Epidemiology*, 15(3):413–419, 1986. doi:10.1093/ije/15.3.413.

Sander Greenland, Judea Pearl, and James M Robins. Causal diagrams for epidemiologic research. *Epidemiology*, 10(1):37–48, 1999a.

Sander Greenland, Judea Pearl, and James M Robins. Confounding and collapsibility in causal inference. *Statistical Science*, 14(1):29–46, 1999b. doi:10.1214/ss/1009211805.

Susan Gruber and Mark J van der Laan. Consistent causal effect estimation under dual misspecification and implications for confounder selection procedures. *Statistical Methods in Medical Research*, 24(6):1003–1008, 2015. doi:10.1177/0962280212437.

Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66(2):315–331, 1998. doi:10.2307/2998560.

Sebastian Haneuse and Andrea Rotnitzky. Estimation of the effect of interventions that modify the received treatment. *Statistics in Medicine*, 32(30):5260–5277, 2013. doi:10.1002/sim.5907.

Sebastien Haneuse, Tyler J VanderWeele, and David Arterburn. Using the E-value to assess the potential effect of unmeasured confounding in observational studies. *JAMA*, 321(6):602–603, 2019. doi:10.1001/jama.2018.21554.

Nima S Hejazi and Mark J van der Laan. Revisiting the propensity score's central role: Towards bridging balance and efficiency in the era of causal machine learning. *Observational Studies*, 9 (1):23–34, 2023. doi:10.1353/obs.2023.0001.

Miguel A Hernán and James M Robins. *Causal Inference: What If*. CRC Press, 2024.

Miguel A Hernán and Sarah L Taubman. Does obesity shorten life? The importance of well-defined interventions to answer causal questions. *International Journal of Obesity*, 32(S3):S8, 2008. doi:10.1038/ijo.2008.82.

Miguel Ángel Hernán, Babette Brumback, and James M Robins. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*, 11 (5):561–570, 2000. doi:10.1097/00001648-200009000-00012.

Oliver Hines, Oliver Dukes, Karla Diaz-Ordaz, and Stijn Vansteelandt. Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 76(3):292–304, 2022. doi:10.1080/00031305.2021.2021984.

Keisuke Hirano and Guido W Imbens. The propensity score with continuous treatments. In Andrew Gelman and Xiao-Li Meng, editors, *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, pages 73–84. John Wiley & Sons, Ltd, 2004. doi:10.1002/0470090456.ch7.

Keisuke Hirano, Guido W Imbens, and Geert Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003. doi:10.1111/1468-0262.00442.

Paul W Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986. doi:10.1080/01621459.1986.10478354.

Alan E Hubbard and Mark J van der Laan. Population intervention models in causal inference. *Biometrika*, 95(1):35–47, 2008. doi:10.1093/biomet/asm097.

Michael G Hudgens and M Elizabeth Halloran. Causal vaccine effects on binary postinfection outcomes. *Journal of the American Statistical Association*, 101(473):51–64, 2006. doi:10.1198/016214505000000970.

Michael G Hudgens and M Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008. doi:10.1198/016214508000000292.

Kosuke Imai and David A Van Dyk. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467):854–866, 2004. doi:10.1198/016214504000001187.

Guido W Imbens. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000. doi:10.1093/biomet/87.3.706.

Guido W Imbens and Joshua D Angrist. Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475, 1994. doi:10.2307/2951620.

Yannis Jemiai, Andrea Rotnitzky, Bryan E Shepherd, and Peter B Gilbert. Semiparametric estimation of treatment effects given base-line covariates on an outcome measured after a post-randomization event occurs. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(5):879–901, 2007. doi:10.1111/j.1467-9868.2007.00615.x.

Cheng Ju, Susan Gruber, Samuel D Lendle, Antoine Chambaz, Jessica M Franklin, Richard Wyss, Sebastian Schneeweiss, and Mark J van der Laan. Scalable collaborative targeted learning for high-dimensional data. *Statistical Methods in Medical Research*, 28(2):532–554, 2019. doi:10.1177/0962280217729845.

Edward H Kennedy. Semiparametric theory and empirical processes in causal inference. In Hua He, Pan Wu, and Ding-Geng (Din) Chen, editors, *Statistical Causal Inferences and Their Applications in Public Health Research*, pages 141–167. Springer, 2016. doi:10.1007/978-3-319-41259-7_8.

Edward H Kennedy. Semiparametric doubly robust targeted double machine learning: A review. In Eric Laber, Bibhas Chakraborty, Erica E M Moodie, Tianxi Cai, and Mark J van der Laan, editors, *Handbook of Statistical Methods for Precision Medicine*, pages 207–236. Chapman and Hall/CRC, 2024. doi:10.48550/arXiv.2203.06469.

Edward H Kennedy, Zongming Ma, Matthew D McHugh, and Dylan S Small. Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1229–1245, 2017. doi:10.1111/rssb.12212.

Edward H Kennedy, Scott Lorch, and Dylan S Small. Robust causal inference with continuous instruments using the local instrumental variable curve. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(1):121–143, 2019. doi:10.1111/rssb.12300.

Michael R Kosorok. *Introduction to Empirical Processes and Semiparametric Inference*. Springer, 2008. doi:10.1007/978-0-387-74978-5.

Eric B Laber, Daniel J Lizotte, Min Qian, William E Pelham, and Susan A Murphy. Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, 8(1): 1225–1272, 2014. doi:10.1214/14-ejs920.

Chanhwa Lee, Donglin Zeng, and Michael G Hudgens. Efficient nonparametric estimation of stochastic policy effects with clustered interference. *arXiv*, arXiv:2212.10959, 2022.

Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: A tool for detecting confounding and bias in observational studies. *Epidemiology*, 21(3):383–388, 2010. doi:10.1097/EDE.0b013e3181d61eeb.

Alexander R Luedtke and Mark J van der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *The Annals of Statistics*, 44(2):713–742, 2016. doi:10.1214/15-AOS1384.

Alexander R Luedtke and Mark J van der Laan. Evaluating the impact of treating the optimal subgroup. *Statistical Methods in Medical Research*, 26(4):1630–1640, 2017. doi:10.1177/0962280217708664.

Alexander R Luedtke, Oleg Sofrygin, Mark J van der Laan, and Marco Carone. Sequential double robustness in right-censored longitudinal models. *arXiv preprint arXiv:1705.02459*, 2017. doi:10.48550/arXiv.1705.02459.

Wang Miao, Zhi Geng, and Eric J Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018. doi:10.1093/biomet/asy038.

Caleb H Miles. On the causal interpretation of randomised interventional indirect effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 85(4):1154–1172, 2023. doi:10.1093/jrsssb/qkad066.

Erica E M Moodie, Thomas S Richardson, and David A Stephens. Demystifying optimal dynamic treatment regimes. *Biometrics*, 63(2):447–455, 2007. doi:10.1111/j.1541-0420.2006.00686.x.

Kelly L Moore and Mark J van der Laan. Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation. *Statistics in Medicine*, 28(1):39–64, 2009. doi:10.1002/sim.3445.

Susan A Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(2):331–355, 2003. doi:10.1111/1467-9868.00389.

Romain Neugebauer and Mark J van der Laan. Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference*, 137(2):419–434, 2007. doi:10.1016/j.jspi.2005.12.008.

Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021. doi:10.1093/biomet/asaa076.

Elizabeth L Ogburn and Tyler J VanderWeele. Causal diagrams for interference. *Statistical Science*, 29(4):559–578, 2014. doi:10.1214/14-STS501.

Elizabeth L Ogburn and Tyler J VanderWeele. Vaccines, contagion, and social networks. *The Annals of Applied Statistics*, 11(2):919–948, 2017. doi:10.1214/17-AOAS1023.

Elizabeth L Ogburn, Andrea Rotnitzky, and James M Robins. Doubly robust estimation of the local average treatment effect curve. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(2):373–396, 2015. doi:10.1111/rssb.12078.

Elizabeth L Ogburn, Oleg Sofrygin, Ivan Diaz, and Mark J van der Laan. Causal inference for social network data. *Journal of the American Statistical Association*, 119(545):597–611, 2024. doi:10.1080/01621459.2022.2131557.

Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995. doi:10.1093/biomet/82.4.669.

Judea Pearl. Direct and indirect effects. In *Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence*, 2001.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2009. doi:10.1017/CBO9780511803161.

Judea Pearl. On the consistency rule in causal inference: Axiom, definition, assumption, or theorem? *Epidemiology*, 21(6):872–875, 2010. doi:10.1097/EDE.0b013e3181f5d3fd.

Judea Pearl. Principal stratification—a goal or a tool? *The International Journal of Biostatistics*, 7(1):1–13, 2011. doi:10.2202/1557-4679.1322.

Judea Pearl and Dana Mackenzie. *The Book of Why: The New Science of Cause and Effect*. Basic Books, 2018.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal Inference in Statistics: A Primer*. John Wiley & Sons, 2016.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.

Maya L Petersen and Mark J van der Laan. Causal models and learning from data: Integrating causal modeling and statistical estimation. *Epidemiology*, 25(3):418, 2014. doi:10.1097/EDE.0000000000000078.

Maya L Petersen, Sandra E Sinisi, and Mark J van der Laan. Estimation of direct causal effects. *Epidemiology*, pages 276–284, 2006. doi:10.1097/01.ede.0000208475.99429.2d.

Maya L Petersen, Kristin E Porter, Susan Gruber, Yue Wang, and Mark J van der Laan. Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21(1):31–54, 2012. doi:10.1177/0962280210386207.

Rachael V Phillips, Mark J van der Laan, Hana Lee, and Susan Gruber. Practical considerations for specifying a super learner. *International Journal of Epidemiology*, 2023. doi:10.1093/ije/dyad023.

Min Qian and Susan A Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180–1210, 2011. doi:10.1214/10-AOS864.

Hongxiang Qiu, Marco Carone, and Alex Luedtke. Individualized treatment rules under stochastic treatment cost constraints. *Journal of Causal Inference*, 10:480–493, 2022. doi:10.1515/jci-2022-0005.

James M Robins. A new approach to causal inference in mortality studies with sustained exposure periods—Application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512, 1986. doi:10.1016/0270-0255(86)90088-6.

James M Robins and Sander Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(2):143–155, 1992.

James M Robins, Miguel Ángel Hernán, and Babette Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000. doi:10.1097/00001648-200009000-00011.

James M Robins, Mariela Sued, Quanhong Lei-Gomez, and Andrea Rotnitzky. Comment: Performance of double-robust estimators when "inverse probability" weights are highly variable. *Statistical Science*, 22(4):544–559, 2007. doi:10.1214/07-STS227D.

James M Robins, Thomas S Richardson, and Ilya Shpitser. An interventionist approach to mediation analysis. In Hector Geffner, Rina Dechter, and Joseph Y Halpern, editors, *Probabilistic and Causal Inference: The Works of Judea Pearl*, pages 713–764. Association for Computing Machinery, 2022. doi:10.1145/3501714.3501754.

Paul R Rosenbaum. Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477):191–200, 2007. doi:10.1198/016214506000001112.

Paul R Rosenbaum and Donald B Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983. doi:10.1093/biomet/70.1.41.

Michael Rosenblum and Mark J van der Laan. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *The International Journal of Biostatistics*, 6(2), 2010. doi:10.2202/1557-4679.1238.

Andrea Rotnitzky, Quanhong Lei, Mariela Sued, and James M Robins. Improved double-robust estimation in missing data and causal inference models. *Biometrika*, 99(2):439–456, 2012. doi:10.1093/biomet/ass013.

Andrea Rotnitzky, James M Robins, and Lucia Babino. On the multiply robust estimation of the mean of the g-functional. *arXiv preprint arXiv:1705.08582*, 2017. doi:10.48550/arXiv.1705.08582.

Daniel B Rubin and Mark J van der Laan. A doubly robust censoring unbiased transformation. *The International Journal of Biostatistics*, 3(1), 2007. doi:10.2202/1557-4679.1052.

Daniel B Rubin and Mark J van der Laan. Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis. *The International Journal of Biostatistics*, 4(1), 2008. doi:10.2202/1557-4679.1084.

Daniel B Rubin and Mark J van der Laan. Targeted ANCOVA estimator in RCTs. In Mark J van der Laan and Sherri Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*, pages 201–215. Springer, 2011. doi:10.1007/978-1-4419-9782-1_12.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974. doi:10.1037/h0037350.

Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005. doi:10.1198/016214504000001880.

Donald B Rubin. Causal inference through potential outcomes and principal stratification: Application to studies with "censoring" due to death. *Statistical Science*, pages 299–309, 2006. doi:10.1214/088342306000000114.

Fredrik Sävje, Peter Aronow, and Michael Hudgens. Average treatment effects in the presence of unknown interference. *The Annals of Statistics*, 49(2):673–701, 2021. doi:10.1214/20-aos1973.

Mireille E Schnitzer, Erica E M Moodie, Mark J van der Laan, Robert W Platt, and Marina B Klein. Modeling the impact of hepatitis C viral clearance on end-stage liver disease in an HIV co-infected cohort with targeted maximum likelihood estimation. *Biometrics*, 70(1):144–152, 2014a. doi:10.1111/biom.12105.

Mireille E Schnitzer, Mark J van der Laan, Erica E M Moodie, and Robert W Platt. Effect of breastfeeding on gastrointestinal infection in infants: A targeted maximum likelihood approach for clustered longitudinal data. *The Annals of Applied Statistics*, 8(2):703, 2014b. doi:10.1214/14-aoas727.

Mireille E Schnitzer, Judith J Lok, and Susan Gruber. Variable selection for confounder control, flexible modeling and collaborative targeted minimum loss-based estimation in causal inference. *The International Journal of Biostatistics*, 12(1):97–115, 2016. doi:10.1515/ijb-2015-0017.

Bryan E Shepherd, Peter B Gilbert, Yannis Jemiai, and Andrea Rotnitzky. Sensitivity analyses comparing outcomes only existing in a subset selected post-randomization, conditional on covariates, with application to HIV vaccine trials. *Biometrics*, 62(2):332–342, 2006. doi:10.1111/j.1541-0420.2005.00495.x.

Dylan S Small, Zhiqiang Tan, Roland R Ramsahai, Scott A Lorch, and M Alan Brookhart. Instrumental variable estimation with a stochastic monotonicity assumption. *Statistical Science*, 32(4): 561–579, 2017. doi:10.1214/17-STS623.

Richard J C M Starmans. The predicament of truth: On statistics, causality, physics, and the philosophy of science. In Mark J van der Laan and Sherri Rose, editors, *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*, pages 561–584. Springer, 2018. doi:10.1007/978-3-319-65304-4_30.

Eric J Tchetgen Tchetgen. Identification and estimation of survivor average causal effects. *Statistics in Medicine*, 33(21):3601–3628, 2014. doi:10.1002/sim.6181.

Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75, 2012. doi:10.1177/0962280210386779.

Anastasios Tsiatis. *Semiparametric Theory and Missing Data*. Springer, 2007. doi:10.1007/0-387-37345-4.

Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Xiaomin Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in Medicine*, 27(23):4658–4677, 2008. doi:10.1002/sim.3113.

Anastasios A Tsiatis, Marie Davidian, and Weihua Cao. Improved doubly robust estimation when data are monotonely coarsened, with application to longitudinal studies with dropout. *Biometrics*, 67(2):536–545, 2011. doi:10.1111/j.1541-0420.2010.01476.x.

Lars van der Laan, Wenbo Zhang, and Peter B Gilbert. Nonparametric estimation of the causal effect of a stochastic threshold-based intervention. *Biometrics*, 79(2):1014–1028, 2023. doi:10.1111/biom.13690.

Lars van der Laan, Marco Carone, and Alex Luedtke. Combining T-learning and DR-learning: A framework for oracle-efficient estimation of causal contrasts. *arXiv*, arXiv:2402.01972, 2024.

Mark J van der Laan. Causal inference for a population of causally connected units. *Journal of Causal Inference*, 2(1):13–74, 2014. doi:10.1515/jci-2013-0002.

Mark J van der Laan and Susan Gruber. Targeted minimum loss-based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1), 2012. doi:10.1515/1557-4679.1370.

Mark J van der Laan and Maya L Petersen. Causal effect models for realistic individualized treatment and intention-to-treat rules. *The International Journal of Biostatistics*, 3(1), 2007. doi:10.2202/1557-4679.1022.

Mark J van der Laan and Maya L Petersen. Direct effect models. *The International Journal of Biostatistics*, 4(1), 2008. doi:10.2202/1557-4679.1064.

Mark J van der Laan and James M Robins. *Unified Methods for Censored Longitudinal Data and Causality*. Springer, 2003. doi:10.1007/978-0-387-21700-0.

Mark J van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, 2011. doi:10.1007/978-1-4419-9782-1.

Mark J van der Laan and Sherri Rose. *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer, 2018. doi:10.1007/978-3-319-65304-4.

Mark J van der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006. doi:10.2202/1557-4679.1043.

Mark J van der Laan, Sandrine Dudoit, and Sunduz Keles. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology*, 3(1), 2004. doi:10.2202/1544-6115.1036.

Mark J van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1), 2007. doi:10.2202/1544-6115.1309.

Aad W van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998. doi:10.1017/CBO9780511802256.

Tyler VanderWeele and Stijn Vansteelandt. Mediation analysis with multiple mediators. *Epidemiologic Methods*, 2(1):95–115, 2014. doi:10.1515/em-2012-0010.

Tyler J VanderWeele. Simple relations between principal stratification and direct and indirect effects. *Statistics & Probability Letters*, 78(17):2957–2962, 2008. doi:10.1016/j.spl.2008.05.029.

Tyler J VanderWeele. Principal stratification–uses and limitations. *The International Journal of Biostatistics*, 7(1):1–14, 2011. doi:10.2202/1557-4679.1329.

Tyler J VanderWeele. *Explanation in Causal Inference: Methods for Mediation and Interaction*. Oxford University Press, 2015.

Tyler J VanderWeele and Peng Ding. Sensitivity analysis in observational research: Introducing the E-value. *Annals of Internal Medicine*, 167(4):268–274, 2017. doi:10.7326/M16-2607.

Tyler J VanderWeele and Stijn Vansteelandt. Conceptual issues concerning mediation, interventions and composition. *Statistics and its Interface*, 2(4):457–468, 2009. doi:10.4310/SII.2009.V2.N4.A7.

Tyler J VanderWeele, Eric J Tchetgen Tchetgen, Marilyn Cornelis, and Peter Kraft. Methodological challenges in Mendelian randomization. *Epidemiology*, 25(3):427, 2014. doi:10.1097/EDE.0000000000000081.

Tyler J VanderWeele, Peng Ding, and Maya Mathur. Technical considerations in the use of the E-value. *Journal of Causal Inference*, 7(2), 2019a. doi:10.1515/jci-2018-0007.

Tyler J VanderWeele, Alex R Luedtke, Mark J van der Laan, and Ronald C Kessler. Selecting optimal subgroups for treatment using many covariates. *Epidemiology*, 30(3):334, 2019b. doi:10.1097/EDE.0000000000000991.

Stijn Vansteelandt. Statistical modelling in the age of data science. *Observational Studies*, 7(1): 217–228, 2021. doi:10.1353/obs.2021.0013.

Stijn Vansteelandt, Maarten Bekaert, and Gerda Claeskens. On model selection and model misspecification in causal inference. *Statistical Methods in Medical Research*, 21(1):7–30, 2010. doi:10.1177/0962280210387717.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. doi:10.1080/01621459.2017.1319839.

Bingkai Wang, Elizabeth L Ogburn, and Michael Rosenblum. Analysis of covariance in randomized trials: More precision and valid confidence intervals, without model assumptions. *Biometrics*, 75 (4):1391–1400, 2019. doi:10.1111/biom.13062.

Ted Westling. Nonparametric tests of the causal null with nondiscrete exposures. *Journal of the American Statistical Association*, 117(539):1551–1562, 2022. doi:10.1080/01621459.2020.1865168.

Ted Westling, Peter Gilbert, and Marco Carone. Causal isotonic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(3):719, 2020. doi:10.1111/rssb.12372.

Richard Wyss, Sebastian Schneeweiss, Mark J van der Laan, Samuel D Lendle, Cheng Ju, and Jessica M Franklin. Using super learner prediction modeling to improve high-dimensional propensity score estimation. *Epidemiology*, 29(1):96–106, 2018. doi:10.1097/EDE.0000000000000762.

Jessica G Young, Miguel A Hernán, and James M Robins. Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. *Epidemiologic Methods*, 3(1):1–19, 2014. doi:10.1515/em-2012-0001.

Baqun Zhang, Anastasios A Tsiatis, Marie Davidian, Min Zhang, and Eric Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1:103–114, 2012. doi:10.1002/sta4.11.

Qingyuan Zhao, Dylan S Small, and Bhaswar B Bhattacharya. Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(4):735–761, 2019. doi:10.1111/rssb.12327.

Yingqi Zhao, Donglin Zeng, A John Rush, and Michael R Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012. doi:10.1080/01621459.2012.695674.

Wenjing Zheng and Mark J van der Laan. Cross-validated targeted minimum-loss-based estimation. In Mark J van der Laan and Sherri Rose, editors, *Targeted Learning: Causal Inference for Observational and Experimental Data*, pages 459–474. Springer, 2011. doi:10.1007/978-1-4419-9782-1_27.

Wenjing Zheng and Mark J van der Laan. Targeted maximum likelihood estimation of natural direct effects. *The International Journal of Biostatistics*, 8(1):1–40, 2012. doi:10.2202/1557-4679.1361.

Paul N Zivich, Michael G Hudgens, Maurice A Brookhart, James Moody, David J Weber, and Allison E Aiello. Targeted maximum likelihood estimation of causal effects with interference: A simulation study. *Statistics in Medicine*, 41(23):4554–4577, 2022. doi:10.1002/sim.9525.