

Annual Review of Sociology

Recent Developments in Causal Inference and Machine Learning

Jennie E. Brand,¹ Xiang Zhou,² and Yu Xie³

¹Department of Sociology, Department of Statistics, California Center for Population Research, and Center for Social Statistics, University of California, Los Angeles, California, USA;
email: brand@soc.ucla.edu

²Department of Sociology, Harvard University, Cambridge, Massachusetts, USA;
email: xiang_zhou@fas.harvard.edu

³Department of Sociology, Princeton University, Princeton, New Jersey, USA;
email: yuxie@princeton.edu

Annu. Rev. Sociol. 2023. 49:81–110

First published as a Review in Advance on
April 26, 2023

The *Annual Review of Sociology* is online at
soc.annualreviews.org

<https://doi.org/10.1146/annurev-soc-030420-015345>

Copyright © 2023 by the author(s). This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited. See credit lines of images or other third-party material in this article for license information.

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

causal inference, counterfactuals, machine learning, treatment effect heterogeneity, mediation, extrapolation, external validity

Abstract

This article reviews recent advances in causal inference relevant to sociology. We focus on a selective subset of contributions aligning with four broad topics: causal effect identification and estimation in general, causal effect heterogeneity, causal effect mediation, and temporal and spatial interference. We describe how machine learning, as an estimation strategy, can be effectively combined with causal inference, which has been traditionally concerned with identification. The incorporation of machine learning in causal inference enables researchers to better address potential biases in estimating causal effects and uncover heterogeneous causal effects. Uncovering sources of effect heterogeneity is key for generalizing to populations beyond those under study. While sociology has long emphasized the importance of causal mechanisms, historical and life-cycle variation, and social contexts involving network interactions, recent conceptual and computational advances facilitate more principled estimation of causal effects under these settings. We encourage sociologists to incorporate these insights into their empirical research.

1. INTRODUCTION

Many important questions in the social sciences, and everyday life, are causal questions. For example, we want to know how parental divorce affects children, how attending college affects job prospects, or how moving to a new neighborhood affects children's academic performance. We ask what would happen if individuals did or did not experience an event, like divorcing or attending college. Since reviews in sociology by Winship & Morgan (1999) and Gangl (2010), the literature on causal inference has developed several new promising directions. Some of the most exciting areas of development lie at the intersection of causal inference with machine learning (Athey & Imbens 2017, 2019; Huber 2023). This review describes several key identification strategies for causal inference and how machine learning methods can enhance our estimation of causal effects. Throughout our review, we describe some empirical applications of these methods in sociology.¹

We emphasize four main principles in our review. First, the plausibility of the assumptions underlying different research designs and identification strategies varies by application. Machine learning methods adapted to causal tasks facilitate estimation, but like other estimation tools, they do not assure the identification of causal effects. Second, causal effect heterogeneity is the norm, and it complicates extrapolation. Researchers may exert considerable effort in establishing a model with high internal validity, or credibility of the causal effect estimator, but with low external validity, or limited generalizability of the causal effect to other populations. We need to assess causal effect heterogeneity to understand the population distribution of causal effects. Machine learning methods can help identify subpopulations most responsive to treatments. Third, when evaluating social mechanisms in sociological research, we need to attend to confounding along the causal pathway, i.e., for the treatment–outcome relationship and the treatment–mediator and mediator–outcome relationships. Fourth, temporal and spatial interference, typical in social settings, complicate the definition, identification, and estimation of causal effects. These complications should be addressed more routinely in sociological research. In the following sections, we discuss (a) effect identification and estimation, (b) effect heterogeneity, (c) effect mediation, and (d) temporal and spatial interference. We conclude with some general remarks.

2. CAUSAL EFFECT IDENTIFICATION AND ESTIMATION

2.1. Notation and Estimands

Empirical work can be descriptive, such that we establish facts through associations between observables. For example, we might observe that college graduates earn higher wages than non-college graduates. But to evaluate causal effects, we draw on counterfactuals, i.e., we ask how much college-educated individuals would have earned without a college degree. The potential-outcomes framework offers a conceptual apparatus for defining causal effects. The framework has roots in research on experiments by Fisher (1935) and Neyman (1923) and research in economics by Roy (1951) and Quandt (1972). Rubin formalized and extended the potential-outcomes framework in a series of papers in statistics in the 1970s and 1980s (e.g., Rubin 1974, 1977, 1986).

Let us define a treatment W , e.g., an event or intervention, applied to unit i , i being a member of a population. A unit exposed to a treatment ($W_i = 1$) at a specific time could have instead been exposed to an alternative treatment (i.e., control, $W_i = 0$) at the same time. For example, a

¹Our review differs from recent reviews in sociology (Lundberg et al. 2022, Molina & Garip 2019) and political science (Grimmer et al. 2021) on machine learning in that we focus on the intersection between causal inference and machine learning. Readers are also directed to Hastie et al. (2017) for a textbook treatment of statistical learning.

person who attended college could have instead not attended college. We assume units assigned to treatment and control groups have potential outcomes in both states, the ones in which they are observed and unobserved. For a binary treatment, let Y be an outcome of interest and Y_{i1} and Y_{i0} the potential outcomes for unit i that would result from exposure to the treatment and control states, respectively. The causal effect of the treatment is thus the difference between the potential outcomes (i.e., $Y_{i1} - Y_{i0}$). The fundamental problem of causal inference is that we cannot observe both potential outcomes (Holland 1986). This framework is often applied to binary treatments, although extending to multicategory treatments is conceptually straightforward. We may also consider continuous treatments, but in this case, the number of potential outcomes becomes infinite, rendering the framework more complex (Gill & Robins 2001).² For each unit, we assume that the treatment status and potential outcomes determine the observed outcome. Let us focus on the case of binary treatment conditions. We have $Y_i = W_i Y_{i1} + (1 - W_i) Y_{i0}$. The stable unit treatment value assumption (SUTVA) (Rubin 1986) implies that the potential outcomes for any unit do not vary with the treatment assigned to other units. In other words, there is no interference between units. However, in many social settings, SUTVA can be problematic. For example, the wages for one college graduate may be affected by the population proportion of workers completing college.

Following Heckman & Robb (1986), we assume that treatment effects are heterogeneous. Using the potential-outcomes notation, we smooth out that heterogeneity and define different estimands for specific populations of interest.³ The average treatment effect (ATE) is the average of individual treatment effects in the population:

$$\tau_{\text{ATE}} = E[Y_1 - Y_0], \quad 1.$$

where we omit the unit subscript i for conciseness. The average treatment effect on the treated (ATT) is the average of individual effects among the treated subpopulation:

$$\tau_{\text{ATT}} = E[Y_1 - Y_0 | W = 1]. \quad 2.$$

Now consider the estimand corresponding to the difference in average outcomes between the treated and control units:

$$\tau = E[Y | W = 1] - E[Y | W = 0]. \quad 3.$$

Following Abadie & Cattaneo (2018), we note that

$$\tau = \tau_{\text{ATE}} + b_{\text{ATE}} = \tau_{\text{ATT}} + b_{\text{ATT}}, \quad 4.$$

where b_{ATE} and b_{ATT} are bias terms given by

$$b_{\text{ATE}} = (E[Y_1 | W = 1] - E[Y_1 | W = 0]) \Pr(W = 0) + (E[Y_0 | W = 1] - E[Y_0 | W = 0]) \Pr(W = 1) \quad 5.$$

and

$$b_{\text{ATT}} = E[Y_0 | W = 1] - E[Y_0 | W = 0]. \quad 6.$$

If the average potential outcomes under both states are identical between treated and control units, the bias terms b_{ATE} and b_{ATT} disappear. This condition is, however, untestable. Confounding arises when pretreatment characteristics correlated with potential outcomes also influence treatment assignment.

²Kennedy et al. (2017) develop methods for estimation of continuous treatment effects.

³Lundberg et al. (2021) and Lundberg (2022) discuss setting theoretical estimands in precise terms, outside of any statistical model.

2.2. Experimental Studies

Randomized experiments, where we randomly assign individuals to treatment and control conditions, offer one strategy to address confounding (Fisher 1935). With successful randomization, experiments generate independence between treatment status and both potential outcomes:

$$(Y_1, Y_0) \perp W, \quad 7.$$

where $\perp$ denotes statistical independence. Consequently, the bias terms in Equation 5 and Equation 6 equal 0, and we can credibly attribute the difference in average outcomes between the treated ($E[Y|W = 1]$) and control ($E[Y|W = 0]$) groups to the treatment. In a traditional experiment, we assign a predetermined number of units to one of two conditions. However, unless researchers conduct an experiment on a population-representative sample, it is not generally possible to derive the population-level ATE from experimental data. We return to the topic of extrapolating study-specific results in Section 3.2.

Recent developments in randomized experiments include adaptive designs for evaluating optimal treatment assignment. For example, multi-armed bandits tailor treatments to individuals when they need treatment. The design aims to balance the goals of exploration (i.e., evaluating the effects of different treatment conditions) and exploitation (i.e., assigning units to treatment conditions with higher payoffs) (Athey & Imbens 2019, Carranza et al. 2022, Offer-Westort et al. 2021, Scott 2010). For example, consider an online setting where treatment is assigned sequentially to different units, and the outcome for each unit is measured quickly after treatment assignment. A multi-armed bandit assigns treatment conditions based on information learned up to the point of the assignment, thus allowing researchers or policy makers to assign more units to conditions with higher payoffs. Sociological applications of multi-armed bandits remain scarce, but it is a promising approach for future studies.

2.3. Observational Studies Under Unconfoundedness

For practical and ethical reasons, sociologists cannot address many interesting social questions using experiments. Some scholars also debate the extent to which randomized experiments should dominate the hierarchy of scientific evidence (Abadie & Cattaneo 2018, Deaton & Cartwright 2018). But in most observational studies, the independence condition (Equation 7) may not hold. In the case of college effects, for example, the simple difference between college graduates' and non-college graduates' wages is not a credible estimate of the causal effect due to pretreatment heterogeneity, i.e., that individuals with higher skills and achievement and advantaged social backgrounds disproportionately complete college. We may observe some confounding factors in our data, while others are unobserved. Researchers may assume that after adjusting for a set of pretreatment covariates X , no additional confounders affect both treatment status and the outcome. That is, they assume unconfoundedness (also called ignorability, selection-on-observables, conditional independence, or exogeneity):

$$(Y_1, Y_0) \perp W | X, \quad 8.$$

which allows for the identification of the causal effect of W on Y by adjusting for X . **Figure 1** is a directed acyclic graph (DAG) representing the causal relationships between W , X , and Y under unconfoundedness.⁴

⁴DAGs represent assumptions about nonparametric causal relationships between variables (in contrast to path diagrams that reflect linear structural equations) [see Pearl (2009) and Morgan & Winship (2014) for




Figure 1

Direct acyclic graph under unconfoundedness. W denotes treatment status, Y denotes the outcome of interest, and X denotes observed pretreatment confounders.

To estimate the ATE, we also assume positivity, meaning that treatment assignment is probabilistic at all covariate values in the population. Positivity is a strong assumption as it rules out the possibility that treatment status has no variation (i.e., at either 0 or 1) at some covariate values. The latter might happen by chance, even if positivity holds in the population. For example, while there may be young adults from families in the highest income decile who did not attend college, such youth may fail to appear in a particular sample. Moreover, near violations of positivity (e.g., very few treated/untreated units at some covariate values) can result in unstable estimates of causal effects for subgroups of the population. In practice, we often trim observations with very high and low estimated treatment probabilities to reduce instability in our estimated effects, or those outside the region of common support, leading to effect estimates that do not fully represent the population. In restricting to the region of common support, we sacrifice a degree of external validity to enhance internal validity.

Under the assumptions of unconfoundedness and positivity, researchers draw on various methods to estimate causal effects, such as regression imputation, propensity score matching (PSM), and inverse probability weighting (IPW) [see Imbens (2004) or Gangl (2010) for a review of these methods]. Using regression imputation, the researcher fits a regression model for the conditional mean of the outcome Y given treatment status W and pretreatment covariates X , $\mu_W(X) = E[Y|W, X]$; imputes the potential outcomes under treatment and control for each unit, $\hat{\mu}_1(X_i) = \hat{E}[Y|W = 1, X_i]$ and $\hat{\mu}_0(X_i) = \hat{E}[Y|W = 0, X_i]$; and estimates the ATE using the average difference between these imputed outcomes:

$$\hat{\tau}_{\text{ATE}} = \frac{1}{n} \sum_{i=1}^n (\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)), \quad 9.$$

where n is the sample size. If the outcome model $\mu_W(X)$ is additive in W and X , $\hat{\tau}_{\text{ATE}}$ will reduce to the coefficient on W in the regression model.

Using PSM and IPW, the researcher fits a model for the propensity score, i.e., the conditional probability of treatment given the pretreatment covariates, $p(X) = \Pr[W = 1|X]$, and obtains the estimated propensity score for each unit, $\hat{p}(X_i)$ (Rosenbaum & Rubin 1983). With PSM, the researcher then matches treated and control units with similar propensity score values and uses their differences to estimate effects (Abadie & Imbens 2016, Caliendo & Kopeinig 2008, Imbens 2015). Matching algorithms differ primarily in how researchers define the distance between units (e.g., propensity scores), select the number of control units, select controls with or without replacement, and weight multiple control units (Austin & Stuart 2017, Morgan & Harding 2006, Morgan & Winship 2014). Decisions regarding how many controls to use and whether to match with or

background on the use of DAGs]. Edges are directed, such that an arrow indicates the effect of one variable on another (e.g., the effect of W on Y). They are also acyclic, such that there are no feedback loops.

without replacement involve a bias-variance trade-off.⁵ With IPW, the researcher estimates the ATE using a weighted difference in means:

$$\hat{\tau}_{\text{ATE}} = \frac{1}{n} \sum_{i=1}^n \left(\frac{W_i Y_i}{\hat{p}(X_i)} - \frac{(1 - W_i) Y_i}{1 - \hat{p}(X_i)} \right). \quad 10.$$

By weighting each unit by the inverse estimated propensity [i.e., $1/\hat{p}(X_i)$ for treated units and $1/(1 - \hat{p}(X_i))$ for control units], researchers create a weighted sample in which treatment status is expected to be independent of all pretreatment covariates. In other words, if the propensity score model is correct, we expect that treated and control units are balanced in their covariate values.

Regression imputation, PSM, and IPW involve modeling different parts of the data distribution. While regression imputation depends on a correctly specified outcome model, PSM and IPW depend on a correctly specified propensity score model. However, correctly specifying either model is difficult, especially when the vector of pretreatment covariates, X , is high-dimensional. When the outcome or propensity score model is misspecified, the corresponding regression imputation, matching, or IPW estimates can be biased. Misspecification may arise either (a) because we have many potential (observed) confounders of the treatment–outcome relationship in the data or (b) because the researcher is agnostic about the functional form in which treatment status and the covariates affect the outcome. Both scenarios are common in sociological research. Given the second scenario, researchers may experiment with higher-order and interaction terms. Imbens & Rubin (2015) propose an iterative approach to produce a flexible specification of the propensity score specification. Scholars have also advocated using flexible machine learning methods to fit the outcome or propensity score models. For example, researchers have used classification and regression trees, random forests, and ensemble methods to estimate propensity scores (e.g., Brand et al. 2021, Lee et al. 2010, McCaffrey et al. 2004, Westreich et al. 2010).⁶ Scholars should draw on theory in selecting covariates to include (Cinelli et al. 2022, Elwert 2015, Elwert & Winship 2014, Pearl 2009).

In each scenario, however, we may face complications because these methods were generally not designed for causal inference. Supervised machine learning methods are designed to minimize prediction errors rather than estimate causal effects. For example, a least absolute shrinkage and selection operator (LASSO) regression for the outcome tends to select a subset of the covariates highly predictive of the outcome. Such a subset, however, may not be the optimal subset for estimating the ATE. Furthermore, if we omit covariates highly predictive of treatment status, even if their correlations with the outcome are modest, substantial bias may arise in our treatment effect estimates (Belloni et al. 2014). Similarly, suppose we use an off-the-shelf machine learning method to fit the propensity score model for matching or IPW. In that case, it will seek a model that minimizes the error of predicting treatment status, which may not be the model that yields the optimal propensity score estimates for balancing covariates between the treated and control units.

Researchers have adapted machine learning methods to estimate causal parameters to mitigate these and other concerns central to causal inference. First, to adapt machine learning to the

⁵ More control units lead to greater efficiency and greater bias, while fewer control units lead to less efficiency and less bias. Allowing replacement increases the average quality of the matches but reduces the number of unique control units used to estimate the counterfactual mean, increasing the estimator's variance [see An & Winship (2017), Imbens (2015), and Imbens & Rubin (2015) for discussion of matching procedures].

⁶ An (2010) describes Bayesian propensity score estimators that model the joint likelihood of both propensity scores and outcomes in one step to incorporate the uncertainty in propensity score estimation. Simulations show that this approach corrects for overly conservative inference based on standard propensity score estimators.

regression-imputation approach, Belloni et al. (2014) propose a double selection procedure in which we fit two LASSO regressions, one for the outcome and one for treatment status. After that, we fit an ordinary least squares regression of the outcome on treatment status and the union of the selected covariates in the first two LASSO regressions. In doing so, researchers adjust for covariates important in predicting either the outcome or treatment status, avoiding the bias resulting from a single LASSO regression of the outcome. Künzel et al. (2019) propose a metalearner that can use any supervised learning algorithm to estimate ATEs. They show that the X-learner performs favorably using random forest and Bayesian additive regression trees (BART) as base learners.

Second, to adapt machine learning for IPW, McCaffrey et al. (2004) proposed fitting the propensity score model using gradient boosting machines. This approach is a precursor to the literature on calibrated propensity scores (e.g., Imai & Ratkovic 2014) and balancing weights (e.g., Athey et al. 2018b, Fong et al. 2018, Hainmueller 2012, Zhou & Wodtke 2020, Zubizarreta 2015). Using optimization methods, researchers choose a set of weights such that in the weighted sample, the treated and control units are either exactly or approximately balanced in pretreatment covariates (by a prespecified balancing metric). This procedure ensures that bias due to covariate imbalance is slight. Zhou (2019), for example, adapts this approach to assess the effect of college completion on intergenerational income mobility.

Finally, machine learning methods are particularly attractive when combined with the so-called doubly robust estimators of ATEs (Robins & Rotnitzky 1995, Robins et al. 1994). Consider the following doubly robust estimator of the ATE:

$$\hat{\tau}_{\text{ATE}} = \frac{1}{n} \sum_{i=1}^n \left(\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i) + \frac{W_i(Y_i - \hat{\mu}_1(X_i))}{\hat{p}(X_i)} - \frac{(1 - W_i)(Y_i - \hat{\mu}_0(X_i))}{1 - \hat{p}(X_i)} \right). \quad 11.$$

Under the assumptions of SUTVA, unconfoundedness, and positivity, it is consistent for the ATE if either the outcome model $\mu_W(X)$ or the propensity score model $p(X)$, but not necessarily both, is correctly specified (Scharfstein et al. 1999). The double-robustness property occurs because the bias of Equation 11 as an estimator of the ATE is governed by the product of two bias terms: (a) the bias of the fitted outcome model $\hat{\mu}_W(X)$ and (b) the bias of the fitted propensity score model $\hat{p}(X)$. Provided one of the two biases converges to zero, the bias of Equation 11 will converge to zero. This property motivates what Chernozhukov et al. (2018) call debiased machine learning (DML) of the ATE, i.e., the use of flexible machine learning methods to construct estimates of $\mu_W(X)$ and $p(X)$ in Equation 11 (see also van der Laan & Rubin 2006). Due to the data-driven nature of machine learning methods, they generally do not provide root- n consistent estimates of the $\mu_W(X)$ and $p(X)$ functions.⁷ However, because of the multiplicative structure of its bias expression, Equation 11 remains a root- n consistent estimator of the ATE under mild conditions.⁸ By contrast, the biases of the regression imputation and IPW estimators (Equations 9 and 10) do not have such a multiplicative structure, preventing root- n consistent estimation of the ATE when researchers use machine learning to estimate the outcome or the propensity score model. In a recent application of this approach, Zhou & Pan (2023) employed a DML approach to assess the heterogeneous effects of college attendance and BA degree completion on earnings for Black and White Americans.

⁷Root- n consistency means that the estimator converges on the true value at a rate of $n^{-1/2}$.

⁸This is true if the product of the convergence rates of the machine learning estimators of $\mu_W(X)$ and $p(X)$ is faster than $n^{-1/2}$. We can achieve this property when, for example, both converge to the truth at a faster-than- $n^{-1/4}$ rate, which is attainable for many machine learning methods.

When researchers use DML, it is advisable to use sample splitting, whereby, for example, we use a portion of the data as a training sample to estimate the outcome and propensity score models and another portion to evaluate Equation 11. This procedure removes the overfitting bias of machine learning estimators of the outcome and propensity score models.⁹ However, a conventional sample splitting procedure would involve a waste of data. To retain efficiency, researchers may draw on cross-fitting, which includes the following steps (Chernozhukov et al. 2018): (a) randomly partition the sample into J folds, $S_1, S_2, \dots, S_J$, where J is a small number such as five; (b) for each j , obtain a fold-specific estimate of the ATE using only data from S_j , but with the outcome and propensity score models estimated from the remainder of the sample ($S \setminus S_j$); and (c) average these fold-specific estimates to form a final estimate of the target parameter.

Finally, researchers should routinely consider how the results obtained under the unconfoundedness assumption would change if we relaxed that assumption. One common approach is to conduct sensitivity analyses by subtracting a bias term from the point estimate and confidence interval of the estimated treatment effects (Gangl 2015, VanderWeele & Arah 2011). The bias term is equal to the product of two parameters:

$$B = \gamma\lambda, \quad 12.$$

where

$$\gamma = E(Y|U = 1, W, X) - E(Y|U = 0, W, X) \quad 13.$$

and

$$\lambda = \Pr(U = 1|W = 1, X) - \Pr(U = 1|W = 0, X). \quad 14.$$

That is, γ is the mean difference in the outcome associated with a unit change in an unobserved binary confounder, U , and λ is the mean difference in the unobserved confounder between treated and control units. Cinelli & Hazlett (2020) provide additional measures and graphical tools for assessing sensitivity to unobserved confounding.

2.4. Quasi-Experimental Designs

In settings where researchers deem the unconfoundedness assumption (Equation 8) implausible, they may seek to identify causal effects using quasi-experimental designs, such as instrumental variables (IVs) or regression discontinuity (RD). IVs are widely used in randomized experiments with imperfect compliance and in natural experiments using observational data (Angrist et al. 1996, Imbens & Angrist 1994). As an example of the latter, several studies have used proximity to a local college as an IV for college attendance to assess the effects of attendance on wages (Card 2001, Deaton 2010). **Figure 2** is a DAG representation of the IV design, where an unobserved confounder U may affect both the treatment W and the outcome Y . The IV Z affects W and can affect Y only indirectly through its effect on W . Researchers use exogenous variation in Z , which induces changes in W , to identify the causal effect of W on Y . An IV analysis is typically implemented using two-stage least squares (2SLS). In the first stage, a linear model is used to predict treatment status given the IV and a set of pretreatment covariates X . In the second stage,

⁹Researchers using machine learning methods often attend to the issue of overfitting more than those using conventional statistical models (Athey & Imbens 2019). The goal is to select flexible models that fit well, but not so well that out-of-sample prediction is compromised. Regularization techniques calibrate machine learning methods to minimize a loss function and avoid overfitting.




Figure 2

Direct acyclic graph under the instrumental variable design. W denotes treatment status, Y denotes the outcome of interest, Z denotes an instrumental variable, and X denotes observed pretreatment confounders.

the outcome Y is regressed on X and the fitted values of W from the first stage, whose coefficient represents the causal effect of W on Y .¹⁰

The IV approach allows for unobserved confounding of the W – Y relationship but relies on other stringent assumptions. First, conditional on the pretreatment covariates X , the instrument must be exogenous. That is, no unobserved confounding exists for the Z – W and Z – Y relationships (i.e., the independence assumption). Second, we assume that the IV affects the likelihood of treatment, even if it does so within a small range (i.e., the relevance assumption). Third, we assume that the IV affects the outcome only indirectly through the treatment (i.e., the exclusion restriction). Finally, allowing for heterogeneous treatment effects, we assume that although the instrument may not affect some people, all those affected are affected in the same direction (i.e., the monotonicity assumption). With these assumptions in place, researchers have suggested that the 2SLS identifies the local average treatment effect (LATE) for a binary treatment W (Angrist & Pischke 2009):

$$\tau_{\text{LATE}} = E[Y_1 - Y_0 | W_1 > W_0], \quad 15.$$

where W_0 and W_1 denote the potential treatment value when the IV Z takes the values of 0 and 1, respectively. In actual social settings, the inducement effect of an IV is often small. Low inducement can be a major limitation in IV analysis because it can subject the causal effect estimate to large variance, substantial finite-sample bias, and high sensitivity to violations of the exclusion restriction (Bound et al. 1995). Felton & Stewart (2022) contend that while sociologists have increasingly adopted IV as a strategy, assumptions underlying the model often go unstated, and robust uncertainty measures are rarely used. Moreover, the 2SLS approach relies on correct specification of the treatment and outcome models, which can be difficult to justify when the pretreatment covariates X are high-dimensional. Blandhol et al. (2022) show that a saturated specification for 2SLS that correctly specifies the relationship between the instruments and the covariates (including interactions) is necessary for researchers to interpret the estimator as an average of covariate-specific LATEs. Chernozhukov et al. (2018) outline a DML approach for estimating the LATE, which involves fitting three models: (a) a model for $E[Z|X]$, (b) a model for $E[W|Z, X]$, and (c) a model for $E[Y|Z, X]$. In contrast to 2SLS, the DML approach allows all these models to be fit using flexible machine learning methods, thus reducing model dependency. This approach offers a more principled method for estimating the LATE.

¹⁰Readers are directed to Steiner et al. (2017) for a discussion of graphical models for quasi-experimental designs.

In an RD design, we determine access to treatment by a cutoff value c on a continuous running variable X (see Cattaneo et al. 2019, Cattaneo & Titiunik 2022). The RD design assumes that the average response of units just below the cutoff provides a good approximation to the average response that we would have observed for units just above the cutoff had they not been assigned to treatment. Under this assumption, a comparison between units just below and above the cutoff mimics a randomized experiment and reveals a local treatment effect, i.e.,

$$\tau_{RD} = E[Y_1 - Y_0 | X = c]. \quad 16.$$

To estimate this quantity, we fit two local linear regressions, one for units below the cutoff and one for units above the cutoff. We use their difference in the predicted outcome at $X = c$ as an estimate of τ_{RD} (Imbens & Lemieux 2008). For example, suppose students were admitted to college based on a minimum score on an admission test. Students just above the minimum score are arguably comparable to those just below the minimum score in terms of other characteristics that predict college-going. Around the test score cutoff, we can compare the outcomes of those who are and are not admitted. Yet, we can imagine situations where not everyone admitted to college would choose to attend, in which case we would have a fuzzy rather than a sharp RD design. This situation allows the cutoff to change treatment status for some, yet not all, units, i.e., compliers. A fuzzy RD design allows the researcher to identify a local treatment effect among compliers (Hahn et al. 2001), i.e.,

$$\tau_{RD(F)} = E[Y_1 - Y_0 | X = c, W_1 > W_0], \quad 17.$$

where W_0 and W_1 denote the potential value of treatment when the running variable X is just below and above the cutoff, respectively. We can estimate Equation 17 using a combination of 2SLS and local linear regressions (Imbens & Lemieux 2008).¹¹ Researchers should assess the validity of the underlying assumptions using supplementary analyses to test for evidence of the manipulation of the cutoff variable and for discontinuities in average covariate values at the threshold. RD methods can have high internal validity for an observational study, but low external validity. Several approaches have been proposed to enable valid extrapolation (Cattaneo & Titiunik 2022).

3. CAUSAL EFFECT HETEROGENEITY

As we note above, individuals differ not only in pretreatment characteristics (i.e., pretreatment heterogeneity) but also in how they respond to a common treatment (i.e., treatment effect heterogeneity). Analyses that estimate heterogeneous treatment effects can yield insights into how scarce social resources are distributed in an unequal society and how events differentially impact populations with different expectations of their occurrence (e.g., Brand 2023, Heckman et al. 2018). In some cases, we may hypothesize that an event has significant consequences for some subgroups but less or no effect among others (e.g., Brand et al. 2019a). Scholars may aim to identify the most responsive subgroups to determine which individuals benefit most from treatment so that policy makers can better assign different treatments to balance competing objectives, such as reducing costs and maximizing outcomes for targeted groups (Athey & Imbens 2019; Manski & Garfinkel 1992; Zhou & Xie 2019, 2020). An important feature of the potential-outcomes framework is that it allows for general heterogeneity in treatment effects from the outset. Attending to treatment effect heterogeneity can also help extrapolate findings to diverse populations and contexts.

¹¹Researchers also need to consider bandwidth selection in RD designs (for a discussion, see Imbens & Lemieux 2008, Lee & Lemieux 2010).

3.1. Estimating Heterogeneous Causal Effects

Social scientists employ a variety of approaches to estimate heterogeneous effects. Researchers commonly partition their samples into subgroups defined by individual characteristics, like gender, race, or social class, to explore variation in treatment effects. Yet, for questions of causal inference, the association between the treatment effect and treatment propensity constitutes a key axis of heterogeneity (Brand & Simon Thomas 2013, Heckman et al. 2006, Xie 2013). One way to identify heterogeneity by selection into treatment is to compare different population parameters. For example, the ATE and ATT may differ. If $ATE > ATT$, those with a lower propensity of treatment have larger estimated treatment effects. If $ATT > ATE$, those with a higher propensity of treatment have larger estimated treatment effects. Or we might directly assess how treatment effects vary by the estimated propensity score (Brand & Simon Thomas 2013, Xie et al. 2012). For example, Cheng et al. (2021) use growth curve models to assess how the effects of college on long-term wages vary across strata of the estimated likelihood that individuals complete a degree. Alternatively, we can obtain matched differences between treated and control units, plot them along a continuous propensity score axis, and then use local polynomial smoothing to observe variation in effects by the likelihood of treatment. For example, Brand & Simon Thomas (2014) use this approach to explore how the effects of maternal job displacement on children's educational attainment vary by the likelihood that mothers lose a job. Economists also often compare IV (LATE) estimates with ordinary least squares estimates to assess differential response patterns. As we indicated above, with treatment effect heterogeneity, the LATE can differ from the ATE and the ATT.¹²

Researchers tend to base decisions of which subgroups to explore in analyses of effect heterogeneity on theoretical priors. For example, researchers may stratify by gender or race because they are interested in sociodemographic variation. In contrast to this approach, emerging machine learning methods allow researchers to explore sources of variation that they may not have previously considered (Lundberg et al. 2022, Shu & Ye 2023). For example, we can search for effect heterogeneity by adapting a variable selection algorithm such as LASSO, which automatically selects the more predictive interactions between the treatment and covariates (Imai & Ratkovic 2013). Social scientists have also employed tree-based methods to uncover differential responses to treatment. Decision trees, a widely used machine learning approach, recursively split data into increasingly smaller subsets where data bear greater similarity (Brand et al. 2020).¹³ Decision trees are attractive for social research because they are easily interpretable. Causal trees, i.e., decision trees adapted for causal inference, partition the data to minimize heterogeneity in within-leaf treatment effects (Athey & Imbens 2016, Brand et al. 2021).¹⁴ We split the data, construct a tree using a training sample, and estimate leaf-specific treatment effects using an estimation sample. This approach allows researchers to uncover subpopulations of interest that they had not prespecified with greater flexibility by searching over high-dimensional functions

¹²Bloome & Schrage (2021) describe an approach for estimating heterogeneous treatment effects using covariance regression models. They demonstrate the approach by analyzing the effects of sharing information about income inequality on redistributive preferences.

¹³At each decision, splits are chosen by selecting a covariate and threshold that minimize an in-sample loss function. This partitioning process is repeated until a regularization penalty selected through cross-validation limits the depth of the tree.

¹⁴Causal trees bear similarity to kernel regression or matching methods. We can think of the leaf as defining the set of nearest neighbors for a given target observation in a leaf, and the estimator from a single tree as a matching estimator with alternative ways of selecting the nearest neighbor to a treated unit (Athey & Imbens 2019).

of covariates.¹⁵ We can then use several methods described above, such as weighting, matching, or machine learning, to estimate leaf-specific effects in the presence of observed confounding. For example, Brand et al. (2021) estimate the effects of college completion on reducing low-wage work with a causal tree. They find that individuals who had the largest effects of college on reducing low-wage work are those with disadvantaged backgrounds and low psychosocial skills.

Single decision trees benefit from interpretability but can be unstable and do not allow causal effects to change more smoothly across covariates. A causal forest builds on the causal tree algorithm by averaging over many trees (Athey et al. 2019, Breiman 2001, Wager & Athey 2018).¹⁶ In principle, every individual has a distinct estimate. Using this strategy, researchers may consider effect heterogeneity by ranking estimated individual treatment effects and then considering the characteristics of groups in the highest- and lowest-ranked categories. Recent approaches also combine supervised learning of the response variable with supervised learning of the propensity score to estimate treatment effect heterogeneity. For example, Nie & Wager (2021) describe a general class of two-step algorithms for heterogeneous treatment effect estimation in observational studies.

Complementing these forest approaches is a DML approach proposed by Semenova & Chernozhukov (2021). Instead of detecting effect heterogeneity from many covariates, this approach allows researchers to directly estimate conditional average treatment effects given a few prespecified covariates. The method is helpful in applications where the researcher wants to see how the treatment effect differs by selected characteristics such as gender, race, or social class categories. For example, Zhou (2022a) adapts the DML approach to study group-based heterogeneity in the total effect of college attendance and its direct and indirect effects via degree completion (see also Zhou & Pan 2023). Closely related to treatment effect heterogeneity is an emerging literature on policy learning (Athey & Wager 2021). In this case, researchers learn in a data-driven way the optimal assignment of treatment to specific subgroups defined in terms of observed characteristics (e.g., parental income categories). Accordingly, policy makers can target those for whom the treatment effects are largest. Policy learning is beneficial in settings where we aim to optimize an outcome for a costly treatment (e.g., a social intervention with limited funds to cover treatment costs). Yadlowsky et al. (2021) propose a rank-weighted ATE to determine treatment prioritization rules based on responsiveness to treatment.

In studies considering treatment effect heterogeneity, researchers should consider how unobserved selection may contribute to heterogeneous response patterns. Localized sensitivity analyses should be routinely performed for analyses that involve effects stratified by unit characteristics, propensity scores, or machine learning-generated categories. Notably, Zhou & Xie (2019, 2020) also describe the relationship between heterogeneity by observed and unobserved selection into treatment and consider the policy implications under different scenarios as the treated population composition shifts.

3.2. Implications of Heterogeneous Causal Effects for Extrapolation

If effects were the same for everyone, it would be easy to generalize an effect estimate from a sample to a population. Effect heterogeneity complicates the generalizability of ATE estimates. Researchers should consider the population of interest when interpreting treatment effect estimates from heterogeneous subgroups. Social scientists who aim to minimize confounding may

¹⁵We also may include the propensity score as an input variable (e.g., Hahn et al. 2020).

¹⁶Building on the comparison to kernel regression or matching, we can think of a causal forest as an average of matching estimators.

draw on experimental or quasi-experimental methods. Yet as a researcher attempts to extrapolate or generalize from a specific group of subjects under study to a target population, the average effects may differ due to compositional differences (Hartman et al. 2015, Kern et al. 2016, Manski & Garfinkel 1992, Westreich et al. 2019). In other words, researchers often face a trade-off between internal and external validity. Internal validity is our degree of confidence that a causal relationship exists between the treatment and the outcome. External validity is our ability to generalize findings to other populations (Manski 1995, Manski & Garfinkel 1992). The literature on causal inference is primarily concerned with the internal validity of a causal relationship, which must be complemented by a focus on generalizable knowledge. Social science and public policy demand greater attention to external validity (Egami & Hartman 2022, Findley et al. 2021).

As described in Section 2, researchers use a variety of strategies to claim the internal validity of their effect estimates. For example, a randomized controlled trial may give us sample ATEs free from pretreatment heterogeneity bias. However, we may be limited in extrapolating and providing estimates of population ATEs (Hartman et al. 2015, Stuart et al. 2015, Xie 2013). Indeed, the population of units for which we credibly assess causal effects might be quite small. For example, some experimental effect estimates may only apply to treated units in the specific geographical setting in which researchers conducted the study, such as the 1962–1967 HighScope Perry Preschool Program conducted in Ypsilanti, Michigan (Xie et al. 2020).

Compositional differences may also arise in a dynamic setting where treatment gradually expands over successive segments (Xie 2013). In this case, units with higher treatment propensity are likely overrepresented when the population treated is small. As the treated population expands, the overrepresentation of high-propensity treated units declines. This compositional shift among newly recruited units, typically from higher- to lower-propensity units, will impact treatment effect estimates from an experimental study that targets those at the margin of being treated. In a social intervention on a graduated schedule where participation is need-based, the poorest individuals may be chosen first and benefit most from the intervention. Researchers may calculate an ATE for the subpopulation subject to the experiment. Yet under these conditions, individuals selected at later stages (i.e., becoming eligible only after the eligibility cut-point is moved up the income distribution) would exhibit lower ATEs. Low external validity is problematic for policy purposes, as policy makers require evidence of the effectiveness of interventions for target populations that may differ from those represented by experimental participants.

Similarly, suppose we adopt an IV design and have a valid IV. In that case, we may have a stronger basis for asserting internal validity than a standard regression approach without an IV, but only for a small segment of the subpopulation induced by the IV. Thus, we may be unable to extrapolate from those induced into treatment to a broader population. For example, we may use college proximity as an instrument for college attendance to assess the effects of college attendance on wages. If those induced into college have different effects than average college-goers, we cannot extrapolate the findings to the broader population (Mogstad & Torgovitsky 2018). With RD designs, researchers also need to consider under what conditions we can extrapolate estimated effects to populations further away from the threshold (Angrist & Rokkanen 2015, Bertanha & Imbens 2020, Dong & Lewbel 2015).¹⁷

Several approaches may help us move from the sample to the population ATE, such as bias-corrected matching (Hotz et al. 2005), propensity score weighting (Cole & Stuart 2010, Stuart et al. 2011), propensity score subclassification (Tipton 2013, Tipton et al. 2014), entropy

¹⁷Regression model estimates from representative samples of the population also face external validity problems, as the units in the sample contribute to the causal effects to differing extents (see Aronow & Samii 2016).

weighting (Hartman et al. 2015), machine-learning-based estimation of heterogeneous treatment effects (e.g., Kern et al. 2016), and calibration methods to generate balancing weights (Josey et al. 2022). A propensity score approach, for example, models membership in the population versus the experimental sample. Then, the propensity scores are used to make the sample subjects resemble the target population (e.g., Xie et al. 2020). This approach can work well when the covariates strongly predict membership in the target population and treatment effect heterogeneity (Pearl & Bareinboim 2014). Machine learning methods can automate the detection of treatment-by-covariate interactions. Kern et al. (2016) show that BART performs reasonably well for extrapolating from the sample to a target population when observed covariates are sufficient for accounting for treatment effect heterogeneity.

4. CAUSAL EFFECT MEDIATION

While traditional sociological approaches to mediation analysis relied on parametric structural equation models to define and estimate direct and indirect effects (e.g., Alwin & Hauser 1975, Baron & Kenny, 1986, Bollen 2014, Duncan 1966), a large body of research has emerged within the causal inference literature that disentangles the tasks of causal definition, identification, and estimation. Causal mediation analysis seeks to uncover whether and how a treatment affects an outcome by quantifying the pathways through which a causal effect operates. Building upon the potential-outcomes framework and graphical causal models (Pearl 2009), a new body of research has provided model-free definitions of direct and indirect effects (Pearl 2001, Robins & Greenland 1992), established the assumptions needed for identifying these effects (Pearl 2001, Robins 2003), and developed an array of estimation strategies (e.g., VanderWeele 2015, 2016). These tools can help researchers discover mechanistic explanations, build theories, and design policy interventions. Sociologists would do well to consider these conceptual and computational tools in studies involving mechanisms. This section briefly reviews the causal approach to mediation analysis and its recent developments.

4.1. Estimating Direct and Indirect Effects

Let M denote a mediator hypothesized to transmit the effect of treatment W on the outcome Y . For example, Wodtke & Parbst (2017) investigated how school poverty mediates the effect of living in a disadvantaged neighborhood on a student's academic achievement. In this context, W is neighborhood disadvantage, Y is academic achievement, M is school poverty, and X is a set of background characteristics. **Figure 3** is a DAG representing the causal relationships involving these variables. Note that the pretreatment covariates X may confound not only the treatment–outcome relationship but also the treatment–mediator and mediator–outcome relationships.




Figure 3

Causal mediation analysis without treatment-induced confounding. W denotes treatment status, Y denotes the outcome of interest, M denotes a putative mediator, and X denotes observed pretreatment confounders.

The most common approach to assessing causal mediation involves decomposing the total effect of W on Y into two components: an indirect effect operating through the mediator M and a direct effect operating through alternative pathways not explicitly considered in the analysis. In **Figure 3**, we capture the indirect and direct effects by the causal paths $W \rightarrow M \rightarrow Y$ and $W \rightarrow Y$, respectively. These effects can be defined more formally using the potential-outcomes notation. Specifically, if we use Y_{wm} to denote the potential outcome under treatment status w and mediator value m , and M_w to denote the potential value of the mediator M under treatment status w , we can write the ATE of W on Y as

$$\tau_{\text{ATE}} = E[Y_{1M_1} - Y_{0M_0}], \quad 18.$$

which we can then decompose into the natural indirect effect (NIE) and natural direct effect (NDE):

$$\tau_{\text{ATE}} = \underbrace{E[Y_{1M_1} - Y_{1M_0}]}_{\tau_{\text{NIE}}} + \underbrace{E[Y_{1M_0} - Y_{0M_0}]}_{\tau_{\text{NDE}}}. \quad 19.$$

The NIE is the expected difference in the outcome if each unit were treated ($W = 1$) and subsequently exposed to the mediator value they experienced as a result of being treated (M_1) rather than the mediator value they would have experienced had they not been treated (M_0). For example, in Wodtke & Parbst's (2017) study, the NIE gauges the effect of neighborhood disadvantage operating through school poverty by fixing the level of neighborhood disadvantage ($W = 1$) for each student. They compare students' academic achievements under the levels of school poverty that they would have naturally experienced with neighborhood disadvantage (M_1) versus without neighborhood disadvantage (M_0). By contrast, the NDE reflects the average treatment effect if the mediator for each unit was fixed at its natural level under the reference treatment level ($W = 0$).

Both the NIE and NDE depend on Y_{1M_0} , a variable in which two different levels of W (0 and 1) are nested within the counterfactual for Y . Consequently, this counterfactual does not correspond to any experimental intervention on W and M . That is, to know the value of M_1 for a unit, it is necessary to set W to 1, but then this precludes setting W to 0 for M . The counterfactual Y_{1M_0} is thus called a cross-world counterfactual (Robins et al. 2022).

To identify the ATE from observational data, we invoke the unconfoundedness assumption, which states that after adjusting for a set of pretreatment covariates X , no additional confounders exist that affect both treatment status and the outcome. To identify the NIE and NDE, we need an unconfoundedness assumption for not only the treatment–outcome relationship but also the treatment–mediator and mediator–outcome relationships. Specifically, the NDE and NIE are nonparametrically identified if, after adjusting for pretreatment covariates X , there is (a) no unobserved treatment–outcome confounding, (b) no unobserved treatment–mediator confounding, and (c) no unobserved mediator–outcome confounding (Imai et al. 2010, VanderWeele & Vansteelandt 2009). Under these assumptions, the mean of the counterfactual $Y_{wM_{w^*}}$ for any $w, w^* \in \{0, 1\}$ can be identified using Pearl's (2001) mediation formula:

$$E[Y_{wM_{w^*}}] = \int E[Y|x, w, m] dP(m|x, w^*) dP(x), \quad 20.$$

where $P(\cdot)$ represents the cumulative distribution function of a random variable. We can use Equation 20 to identify the NIE and NDE by setting w and w^* at different values.

Given identification assumptions *a–c* and the mediation formula (Equation 20), we can use a variety of strategies to estimate the NIE and NDE. Imai et al. (2010) propose a regression-simulation estimator that involves first modeling the conditional mean of the outcome ($E[Y|x, w, m]$) and the conditional distribution of the mediator ($P(m|x, w)$) and then evaluating Equation 20 through

Monte Carlo draws from the estimated conditional distribution of the mediator. We can view this estimator as a plug-in estimator of Equation 20. Alternatively, one can rewrite Equation 20 as $E[E[Y|x, w, M]|x, w^*]$, which leads to a regression imputation estimator that involves modeling only the conditional means of the outcome (Vansteelandt et al. 2012), or as $E[\frac{I(W=w)}{p(w|X)} \frac{p(M|X, w^*)}{p(M|X, w)} Y]$, which leads to a weighting estimator that models the treatment's and the mediator's conditional distributions (VanderWeele 2009). Finally, drawing on semiparametric theory, Tchetgen Tchetgen & Shpitser (2012) develop a triply robust estimator of Equation 20 that involves fitting three models: first, a model for the conditional distribution of the treatment given pretreatment covariates (i.e., a propensity score model); second, a model for the conditional distribution of the mediator given the treatment and pretreatment covariates; and third, a model for the conditional mean of the outcome given the treatment, mediator, and pretreatment covariates. The resulting estimator is triply robust in that it is consistent if any two of the three models are correctly specified. Moreover, like the doubly robust estimator for the ATE, this triply robust estimator is particularly suitable for using flexible machine learning methods to estimate its nuisance functions (e.g., the treatment, mediator, and outcome models). This fact makes it highly attractive in high-dimensional settings.

The identification assumptions $a-c$ are strong and unverifiable, and the estimated NIE and NDE can be biased whenever unobserved confounding exists for any of the causal relationships involved. In practice, to assess the robustness of mediation analysis results to different forms of unobserved confounding, one can employ a set of general-purpose bias formulas developed by VanderWeele (2010) and VanderWeele & Arah (2011) (for a recent sociological application, see Brand et al. 2019b).

4.2. Treatment-Induced Confounding

Among identification assumptions $a-c$, c is especially restrictive because it requires that there must not be any observed or unobserved confounders of the mediator–outcome relationship that are affected by the treatment. This assumption is plausible if the treatment and mediator are temporally and mechanistically proximate to each other but is likely violated in other settings. For example, Klein & Kühhirt (2021) investigated the role of parental cognitive ability in mediating the effect of grandparents' education on grandchildren's cognitive ability. In this case, it is likely that some posttreatment variables, such as grandparents' income and occupational status, are affected by the treatment (grandparents' education) and affect both the mediator (parental cognitive ability) and the outcome (children's cognitive ability). **Figure 4** depicts a DAG where L is a treatment-induced confounder of the mediator–outcome relationship.




Figure 4

Causal mediation analysis with treatment-induced confounding. W denotes treatment status, Y denotes the outcome of interest, M denotes a putative mediator, X denotes observed pretreatment confounders, and L denotes treatment-induced confounders.

Treatment-induced confounders pose a dilemma for causal mediation analysis. If they were omitted, our estimated effects of the mediator on the outcome, and by extension, the estimated NIE and NDE, would be biased. However, controlling for treatment-induced confounders is also problematic because it blocks causal pathways and potentially unblocks noncausal pathways from the treatment to the outcome, leading to biased estimates of the NIE and NDE (Elwert & Winship 2014). In fact, the NIE and NDE are not nonparametrically identified in the presence of treatment-induced confounding. Scholars have proposed several strategies to address this challenge. First, the NIE and NDE can be identified in the presence of treatment-induced confounding if we impose an additional assumption positing that the treatment and mediator have no interaction effect (INE) on the outcome for each unit (Imai & Yamamoto 2013, Robins 2003). This assumption, however, is implausible in most applications because the no-interaction assumption must hold for every unit. To overcome this limitation, these scholars have developed sensitivity analysis methods for assessing the robustness of findings to potential violations of the no-interaction assumption.

Second, scholars have proposed an alternative class of estimands known as interventional direct and indirect effects (for a review, see Nguyen et al. 2021). Unlike the NIE and NDE, interventional effects can still be nonparametrically identified in the presence of treatment-induced confounding. Among interventional effects, a special case is the controlled direct effect (CDE), which measures the strength of the treatment–outcome relationship when the mediator is fixed at a given value for all units (Acharya et al. 2018, Pearl 2001, Robins 2003). A nonzero CDE thus implies that the effect of the treatment on the outcome does not operate exclusively through the mediator of interest. For example, in Klein & Kühhirt's (2021) study, a nonzero CDE would imply the effect of grandparent education on grandchildren's cognitive ability does not operate solely through parental cognitive ability.

Apart from the CDE, another set of interventional effects are the so-called randomized interventional analogs to the NDE (rNDE) and the NIE (rNIE) (Didelez et al. 2006, Geneletti 2007, VanderWeele et al. 2014). The rNDE and rNIE are like the NDE and NIE except that, instead of setting the mediator to the level it would have naturally been for each unit under a particular treatment status, these estimands involve setting the mediator to a value randomly drawn from its population distribution under a given treatment status. The rNDE and rNIE thus evaluate the effects of a hypothetical intervention on the distribution of a putative mediator. For example, Wodtke et al. (2020b) used the rNDE and rNIE to assess the extent to which school quality mediates the effect of neighborhood disadvantage on children's academic achievement.

Researchers can estimate interventional effects such as the CDE, rNDE, and rNIE via several alternative methods, such as sequential g-estimation (Vansteelandt 2009) and IPW (VanderWeele et al. 2014). More recently, Zhou & Wodtke (2019) proposed the regression-with-residuals (RWR) method, which is algebraically equivalent to sequential g-estimation in special cases but, unlike the latter, can accommodate several types of effect moderation (see also Wodtke & Zhou 2020). RWR has been applied in several sociological studies (e.g., Klein & Kühhirt 2021; Levy et al. 2019; Wodtke et al. 2020b, 2022). Nonetheless, as with sequential g-estimation and IPW, RWR is premised on a set of strong modeling assumptions, which, when violated, can lead to biased estimates. Scholars have recently leveraged semiparametric theory to reduce model dependence and develop more robust estimators of interventional direct and indirect effects (Díaz et al. 2021, Xia & Chan 2021). Researchers can combine these estimators with machine learning to yield optimal performance, like the doubly robust estimator for the ATE and the triply robust estimator for the NDE and NIE.

4.3. Causal Mediation Analysis with Multiple Mediators

Researchers often aim to test several competing hypotheses of underlying processes when analyzing causal mechanisms, leading to multiple mediators of interest. In the presence of multiple mediators, the prevailing practice is to treat different mediators as causally independent (i.e., assuming they do not affect each other) and then estimate the NIE for each mediator separately. In many applications, however, the mediators are likely causally dependent. In general, if two mediators are present and one mediator affects both the other mediator and the outcome, treating these mediators as causally independent may lead to biased estimates of the NIE for the second mediator. This approach would lead to bias because it fails to account for the first mediator as a potential confounder of the relationship between the second mediator and the outcome. However, to the extent that the first mediator is affected by the treatment, it is a treatment-induced confounder, which renders the NIE for the second mediator nonidentifiable without functional form assumptions. In such cases, we could attempt to evaluate the NIE via additional assumptions and sensitivity analysis (Imai & Yamamoto 2013) or consider interventional effects, such as the rNIE.¹⁸

Apart from interventional effects, other mediation estimands that can still be identified in the presence of multiple causally dependent mediators are path-specific effects (PSEs) (Avin et al. 2005). Specifically, suppose we have K causally ordered mediators $M_1, M_2, \dots, M_K$ that lie on the causal paths from W to Y . Then, under the assumption that no unobserved confounding exists for any of the treatment–mediator, treatment–outcome, and mediator–outcome relationships, the ATE can be decomposed into $K + 1$ PSEs: one direct effect and K mutually exclusive indirect effects that each reflect the contribution of a specific mediator beyond the contributions of its preceding mediators (Daniel et al. 2015, Zhou & Yamamoto 2022). Like the NIE and NDE, researchers can estimate these PSEs via regression–simulation (Miles et al. 2017), regression–imputation (Zhou & Yamamoto 2022), IPW (VanderWeele et al. 2014), or multiply robust methods that are amenable to machine learning estimation of its nuisance functions (Miles et al. 2020, Zhou 2022b). In a recent study, Ahearn et al. (2022) investigated the pathways through which college attendance increases voting, focusing on three sets of causally ordered mediators: degree completion, family formation and stability, and socioeconomic status. Using the regression–imputation approach, they estimated the corresponding PSEs.

5. TEMPORAL AND SPATIAL INTERFERENCE

Many sociological questions involve the study of effects over time or interactions within networks. Indeed, historical or life-cycle variation and network interactions lie at the center of sociological inquiry. But these settings complicate the definition and identification of causal effects. Just as sociologists studying temporal variation or network settings should consider causal processes, causal inference scholars should consider the complications involved in allowing treatments and effects to vary over time and interference between units under study. SUTVA posits that one unit's outcome is not affected by the treatment status of other units in the population. However, we often face temporal or spatial interference that renders SUTVA untenable. This section briefly reviews causal inference methods developed to study temporal and spatial interference.

5.1. Estimating Treatment Effects in the Presence of Temporal Interference

Temporal interference may arise in settings with time-varying treatments in which treatment status at a given time has not only contemporaneous effects, i.e., effects on outcomes measured

¹⁸Readers are directed to Reardon & Raudenbush (2013) for discussion of multisite, multiple-mediator IV models.




Figure 5

Causal inference with temporal interference. W_t denotes treatment status at time t , Y denotes the outcome of interest, X denotes baseline confounders, and L_t denotes time-varying confounders at time t , where $t \in \{1, 2\}$.

immediately thereafter, but also carry-over effects, i.e., effects on outcomes at later time points. For example, exposure to family instability in early childhood may differ from exposure in adolescence, and exposure may have both short-term and long-term effects on a child's cognitive and socioemotional development (Lee & McLanahan 2015). A common strategy to incorporate temporal interference in causal analysis is through Robins's (1986, 1997) extension of the potential-outcomes framework to time-varying treatments. Consider a study with $T \geq 2$ time points, where we are interested in the effect of a time-varying treatment W_t on an end-of-study outcome Y . Apart from a set of baseline or time-invariant confounders X , there is also a vector of observed time-varying confounders, L_t , that may be affected by prior treatments. Note that L_t may also include current or past outcome measures (Brand & Xie 2007). **Figure 5** shows a DAG representation of this setting when $T = 2$. In Lee & McLanahan's (2015) study of the relationship between family instability and child development, W_t denotes a family transition at time t , Y denotes a child's developmental outcome at time T , X includes a set of time-invariant covariates (e.g., mother's education), and L_t includes a set of time-varying covariates (e.g., poverty status). Following Robins et al. (2000), we use overbars to denote treatment histories such that $\bar{W} = (W_1, \dots, W_T)$ represents the observed treatment history until the end of the study and $Y_{\bar{w}}$ represents the potential outcome under a given treatment history $\bar{w}$. This notation allows us to consider various treatment effects based on contrasts between potential outcomes. For instance, with two time points ($T = 2$), we could consider the distal treatment effect (DTE), defined as

$$\tau_{\text{DTE}} = E[Y_{1,0} - Y_{0,0}], \quad 21.$$

which captures the average effect of receiving treatment only at time 1 rather than never. Alternatively, we could consider the following treatment effects (Wodtke et al. 2020a):

$$\tau_{\text{PTE}} = E[Y_{w_1,1} - Y_{w_1,0}], \quad 22.$$

$$\tau_{\text{CTE}} = E[Y_{1,1} - Y_{0,0}], \quad 23.$$

and

$$\tau_{\text{INE}} = E[(Y_{1,1} - Y_{1,0}) - (Y_{0,1} - Y_{0,0})], \quad 24.$$

where τ_{PTE} is the proximal treatment effect (PTE), τ_{CTE} is the cumulative treatment effect (CTE), and τ_{INE} is the interaction effect of treatments at time points 1 and 2. Note that if there is no

temporal interference, i.e., if the potential outcome Y_{w_1, w_2} depends only on treatment status at time 2, the DTE and the INE will both be zero, and the PTE will equal the CTE.

To identify the various causal contrasts considered above, it suffices to identify the expected potential outcome $E[Y_{\bar{w}}]$ for every treatment sequence $\bar{w}$. A key identification assumption for this quantity is sequential ignorability, which states that treatment at each time point is unconfounded conditional on past treatments and observed confounders. Although it does not allow for unobserved confounding, the assumption of sequential ignorability allows for both carryover effects (i.e., past treatments affect current outcomes) and feedback effects (i.e., past outcomes affect current treatments). These are typically assumed away in fixed-effects models (Imai & Kim 2019).¹⁹ Under sequential ignorability, we can estimate the expected potential outcome $E[Y_{\bar{w}}]$ via various parametric and semiparametric methods. A common method is the IPW estimation of marginal structural models (MSMs) (Robins et al. 2000; see Wodtke et al. 2011 for a sociological application).²⁰ Apart from MSMs, researchers can also assess time-varying treatment effects through structural nested mean models and their associated estimators, such as the g-estimator (e.g., Naimi et al. 2017, Vansteelandt 2009, Vansteelandt & Sjolander 2016) and the RWR estimator (Wodtke 2020, Wodtke et al. 2020a). These estimators involve modeling the conditional mean of the outcome as well as the conditional means and distributions of time-varying confounders. As with IPW, these methods are based on a set of strong modeling assumptions, which, when violated, can lead to biased estimates. To reduce model dependence, Bang & Robins (2005) propose a semiparametric estimator for the expected potential outcome $E[Y_{\bar{w}}]$. This estimator involves fitting $2 \cdot T$ models: a propensity score model at each time point and a model for an iteratively imputed outcome at each time point. This estimator is multiply robust in that it is consistent whenever the first k propensity score models and the last $K - k$ outcome models are correctly specified, where k can be any integer from 0 to K (Rotnitzky et al. 2017). The estimating equations are amenable to using DML.²¹ Given its reduced dependence on model specification and complementarity with machine learning, we encourage sociologists to use this semiparametric estimator (Bang & Robins 2005) and its variants (e.g., van der Laan & Rose 2018) more widely in future research.²²

5.2. Estimating Treatment Effects in the Presence of Spatial Interference

Spatial interference may arise in settings where units under consideration are not isolated but are connected by a common physical or social space, such as schools, neighborhoods, and friendship

¹⁹Elwert & Pfeffer (2019) incorporate future treatments as a proxy for an unmeasured confounder to address selection bias and discuss the conditions under which future values of the treatment can reduce or fully remove bias.

²⁰The method of IPW involves modeling the conditional distribution of treatment at each time point t given past treatments and observed confounders. It is difficult to use when the treatment is continuous, and is often inefficient and susceptible to large finite sample biases. To overcome these limitations, Zhou & Wodtke (2020) propose an alternative method of constructing weights for MSMs called residual balancing. It can be viewed as an extension of balancing weights (see Section 2) to longitudinal settings with temporal interference.

²¹Specifically, if cross-fitting is used and the estimators of the propensity score and outcome models all converge to the truth at a faster-than- $n^{-1/4}$ rate, the Bang–Robins estimator will be root- n consistent, asymptotically normal, and semiparametrically efficient.

²²We may also be interested in situations in which both the treatment and an effect moderator vary over time. Wodtke & Almirall (2017) describe moderated intermediate causal effects and structural nested mean models for analyzing effect moderation in a longitudinal setting. Using this approach, they examine whether the effects of time-varying exposure to poor neighborhoods on the risk of adolescent childbearing are moderated by time-varying family income.

networks, leading to spillover effects. In such settings, one unit's potential outcome is a function of not only its treatment status but also the treatment status of other related units (Aronow & Samii 2017, Athey et al. 2018a, Tchetgen Tchetgen & VanderWeele 2012, VanderWeele 2015). Such interferences, or interactions, are prevalent in social settings (An 2018, An & VanderWeele 2022, Egami 2021). For example, encouraging an individual to vote through some intervention can increase the turnout for household members (Imai & Jiang 2020). In some cases, such interactions are the focus of analysis; in other cases, they are considered a nuisance to estimating treatment effects (given the assumption of no interference) (Hong & Raudenbush 2015, Ogburn et al. 2022). Yet ignoring interference can lead to biased estimates of causal effects and incorrect statistical inferences (An 2018, Basse & Airolidi 2018, Lee & Ogburn 2021).

When the pattern of interference is unconstrained, spillover effects are hard to study because (a) the number of counterfactuals for each unit increases exponentially as the number of units increases, leading to many causal contrasts that can be hard to estimate nonparametrically, and (b) the outcomes of different units will be dependent, complicating statistical inference. Given these challenges, researchers often study spillover effects under two simplifying assumptions. First, the partial interference assumption posits that individuals are clustered in groups so that interference is limited to individuals within the same group (Sobel 2006). Second, the stratified inference assumption posits that within the same group, the effect of other units' treatment statuses on a focal unit's outcome operates through a known summary function g (e.g., the mean treatment status among other units in the same group) (Hudgens & Halloran 2008). The stratified inference assumption is quite strong, but it helps simplify the analysis, especially when there are more than a few units within each group. For example, when studying the spillover effect of grade retention on a child's test scores, Hong & Raudenbush (2015) invoke both assumptions by specifying a student's test score to be a function of their retention status and the retention rate of their peers in the same school.

Under the assumptions of partial interference and stratified interference, we can denote a unit's potential outcome as $Y_{w,g}$, where w denotes the unit's treatment status and g denotes a summary value of peer treatment status. The average individual effect can be defined as

$$E[Y_{1,g} - Y_{0,g}], \quad 25.$$

and the spillover effect can be defined as

$$E[Y_{w,g^*} - Y_{w,g}], \quad 26.$$

where g^* denotes an alternative value of peer treatment status (VanderWeele 2015). As shown by Hudgens & Halloran (2008), these effects can be identified and unbiasedly estimated using an experimental design with a two-stage randomization procedure (i.e., first at the group level and then at the individual level within groups). To identify these effects in observational studies, one needs to invoke a group-level unconfoundedness assumption—i.e., conditional on a set of group-level covariates (which may include their individual-level components), the treatment assignments of all units within a group are independent of their potential outcomes (Tchetgen Tchetgen & VanderWeele 2012). Under this assumption, researchers can estimate the average individual and spillover effects through various strategies such as IPW, regression imputation, and doubly robust methods (Liu et al. 2019). They can combine the doubly robust approach with DML to yield optimal performance (Park & Kang 2022).

In many social settings, people interact with each other through multiple channels and networks, such as friends, family, neighbors, and others. It is important to estimate the spillover effects that arise through each network; however, those network interactions are often unobserved, rendering unbiased estimation of spillover effects difficult. Egami (2021) develops sensitivity analysis

methods for assessing the potential influence of unobserved networks on causal findings. Relatedly, An (2018) emphasizes the importance of collecting data on treatment diffusion to measure treatment interference properly and then to estimate the direct treatment effect, treatment interference effect, and treatment effect on interference.

6. CONCLUSION

Over the past three decades, causal inference has been an active research area in sociology and related disciplines such as economics, statistics, computer science, and political science. While earlier developments in causal analysis, in the form of path analysis and structural equations, were developed primarily in sociology (e.g., Duncan 1966) and then exported to other fields, much of what constitutes today's sociological methodology on causal inference has heavily borrowed knowledge from other disciplines.

Our review updates the latest advances in causal inference methodology. Given the large size of this literature, we chose to focus on four topics: causal effect identification and estimation in general, causal effect heterogeneity, causal effect mediation, and temporal and spatial interference. Our choice reflects long-standing sociological interests in these topics: population heterogeneity (e.g., Brand & Xie 2010, Xie 2013, Xie et al. 2012), causal mechanisms (e.g., Duncan 1966), and the importance of historical or life-cycle variation and social context (e.g., Mason et al. 1983). As we reviewed, identifying and estimating causal components—a perennial objective in sociology—is no easy task with a counterfactual framework. There is no simple, one-size-fits-all solution. Causal inference with observational data, including quasi-experimental data, is a proposition specific to each research context. Often, researchers must invoke unverifiable assumptions and make consequential research decisions to draw causal conclusions. What makes sense for one research setting may not make sense for another. In applying new methods, we recommend that researchers thoroughly understand their underlying assumptions and trade-offs to apply them judiciously. Researchers should also assess how effects vary across the population and whether the results of their study and sample generalize to a broader population. Sociological inquiry also invites the careful analysis of mechanisms linking treatments to outcomes.

The past literature on causal inference has primarily been concerned with identification issues, while machine learning is often tasked with executing heavy computations with large data sets. The merge of the two strands of literature is facilitated by a long-recognized insight we discussed in the article: Causal effects can be highly heterogeneous across different units. We review the latest developments in causal inference that utilize machine-learning methods to learn about heterogeneous treatment effects. We also describe how researchers can use machine learning to minimize biases in estimating population-level quantities of interest, including direct and indirect effects and effects with temporal and spatial interference.

Moving forward, we expect the continuation of fruitful cross-disciplinary fertilization in this area. We also anticipate increased use of machine-learning methods in causal inference to reduce estimation biases and detect causal effect heterogeneity. Machine-learning methods are particularly attractive and feasible considering future improvements in computational power and increasing availability of large administrative, commercial, and digital trace data (often called big data) for social science research. However, we caution the reader that no computational method, machine-learning methods included, can solve what Holland (1986) called “the fundamental problem in causal inference”—i.e., we never observe counterfactual outcomes. Good research design is primary. Computation is useful but only secondary. Hence, the bridge between machine-learning methods and causal inference can be productive only with innovative and appropriate research designs to address social-scientifically sensible research questions.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

National Institutes of Health Grant R01 HD07460301A1 provided financial support for this research. The first author benefited from facilities and resources provided by the California Center for Population Research at UCLA (CCPR), which receives core support (P2C-HD041022) from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD). We thank affiliates of the Social Inequality Data Science Lab (<https://www.sidatasciencelab.org>), Ian Lundberg, Nathan Hoffmann, and an anonymous reviewer from the *Annual Review of Sociology* for helpful comments and suggestions. The ideas expressed herein are those of the authors.

LITERATURE CITED

- Abadie A, Cattaneo MD. 2018. Econometric methods for program evaluation. *Annu. Rev. Econ.* 10:465–503
- Abadie A, Imbens GW. 2016. Matching on the estimated propensity score. *Econometrica* 84:781–807
- Acharya A, Blackwell M, Sen M. 2018. Analyzing causal mechanisms in survey experiments. *Political Anal.* 26:357–78
- Ahearn C, Brand JE, Zhou X. 2022. How, and for whom, does higher education increase voting? *Res. High. Educ.* <https://doi.org/10.1007/s11162-022-09717-4>
- Alwin DF, Hauser RM. 1975. The decomposition of effects in path analysis. *Am. Sociol. Rev.* 40:37–47
- An W. 2010. Bayesian propensity score estimators: incorporating uncertainties in propensity scores into causal inference. *Sociol. Methodol.* 40:151–89
- An W. 2018. Causal inference with networked treatment diffusion. *Sociol. Methodol.* 48:152–81
- An W, VanderWeele TJ. 2022. Opening the blackbox of treatment interference: tracing treatment diffusion through network analysis. *Sociol. Methods Res.* 51:141–64
- An W, Winship C. 2017. Causal inference in panel data with application to estimating race-of-interviewer effects in the general social survey. *Sociol. Methods Res.* 46:68–102
- Angrist JD, Imbens GW, Rubin DB. 1996. Identification of causal effects using instrumental variables. *J. Am. Stat. Assoc.* 91:444–55
- Angrist JD, Pischke J-S. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton Univ. Press
- Angrist JD, Rokkanen M. 2015. Wanna get away? Regression discontinuity estimation of exam school effects away from the cutoff. *J. Am. Stat. Assoc.* 110:1331–44
- Aronow PM, Samii C. 2016. Does regression produce representative estimates of causal effects? *Am. J. Political Sci.* 60:250–67
- Aronow PM, Samii C. 2017. Estimating average causal effects under general interference, with application to a social network experiment. *Ann. Appl. Stat.* 11:1912–47
- Athey S, Eckles D, Imbens GW. 2018a. Exact p -values for network interference. *J. Am. Stat. Assoc.* 113:230–240
- Athey S, Imbens G. 2016. Recursive partitioning for heterogeneous causal effects. *PNAS* 113:7353–60
- Athey S, Imbens G. 2017. The state of applied econometrics: causality and policy evaluation. *J. Econ. Perspect.* 31:3–32
- Athey S, Imbens G. 2019. Machine learning methods that economists should know about. *Annu. Rev. Econ.* 11:685–725
- Athey S, Imbens G, Wager S. 2018b. Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *J. R. Stat. Soc. Ser. B* 80:597–623
- Athey S, Tibshirani J, Wager S. 2019. Generalized random forests. *Ann. Stat.* 47(2):1148–78
- Athey S, Wager S. 2021. Policy learning with observational data. *Econometrica* 89(1):133–61

- Austin PC, Stuart EA. 2017. Estimating the effect of treatment on binary outcomes using full matching on the propensity score. *Stat. Methods Med. Res.* 26:2505–25
- Avin C, Shpitser I, Pearl J. 2005. *Identifiability of path-specific effects*. Tech. Rep. R-321, Dep. Stat., Univ. Calif. Los Angeles. <https://escholarship.org/uc/item/45x689gq>
- Bang H, Robins JM. 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics* 61:962–73
- Baron RM, Kenny DA. 1986. The moderator–mediator variable distinction in social psychological research: conceptual, strategic, and statistical considerations. *J. Pers. Soc. Psychol.* 51:1173–82
- Basse GW, Airolidi EM. 2018. Limitations of design-based causal inference and A/B testing under arbitrary and network interference. *Sociol. Methodol.* 48:136–51
- Belloni A, Chernozhukov V, Hansen C. 2014. High-dimensional methods and inference on structural and treatment effects. *J. Econ. Perspect.* 28:29–50
- Bertanha M, Imbens GW. 2020. External validity in fuzzy regression discontinuity designs. *J. Bus. Econ. Stat.* 38:593–612
- Blandhol C, Bonney J, Mogstad M, Torgovitsky A. 2022. *When is TSLS actually LATE?* NBER Work. Pap. w29709
- Bloome D, Schrage D. 2021. Covariance regression models for studying treatment effect heterogeneity across one or more outcomes: understanding how treatments shape inequality. *Sociol. Methods Res.* 50:1034–72
- Bollen KA. 2014. *Structural Equations with Latent Variables*. New York: John Wiley & Sons
- Bound J, Jaeger DA, Baker RM. 1995. Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *J. Am. Stat. Assoc.* 90:443–50
- Brand JE. 2023. *Overcoming the Odds: The Benefits for Unlikely College Graduates*. New York: Russell Sage Found.
- Brand JE, Moore R, Song X, Xie Y. 2019a. Parental divorce is not uniformly disruptive to children's educational attainment. *PNAS* 116:7266–71
- Brand JE, Moore R, Song X, Xie Y. 2019b. Why does parental divorce lower children's educational attainment? A causal mediation analysis. *Sociol. Sci.* 6:264–92
- Brand JE, Simon Thomas J. 2013. Causal effect heterogeneity. In *Handbook of Causal Analysis for Social Research*, ed. SL Morgan, pp. 189–214. New York: Springer
- Brand JE, Simon Thomas J. 2014. Job displacement among single mothers: effects on children's outcomes in young adulthood. *Am. J. Sociol.* 119:955–1001
- Brand JE, Xie Y. 2007. Identification and estimation of causal effects with time-varying treatments and time-varying outcomes. *Sociol. Methodol.* 37:393–434
- Brand JE, Xie Y. 2010. Who benefits most from college? Evidence for negative selection in heterogeneous economic returns to higher education. *Am. Sociol. Rev.* 75:273–302
- Brand JE, Xu J, Koch B. 2020. Machine learning. In *Research Methods in the Social Sciences Foundation*, ed. P Atkinson, S Delamont, A Cernat, JW Sakshaug, RA Williams, pp. 1–27. Thousand Oaks, CA: SAGE
- Brand JE, Xu J, Koch B, Geraldo P. 2021. Uncovering sociological effect heterogeneity using machine-learning. *Sociol. Methodol.* 51:189–223
- Breiman L. 2001. Random forests. *Int. J. Mach. Learn. Cybern.* 45:5–32
- Caliendo M, Kopeinig S. 2008. Some practical guidance for the implementation of propensity score matching. *J. Econ. Surv.* 22:31–72
- Card D. 2001. Estimating the return to schooling: progress on some persistent econometric problems. *Econometrica* 69:1127–60
- Carranza AG, Krishnamurthy SK, Athey S. 2022. Flexible and efficient contextual bandits with heterogeneous treatment effect oracles. arXiv:2203.16668 [cs.LG]
- Cattaneo MD, Idrobo N, Titiunik R. 2019. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge, UK: Cambridge Univ. Press
- Cattaneo MD, Titiunik R. 2022. Regression discontinuity designs. *Annu. Rev. Econ.* 14:821–51
- Cheng S, Brand JE, Zhou X, Xie Y, Hout M. 2021. Heterogeneous returns to college over the life course. *Sci. Adv.* 7:eabg7641
- Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, Robins J. 2018. Double/debiased machine learning for treatment and structural parameters. *Econom. J.* 21:C1–68

- Cinelli C, Forney A, Pearl J. 2022. A crash course in good and bad controls. *Sociol. Methods Res.* In press. <https://doi.org/10.1177/00491241221099552>
- Cinelli C, Hazlett C. 2020. Making sense of sensitivity: extending omitted variable bias. *J. R. Stat. Soc. Ser. B* 82:39–67
- Cole SR, Stuart EA. 2010. Generalizing evidence from randomized clinical trials to target populations: the ACTG-320 trial. *Am. J. Epidemiol.* 172:107–15
- Daniel RM, De Stavola BL, Cousens SN, Vansteelandt S. 2015. Causal mediation analysis with multiple mediators. *Biometrics* 71(1):1–14
- Deaton A. 2010. Instruments, randomization, and learning about development. *J. Econ. Lit.* 48:424–55
- Deaton A, Cartwright N. 2018. Understanding and misunderstanding randomized controlled trials. *Soc. Sci. Med.* 210:2–21
- Díaz I, Hejazi NS, Rudolph KE, van Der Laan MJ. 2021. Nonparametric efficient causal mediation with intermediate confounders. *Biometrika* 108:627–41
- Didelez V, Dawid AP, Geneletti S. 2006. Direct and indirect effects of sequential treatments. In *UAI'06: Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence*, ed. R Dechter, T Richardson, pp. 138–46. Arlington, VA: AUAI
- Dong Y, Lewbel A. 2015. Identifying the effect of changing the policy threshold in regression discontinuity models. *Rev. Econ. Stat.* 97:1081–1092
- Duncan OD. 1966. Path analysis: sociological examples. *Am. J. Sociol.* 72:1–16
- Egami N. 2021. Spillover effects in the presence of unobserved networks. *Political Anal.* 29(3):287–316
- Egami N, Hartman E. 2022. Elements of external validity: framework, design, and analysis. *Am. Political Sci. Rev.* In press. <https://doi.org/10.1017/S0003055422000880>
- Elwert F. 2015. Graphical causal models. In *Handbook of Causal Analysis for Social Research*, ed. SL Morgan, pp. 245–73. New York: Springer
- Elwert F, Pfeffer FT. 2019. The future strikes back: using future treatments to detect and reduce hidden bias. *Sociol. Methods Res.* 51:1014–51
- Elwert F, Winship C. 2014. Endogenous selection bias: the problem of conditioning on a collider variable. *Annu. Rev. Sociol.* 40:31–53
- Felton C, Stewart B. 2022. Handle with care: a sociologist's guide to causal inference with instrumental variables. SocArXiv. <https://doi.org/10.31235/osf.io/3ua7q>
- Findley MG, Kikuta K, Denley M. 2021. External validity. *Annu. Rev. Political Sci.* 24:365–93
- Fisher RA. 1935. *The Design of Experiments*. Edinburgh: Oliver & Boyd
- Fong C, Hazlett C, Imai K. 2018. Covariate balancing propensity score for a continuous treatment: application to the efficacy of political advertisements. *Ann. Appl. Stat.* 12:156–77
- Gangl M. 2010. Causal inference in sociological research. *Annu. Rev. Sociol.* 36:21–47
- Gangl M. 2015. Partial identification and sensitivity analysis. In *Handbook of Causal Analysis for Social Research*, ed. SL Morgan, pp. 377–402. New York: Springer
- Geneletti S. 2007. Identifying direct and indirect effects in a non-counterfactual framework. *J. R. Stat. Soc. Ser. B* 69:199–215
- Gill RD, Robins JM. 2001. Causal inference for complex longitudinal data: the continuous case. *Ann. Stat.* 29:1785–811
- Grimmer J, Roberts ME, Stewart BE. 2021. Machine learning for social science: an agnostic approach. *Annu. Rev. Political Sci.* 24:395–419
- Hahn J, Todd P, Van der Klaauw W. 2001. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* 69:201–9
- Hahn PR, Murray J, Carvalho C. 2020. Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects. *Bayesian Anal.* 15:965–1056
- Hainmueller J. 2012. Entropy balancing for causal effects: a multivariate reweighting method to produce balanced samples in observational studies. *Political Anal.* 20:25–46
- Hartman E, Grieve R, Ramsahai R, Sekhon JS. 2015. From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects. *J. R. Stat. Soc. Ser. A* 178:757–78

- Hastie T, Tibshirani R, Friedman JH. 2017. *The Elements of Statistical Learning*. Berlin: Springer. 2nd ed.
- Heckman JJ, Humphries JE, Veramendi G. 2018. The nonmarket benefits of education and ability. *J. Hum. Cap.* 12:282–304
- Heckman JJ, Robb R. 1986. Alternative methods for solving the problem of selection bias in evaluating the impact of treatments on outcomes. In *Drawing Inferences from Self-Selected Samples*, ed. H Wainer, pp. 63–107. Mahwah, NJ: Lawrence Erlbaum
- Heckman JJ, Urzua S, Vytlacil E. 2006. Understanding instrumental variables in models with essential heterogeneity. *Rev. Econ. Stat.* 88:389–432
- Holland PW. 1986. Statistics and causal inference. *J. Am. Stat. Assoc.* 81:945–60
- Hong G, Raudenbush S. 2015. Heterogeneous agents, social interactions, and causal inference. In *Handbook of Causal Analysis for Social Research*, ed. SL Morgan, pp. 331–52. New York: Springer
- Hotz JV, Imbens GW, Mortimer JH. 2005. Predicting the efficacy of future training programs using past experiences at other locations. *J. Econom.* 125:241–70
- Huber M. 2023. *Causal Analysis: Impact Evaluation and Causal Machine Learning with Applications in R*. Cambridge, MA: MIT Press
- Hudgens MG, Halloran ME. 2008. Toward causal inference with interference. *J. Am. Stat. Assoc.* 103:832–42
- Imai K, Jiang Z. 2020. Identification and sensitivity analysis of contagion effects in randomized placebo-controlled trials. *J. R. Stat. Soc. Ser. A* 183:1637–57
- Imai K, Keele L, Yamamoto T. 2010. Identification, inference and sensitivity analysis for causal mediation effects. *Stat. Sci.* 25:51–71
- Imai K, Kim IS. 2019. When should we use unit fixed effects regression models for causal inference with longitudinal data? *Am. J. Political Sci.* 63:467–90
- Imai K, Ratkovic M. 2013. Estimating treatment effect heterogeneity in randomized program evaluation. *Ann. Appl. Stat.* 7:443–70
- Imai K, Ratkovic M. 2014. Covariate balancing propensity score. *J. R. Stat. Soc. Ser. B* 76:243–63
- Imai K, Yamamoto T. 2013. Identification and sensitivity analysis for multiple causal mechanisms: revisiting evidence from framing experiments. *Political Anal.* 21:141–71
- Imbens GW. 2004. Nonparametric estimation of average treatment effects under exogeneity: a review. *Rev. Econ. Stat.* 86:4–29
- Imbens GW. 2015. Matching methods in practice. *J. Hum. Resour.* 50:373–419
- Imbens GW, Angrist JD. 1994. Identification and estimation of local average treatment effects. *Econometrica* 62:467–75
- Imbens GW, Lemieux T. 2008. Regression discontinuity designs: a guide to practice. *J. Econom.* 142:615–35
- Imbens GW, Rubin D. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge, UK: Cambridge Univ. Press
- Josey KP, Yang F, Ghosh D, Raghavan S. 2022. A calibration approach to transportability and data-fusion with observational data. *Stat. Med.* 41:4511–31
- Kennedy EH, Ma Z, McHugh MD, Small DS. 2017. Non-parametric methods for doubly robust estimation of continuous treatment effects. *J. R. Stat. Soc. Ser. B* 79:1229–45
- Kern HL, Stuart EA, Hill J, Green DP. 2016. Assessing methods for generalizing experimental impact estimates to target populations. *J. Res. Educ. Eff.* 9:103–27
- Klein M, Kühhirt M. 2021. Direct and indirect effects of grandparent education on grandchildren’s cognitive development: the role of parental cognitive ability. *Sociol. Sci.* 8:265–84
- Künzel SR, Sekhon JS, Bickel PJ, Yu B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *PNAS* 116:4156–65
- Lee BK, Lessler J, Stuart E. 2010. Improving propensity score weighting using machine learning. *Stat. Med.* 29:337–46
- Lee D, McLanahan S. 2015. Family structure transitions and child development: instability, selection, and population heterogeneity. *Am. Sociol. Rev.* 80:738–63
- Lee DS, Lemieux T. 2010. Regression discontinuity designs in economics. *J. Econ. Lit.* 48:281–355
- Lee Y, Ogburn EL. 2021. Network dependence can lead to spurious associations and invalid inference. *J. Am. Stat. Assoc.* 116:1060–74

- Levy BL, Owens A, Sampson RJ. 2019. The varying effects of neighborhood disadvantage on college graduation: moderating and mediating mechanisms. *Sociol. Educ.* 92:269–92
- Liu L, Hudgens MG, Saul B, Clemens JD, Ali M, Emch ME. 2019. Doubly robust estimation in observational studies with partial interference. *Stat* 8:e214
- Lundberg I. 2022. The gap-closing estimand: a causal approach to study interventions that close disparities across social categories. *Sociol. Methods Res.* In press. <https://doi.org/10.1177/00491241211055769>
- Lundberg I, Brand JE, Jeon N. 2022. Researcher reasoning meets computational capacity: machine learning for social science. *Soc. Sci. Res.* 108:102807
- Lundberg I, Johnson R, Stewart BM. 2021. What is your estimand? Defining the target quantity connects statistical evidence to theory. *Am. Sociol. Rev.* 86:532–65
- Manski CF. 1995. *Identification Problems in the Social Sciences*. Cambridge, MA: Harvard Univ. Press
- Manski CF, Garfinkel I. 1992. Introduction. In *Evaluating Welfare and Training Programs*, ed. CF Manski, I Garfinkel, pp. 1–21. Cambridge, MA: Harvard Univ. Press
- Mason WM, Wong GY, Entwistle B. 1983. Contextual analysis through the multilevel linear model. *Sociol. Methodol.* 14:72–103
- McCaffrey DF, Ridgeway G, Morral AR. 2004. Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychol. Methods* 9:403
- Miles CH, Shpitser I, Kanki P, Meloni S, Tchetgen Tchetgen EJ. 2017. Quantifying an adherence path-specific effect of antiretroviral therapy in the Nigeria PEPFAR program. *J. Am. Stat. Assoc.* 112:1443–52
- Miles CH, Shpitser I, Kanki P, Meloni S, Tchetgen Tchetgen EJ. 2020. On semiparametric estimation of a path-specific effect in the presence of mediator-outcome confounding. *Biometrika* 107:159–72
- Mogstad M, Torgovitsky A. 2018. Identification and extrapolation of causal effects with instrumental variables. *Annu. Rev. Econ.* 10:577–613
- Molina M, Garip F. 2019. Machine learning for sociology. *Annu. Rev. Sociol.* 45:27–45
- Morgan S, Harding D. 2006. Matching estimators of causal effects: prospects and pitfalls in theory and practice. *Sociol. Methods Res.* 35:3–60
- Morgan S, Winship C. 2014. *Counterfactuals and Causal Inference: Methods and Principles for Social Research*. Cambridge, UK: Cambridge Univ. Press
- Naimi AJ, Cole SR, Kennedy EH. 2017. An introduction to g methods. *Int. J. Epidemiol.* 46:756–62
- Neyman J. 1923. On the application of probability theory to agricultural experiments. *Stat. Sci.* 5:465–80
- Nguyen TQ, Schmid I, Stuart EA. 2021. Clarifying causal mediation analysis for the applied researcher: defining effects based on what we want to learn. *Psychol. Methods* 26:255–71
- Nie X, Wager S. 2021. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika* 108:299–319
- Offer-Westort M, Coppock A, Green DP. 2021. Adaptive experimental design: prospects and applications in political science. *Am. J. Political Sci.* 65:826–44
- Ogburn EL, Sofrygin O, Díaz I, Van Der Laan MJ. 2022. Causal inference for social network data. *J. Am. Stat. Assoc.* In press. <https://doi.org/10.1080/01621459.2022.2131557>
- Park C, Kang H. 2022. Efficient semiparametric estimation of network treatment effects under partial interference. *Biometrika* 109:1015–31
- Pearl J. 2001. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, ed. J Breese, D Koller, pp. 411–20. Burlington, MA: Morgan Kaufmann
- Pearl J. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge Univ. Press
- Pearl J, Bareinboim E. 2014. External validity: from do-calculus to transportability across populations. *Stat. Sci.* 29:579–95
- Quandt R. 1972. A new approach to estimating switching regression. *J. Am. Stat. Assoc.* 67:306–10
- Reardon SF, Raudenbush SW. 2013. Under what assumptions do site-by-treatment instruments identify average causal effects? *Sociol. Methods Res.* 42:143–63
- Robins JM. 1986. A new approach to causal inference in mortality studies with a sustained exposure period-application to control of the healthy worker survivor effect. *Math. Model.* 7:1393–512
- Robins JM. 1997. Causal inference from complex longitudinal data. In *Latent Variable Modeling and Applications to Causality*, ed. M Berkane, pp. 69–117. New York: Springer

- Robins JM. 2003. Semantics of causal DAG models and the identification of direct and indirect effects. In *Highly Structured Stochastic Systems*, ed. PJ Green, NL Hjort, S Richardson, pp. 70–81. Oxford, UK: Oxford Univ. Press
- Robins JM, Greenland S. 1992. Identifiability and exchangeability for direct and indirect effects. *Epidemiology* 3:143–55
- Robins JM, Hernan MA, Brumback B. 2000. Marginal structural models and causal inference in epidemiology. *Epidemiology* 11:550–60
- Robins JM, Richardson TS, Shpitser I. 2022. An interventionist approach to mediation analysis. In *Probabilistic and Causal Inference: The Works of Judea Pearl*, ed. H Geffner, R Dechter, JY Halpern, pp. 713–64. New York: ACM
- Robins JM, Rotnitzky A. 1995. Semiparametric efficiency in multivariate regression models with missing data. *J. Am. Stat. Assoc.* 90:122–29
- Robins JM, Rotnitzky A, Zhao LP. 1994. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.* 89:846–66
- Rosenbaum PR, Rubin DB. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70:41–55
- Rotnitzky A, Robins JM, Babino L. 2017. On the multiply robust estimation of the mean of the g-functional. arXiv:1705.08582 [stat.ME]
- Roy AD. 1951. Some thoughts on the distribution of earnings. *Oxf. Econ. Pap.* 3:135–46
- Rubin DB. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* 66: 688–701
- Rubin DB. 1977. Assignment to treatment group on the basis of a covariate. *J. Educ. Stat.* 2:1–26
- Rubin DB. 1986. Which ifs have causal answers? Discussion of “Statistics and Causal Inference” by Holland. *J. Am. Stat. Assoc.* 83:396
- Scharfstein DO, Rotnitzky A, Robins JM. 1999. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J. Am. Stat. Assoc.* 94:1096–120
- Scott SL. 2010. A modern Bayesian look at the multi-armed bandit. *Appl. Stoch. Models Bus. Ind.* 26:639–58
- Semenova V, Chernozhukov V. 2021. Debiased machine learning of conditional average treatment effects and other causal functions. *Econom. J.* 24:264–89
- Shu X, Ye Y. 2023. Knowledge discovery: methods from data mining and machine learning. *Soc. Sci. Res.* 110:102817
- Sobel ME. 2006. What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *J. Am. Stat. Assoc.* 101:1398–407
- Steiner PM, Kim Y, Hall CE, Su D. 2017. Graphical models for quasi-experimental designs. *Sociol. Methods Res.* 46:155–88
- Stuart EA, Bradshaw CP, Leaf PJ. 2015. Assessing the generalizability of randomized trial results to target populations. *Prev. Sci.* 16:475–85
- Stuart EA, Cole SR, Bradshaw CP, Leaf PJ. 2011. The use of propensity scores to assess the generalizability of results from randomized trials. *J. R. Stat. Soc. Ser. A* 174:369–86
- Tchetgen Tchetgen EJ, Shpitser I. 2012. Semiparametric theory for causal mediation analysis: efficiency bounds, multiple robustness, and sensitivity analysis. *Ann. Stat.* 40:1816–45
- Tchetgen Tchetgen EJ, VanderWeele TJ. 2012. On causal inference in the presence of interference. *Stat. Methods Med. Res.* 21:55–75
- Tipton E. 2013. Improving generalizations from experiments using propensity score subclassification: assumptions, properties, and contexts. *J. Educ. Behav. Stat.* 38:239–66
- Tipton E, Hedges L, Vaden-Kiernan M, Borman G, Sullivan K, Caverly S. 2014. Sample selection in randomized experiments: a new method using propensity score stratified sampling. *J. Res. Educ. Eff.* 7:114–35
- van der Laan MJ, Rose S. 2018. *Targeted Learning in Data Science*. New York: Springer
- van der Laan MJ, Rubin D. 2006. Targeted maximum likelihood learning. *Int. J. Biostat.* 2:11
- VanderWeele TJ. 2009. Marginal structural models for the estimation of direct and indirect effects. *Epidemiology* 20:18–26

- VanderWeele TJ. 2010. Bias formulas for sensitivity analysis for direct and indirect effects. *Epidemiology* 21:540–51
- VanderWeele TJ. 2015. *Explanation in Causal Inference: Methods for Mediation and Interaction*. Oxford, UK: Oxford Univ. Press
- VanderWeele TJ. 2016. Mediation analysis: a practitioner's guide. *Annu. Rev. Public Health* 37:17–32
- VanderWeele TJ, Arah OA. 2011. Unmeasured confounding for general outcomes, treatments, and confounders: bias formulas for sensitivity analysis. *Epidemiology* 22:42–52
- VanderWeele TJ, Vansteelandt S. 2009. Conceptual issues concerning mediation, interventions and composition. *Stat. Interface* 2:457–68
- VanderWeele TJ, Vansteelandt S, Robins JM. 2014. Effect decomposition in the presence of an exposure-induced mediator-outcome confounder. *Epidemiology* 25:300–6
- Vansteelandt S. 2009. Estimating direct effects in cohort and case-control studies. *Epidemiology* 20:851–60
- Vansteelandt S, Bekaert M, Lange T. 2012. Imputation strategies for the estimation of natural direct and indirect effects. *Epidemiol. Methods* 1:131–158
- Vansteelandt S, Sjolander A. 2016. Revisiting g-estimation of the effect of a time-varying exposure subject to time-varying confounding. *Epidemiol. Methods* 5:37–56
- Wager S, Athey S. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *J. Am. Stat. Assoc.* 113:1228–42
- Westreich D, Edwards JK, Lesko CR, Cole SR, Stuart EA. 2019. Target validity and the hierarchy of study designs. *Am. J. Epidemiol.* 188:438–43
- Westreich D, Lessler J, Funk MJ. 2010. Propensity score estimation: neural networks, support vector machines, decision trees (CART), and meta-classifiers as alternatives to logistic regression. *J. Clin. Epidemiol.* 63:826–33
- Winship C, Morgan SL. 1999. The estimation of causal effects from observational data. *Annu. Rev. Sociol.* 25:659–707
- Wodtke GT. 2020. Regression-based adjustment for time-varying confounders. *Sociol. Methods Res.* 49:906–46
- Wodtke GT, Alaca Z, Zhou X. 2020a. Regression-with-residuals estimation of marginal effects: a method of adjusting for treatment-induced confounders that may also be effect modifiers. *J. R. Stat. Soc. Ser. A* 183:311–32
- Wodtke GT, Almirall D. 2017. Estimating moderated causal effects with time-varying treatments and time-varying moderators: structural nested mean models and regression with residuals. *Sociol. Methodol.* 47:212–45
- Wodtke GT, Harding DJ, Elwert F. 2011. Neighborhood effects in temporal perspective: the impact of long-term exposure to concentrated disadvantage on high school graduation. *Am. Sociol. Rev.* 76:713–36
- Wodtke GT, Parbst M. 2017. Neighborhoods, schools, and academic achievement: a formal mediation analysis of contextual effects on reading and mathematics abilities. *Demography* 54:1653–76
- Wodtke GT, Ramaj S, Schachner J. 2022. Toxic neighborhoods: the effects of concentrated poverty and environmental lead contamination on early childhood development. *Demography* 59:1275–98
- Wodtke GT, Yildirim U, Harding DJ, Elwert F. 2020b. *Are neighborhood effects explained by differences in school quality?* Work. Pap. 102-20, Inst. Res. Labor Employ., Univ. Calif., Berkeley, CA
- Wodtke GT, Zhou X. 2020. Effect decomposition in the presence of treatment-induced confounding: a regression-with-residuals approach. *Epidemiology* 31:369–75
- Xia F, Chan KCG. 2021. Identification, semiparametric efficiency, and quadruply robust estimation in mediation analysis with treatment-induced confounding. *J. Am. Stat. Assoc.* <https://doi.org/10.1080/01621459.2021.1990765>
- Xie Y. 2013. Population heterogeneity and causal inference. *PNAS* 110:6262–68
- Xie Y, Brand JE, Jann B. 2012. Estimating heterogeneous treatment effects with observational data. *Sociol. Methodol.* 42:314–47
- Xie Y, Near C, Xu H, Song X. 2020. Heterogeneous treatment effects on children's cognitive/non-cognitive skills: a reevaluation of an influential early childhood intervention. *Soc. Sci. Res.* 86:102389
- Yadlowsky S, Fleming S, Shah N, Brunskill E, Wager S. 2021. Evaluating treatment prioritization rules via rank-weighted average treatment effects. arXiv:2111.07966 [stat.ME]

- Zhou X. 2019. Equalization or selection? Reassessing the “meritocratic power” of a college degree in intergenerational income mobility. *Am. Sociol. Rev.* 84:459–85
- Zhou X. 2022a. Attendance, completion, and heterogeneous returns to college: a causal mediation approach. *Sociol. Methods Res.* In press. <https://doi.org/10.1177/00491241221113876>
- Zhou X. 2022b. Semiparametric estimation for causal mediation analysis with multiple causally ordered mediators. *J. R. Stat. Soc. Ser. B* 84:794–821
- Zhou X, Pan G. 2023. Higher education and the black-white earnings gap. *Am. Sociol.* 88(1):154–88
- Zhou X, Wodtke GT. 2019. A regression-with-residuals method for estimating controlled direct effects. *Political Anal.* 27:360–69
- Zhou X, Wodtke GT. 2020. Residual balancing: a method of constructing weights for marginal structural models. *Political Anal.* 28:487–506
- Zhou X, Xie Y. 2019. Marginal treatment effects from a propensity score perspective. *J. Political Econ.* 127:3070–84
- Zhou X, Xie Y. 2020. Heterogeneous treatment effects in the presence of self-selection: a propensity score perspective. *Sociol. Methodol.* 50:350–85
- Zhou X, Yamamoto T. 2022. Tracing causal paths from experimental and observational data. *J. Politics* 85(1):250–65
- Zubizarreta JR. 2015. Stable weights that balance covariates for estimation with incomplete outcome data. *J. Am. Stat. Assoc.* 110:910–22