

Causal Inference in Epidemiology

Douglas Ezra Morrison

Last modified: 2026-05-07: 6:37:20 (UTC)

Contents

Preface	1
UC Davis courses	1
Online Videos	1
UC Davis Datalab learning group	2
Books	2
1 Introduction to causal inference	3
1.1 Introduction	3
1.2 Individual causal effects	3
2 Collider bias	6
2.1 Definition	6
2.2 Classic example: Sackett’s hospitalized patients	6
2.3 Modern example: The obesity paradox	7
2.4 Prevention	7
3 Bias analyses	8
4 Difference in differences analyses	9
4.1 Change in Changes	9
5 Further resources	12
5.1 The Catalog of Bias	12
6 Summary	13
References	14

Preface

This will be an online book about causal inference.

Here are some other resources for learning causal inference:

UC Davis courses

- EPI 205¹ “Principles of Epidemiology”
- EPI 206² “Epidemiologic Study Design”
- EPI/SPH 207³ “Advanced Epidemiologic Methodology”
- EPI 225⁴ “Advanced Topics in Epidemiology Methods”
- POL 285⁵ “Statistics of Causal Inference in Political Science”
- MGB/MGP/MGT 454A⁶ “Causal Inference and Statistical Experiments”
 - syllabus: <https://webapps.aws.ucdavis.edu/public/documents/4861649/Syllabus><https://schedule.aws.ucdavis.edu/public/documents/5319910/Syllabus>
- PSC 204B⁷ “Causal Modeling of Correlational Data”
- PSC 205C⁸ “Structural Equation Modeling”

Course search options:

- <https://schedule.aws.ucdavis.edu/courseScheduling>
- <https://catalog.ucdavis.edu/course-search/>
- <https://catalog.ucdavis.edu/courses-subject-code/>

Online Videos

- “Introduction to Causal Inference”⁹ (slides here¹⁰)
- Online Causal Inference Seminar series¹¹

¹<https://catalog.ucdavis.edu/search/?q=EPI+205>

²<https://catalog.ucdavis.edu/search/?q=EPI+206>

³<https://catalog.ucdavis.edu/search/?q=EPI+207>

⁴<https://catalog.ucdavis.edu/search/?q=EPI+225>

⁵<https://catalog.ucdavis.edu/search/?q=POL+285>

⁶<https://catalog.ucdavis.edu/search/?q=MGB+454A>

⁷<https://catalog.ucdavis.edu/search/?q=PSC+204B>

⁸<https://catalog.ucdavis.edu/search/?q=PSC+205C>

⁹<https://ucdhs.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=43e9eb6f-3ed9-41ac-8ad9-ae22016572c8%20>

¹⁰<https://health.ucdavis.edu/media-resources/ctsc/documents/pdfs/causal-inference-intro-2022.pdf>

¹¹<https://www.youtube.com/channel/UCiiOj5GSES6uw21kfXnxj3A/videos>

UC Davis Datalab learning group

- <https://datalab.ucdavis.edu/causal-inference/>
 - Reading list¹²

Other links:

- <https://cameron.econ.ucdavis.edu/causal/>
- https://datalab-icmat.github.io/causal_reading_group.html
- Lab exercises by Ben Noble¹³: <https://github.com/bennoble/causal-inference-2022>

Books

- Judea Pearl (2016)
- Hernán and Robins (2020)

¹²https://docs.google.com/document/d/1K0QZFSjQIYnOTahpRK7Q83eaiIIFjfa-clSbj_ifgco/edit?tab=t.0#heading=h.farbmh6n76gq

¹³<https://benjaminncoble.org/>

1 Introduction to causal inference

1.1 Introduction

Hernán and Robins (2020) (page vii):

Unfortunately, the scientific literature is plagued by studies in which the causal question is not explicitly stated and the investigators' unverifiable assumptions are not declared. This casual attitude towards causal inference has led to a great deal of confusion.

Rigorously defining cause and effect is difficult. Fortunately, many humans have strong intuitions about these concepts. We will make cursory attempts at definitions for the basic terms, and leave the finer points to philosophers.

1.2 Individual causal effects

Definition 1.1 (Action, intervention, exposure, policy, treatment). An **action** (also called an **intervention**, **exposure**, **policy**, or **treatment**) is a choice that we consider making.

Definition 1.2 (Random variable). A **random variable** is a variable that may have different values for different individuals and/or for different actions/exposures that an individual might experience.¹

Definition 1.3 (potential outcome, consequence). A **potential outcome** of action a (also called a **consequence** of a) is the value of a random variable Y that would occur if we were to take action a . The potential outcome of action a on random variable Y is often denoted $Y(a)$, Y^a , or $Y^{A=a}$. We will use notation $Y(a)$.

¹adapted from Hernán and Robins (2020) p3

Definition 1.4 (Factual outcome, observed outcome). A *factual outcome* (or *observed outcome*) is the potential outcome corresponding to an action that was actually taken.

If we consider taking action a or an alternative a' , and we actually take action $A = a$, then $Y(a)$ is the factual outcome.

Definition 1.5 (Counterfactual outcome). A counterfactual outcome is a potential outcome corresponding to an action that was not actually taken.

If we consider taking action a or an alternative a' , and we actually take action $A = a$, then $Y(a')$ is a counterfactual outcome.

There might be more than one counterfactual outcome, depending on how many action options were considered, but there can only ever be one factual outcome per random variable.

Definition 1.6 (Cause). Action a **causes** outcome y (or is a cause of y) if:

- outcome y would occur if we were to take action a
and
- outcome y would not occur if we did not take action a .

In other words, if:

- $Y(a) = y$
and
- $\exists a' \neq a : Y(a') \neq y$

Definition 1.7 (effect). The effect of action a on outcome Y , relative to a given alternative action a' , is the contrast in potential outcomes, $Y(a)$ versus $Y(a')$.

Definition 1.8 (Consistency). Consistency is the assumption that if we observe an action a , then the observed outcome Y is equal to the “factual potential outcome” corresponding to action a ; in other words, if $A = a$, then $Y(a) = Y$.

Definition 1.9 (Exchangeability). Subpopulations defined by exposure X are exchangeable with respect to a potential outcome $Y(x)$ if the distribution of $Y(x)$ does not depend on the observed exposure X :

$$Y(x) \perp\!\!\!\perp X$$

Theorem 1.1. *If subpopulations defined by values of exposure X are exchangeable with respect to potential outcome $Y(x)$, then the expected value of $Y(x)$ does not depend on the observed value of X :*

$$E[Y(x)|X = x'] = E[Y(x)|X = x]$$

Definition 1.10 (Conditional exchangeability). Subpopulations defined by exposure X are exchangeable with respect to a potential outcome $Y(x)$ if the distribution of $Y(x)$ does not depend on the observed exposure X , conditional on covariate(s) Z :

$$Y(x) \perp\!\!\!\perp X | \tilde{Z}$$

Theorem 1.2. *If subpopulations defined by values of exposure X are conditionally exchangeable with respect to potential outcome $Y(x)$ given covariate $\tilde{Z}$, then the expected value of $Y(x)$ does not depend on the observed value of X :*

$$E[Y(x)|X = x', \tilde{Z} = \tilde{z}] = E[Y(x)|X = x, \tilde{Z} = \tilde{z}]$$

Definition 1.11 (Fundamental problem of causal inference). The *fundamental problem of causal inference* is that only one potential outcome (the factual outcome) can be observed per person (or per sampling unit, more generally) (Holland 1986). The other, counterfactual outcomes, are all missing data, and thus, the causal effects are all missing data as well.

2 Collider bias

2.1 Definition

A **collider** is a variable that is caused by two or more other variables in a causal diagram (directed acyclic graph or DAG) (Catalog of Bias Collaboration 2024b). The term “collider” comes from the visual representation where multiple arrows “collide” into the variable.

The basic collider structure is an **inverted fork (V-structure)**:

Exposure $\rightarrow$ Collider $\leftarrow$ Outcome

More complex patterns exist, such as **M-structures** where multiple colliders are connected in chains, creating additional opportunities for bias when conditioning on any collider in the structure.

Inappropriately controlling for a collider variable, by study design or statistical analysis, results in **collider bias** (Catalog of Bias Collaboration 2024b).

Controlling for a collider can induce a distorted association between the exposure and outcome when in fact none exists. This bias predominantly occurs in observational studies.

Because collider bias can be induced by sampling (restricting the study to certain values of the collider), selection bias can sometimes be considered to be a form of collider bias (Catalog of Bias Collaboration 2024b).

2.2 Classic example: Sackett’s hospitalized patients

A classic example of collider bias was provided by Sackett (1979) and discussed by Catalog of Bias Collaboration (2024b). He analyzed data from 257 hospitalized individuals and detected an association between locomotor disease and respiratory disease (odds ratio 4.06). However, when he repeated the analysis in a sample of 2,783 individuals from the general population, he found no association (odds ratio 1.06). The original analysis of hospitalized individuals was biased because both diseases caused individuals to be hospitalized. By looking only within the stratum of hospitalized individuals, Sackett had observed a distorted association. In this example, locomotor disease and respiratory disease are independent causes of hospitalization—the collider. Controlling for the collider by study design (selection bias) induced a distorted association between the two diseases.

2.3 Modern example: The obesity paradox

A more recent example can be seen in the “obesity paradox”—an apparent protective effect of obesity on mortality in individuals with chronic conditions such as cardiovascular disease (CVD) (Catalog of Bias Collaboration 2024b). In fact, obesity increases mortality rates in the general population. The collider bias occurs when investigators condition on CVD (by design or analysis), resulting in a distorted association. Consequently, in a sample that includes only patients with CVD, obesity falsely appears to protect against mortality, whereas in the wider population, obesity increases the risk of early death. Banack and Kaufman (2014) showed using the third US National Health and Nutrition Examination Survey (NHANES III) that the unbiased mortality risk ratio for the entire cohort was 1.24 [95% CI = 1.11, 1.39] (harmful), but the biased stratum-specific mortality risk ratio in patients with CVD was 0.79 [95% CI = 0.68, 0.91] (protective).

2.4 Prevention

Collider bias can be prevented by carefully applying appropriate inclusion criteria—making sure that the exposure and outcome of interest do not drive inclusion or selective retention in a study (Catalog of Bias Collaboration 2024b). Causal diagrams (directed acyclic graphs or DAGs) can help identify colliders that should be left uncontrolled and confounders that should be controlled.

For more detailed information about collider bias, including causal diagrams and additional examples, see Catalog of Bias Collaboration (2024b).

3 Bias analyses

“Continuation of Annual Screening Mammography and Breast Cancer Mortality in Women Older Than 70 Years” (2020), supplementary materials p16:

A sensitivity analysis for unmeasured confounding can only conclude that, if there is much unmeasured confounding, there will be much bias in the estimates and, conversely, if there is little unmeasured confounding, there will be little bias in the estimates. But, if the magnitude of unmeasured confounding is unknown (as it often is), then the magnitude of confounding bias is also unknown. That is, sensitivity analysis for unmeasured confounding is, by definition, uninformative. The situation is different when investigators have some information about the unmeasured confounder(s). For example, our analysis did not adjust for cigarette smoking, which may be associated with both breast cancer screening (because smokers are less likely to use preventive services) and breast cancer death. If we knew the prevalence of smoking in screening-defined groups and the association between smoking and breast cancer death, then we could use published methods to correct the confounded effect estimates.

Unfortunately, these methods assume that the residual confounding can be summarized by a single variable that happens to be dichotomous, time-fixed, and unassociated with all measured variables that were adjusted for in the analysis (Lash et al. 2009; Schneeweiss 2006).

See also Fox et al. (2022) for an updated reference.

4 Difference in differences analyses

Many approaches to causal inference assume exchangeability (Definition 1.10) and exploit its consequence (Theorem 1.1):

$$E[Y(x)|X = x'] = E[Y(x)|X = x]$$

Difference-in-differences makes a weaker exchangeability assumption:

$$E[Y_t(0) - Y_{t'}(0)|X = 1] = E[Y_t(0) - Y_{t'}(0)|X = 0]$$

4.1 Change in Changes

The **Change in Changes (CiC)** model (Athey and Imbens 2006) is a DiD-type method that estimates the **Quantile Treatment Effect on the Treated (QTET)** — that is, treatment effects across the full distribution of outcomes, not just the mean (Callaway 2024).

CiC requires two periods of data (a pre-treatment and a post-treatment period). The data can be either repeated cross-sections or panel data (Callaway 2024).

4.1.1 Assumption

Rather than the standard parallel trends assumption (that average outcomes would have followed parallel paths absent treatment), CiC assumes that the *distribution* of untreated potential outcomes evolves over time in the same way for both treated and control groups (Athey and Imbens 2006).

4.1.2 Covariate adjustment

CiC can condition on covariates by first fitting a linear model for outcomes conditional on group–time indicators and covariates, then residualizing (removing predicted values), and finally applying the CiC estimator to these quasi-residuals (Athey and Imbens 2006; Callaway 2024).

4.1.3 R implementation

The `CiC()` function in the `qte` R package (Callaway 2024) implements this estimator. Key arguments are (note: `formula` and `xformula` are the actual argument names in the package):

Argument	Description
<code>formula</code>	$y \sim d$ where y is the outcome and d is a binary treatment indicator
<code>xformula</code>	Optional one-sided formula for additional covariates (e.g., $\sim \text{age} + \text{education}$)
<code>t</code>	Post-treatment time period
<code>tmin1</code>	Pre-treatment time period
<code>tname</code>	Name of the column containing the time variable
<code>data</code>	Data frame containing all variables
<code>panel</code>	TRUE if the dataset is panel data
<code>probs</code>	Vector of quantile levels at which to estimate the QTET
<code>iters</code>	Number of bootstrap iterations for standard errors

The function returns a `QTE` object with QTET estimates and (optionally) bootstrap confidence intervals at each quantile in `probs`.

4.1.4 Example

The following example from Callaway (2024) estimates the QTET for the National Supported Work Demonstration job-training program using the `lalonge.psid.panel` dataset, conditioning on several pre-treatment characteristics:

```
library(qte)
data(lalonge)

c1 <- CiC(
  re ~ treat,
  t = 1978, tmin1 = 1975, tname = "year",
  xformula = ~ age + I(age^2) + education + black + hispanic + married + nodegree,
  data = lalonge.psid.panel, idname = "id",
  se = FALSE, probs = seq(0.05, 0.95, 0.05)
)

summary(c1)
#>
#> Quantile Treatment Effect:
#>
#> tau   QTE
#> 0.05    0.00
#> 0.1     0.00
#> 0.15    0.00
#> 0.2     0.00
#> 0.25  485.23
#> 0.3    929.88
#> 0.35 1460.36
#> 0.4   2321.11
```

```
#> 0.45 3462.56
#> 0.5  4232.31
#> 0.55 5010.34
#> 0.6  6210.67
#> 0.65 7458.11
#> 0.7  7508.97
#> 0.75 8210.84
#> 0.8  8332.69
#> 0.85 8515.98
#> 0.9  8822.00
#> 0.95 8273.74
#>
#> Average Treatment Effect:    4629.24
```

The resulting QTET estimates suggest that the treatment had no measurable effect on the lower quantiles of the earnings distribution but large positive effects at higher quantiles, with an estimated average treatment effect of roughly \$4,600 (Callaway 2024).

5 Further resources

5.1 The Catalog of Bias

The Catalog of Bias¹ (Catalog of Bias Collaboration 2024a) is a collaborative project mapping all the biases that affect health evidence. The catalog provides:

- Clear definitions and explanations of various biases
- Real-world examples from health research
- Directed acyclic graphs (DAGs) illustrating bias mechanisms
- Guidance on study design and analysis to prevent or minimize bias

The catalog is an invaluable resource for researchers, clinicians, policymakers, journalists, and the public who seek to better understand threats to research validity.

¹<https://catalogofbias.org/>

6 Summary

In summary, this book has no content whatsoever.

```
1 + 1  
#> [1] 2
```

References

- Athey, Susan, and Guido Imbens. 2006. “Identification and Inference in Nonlinear Difference-in-Differences Models.” *Econometrica* 74 (2): 431–97. <https://doi.org/10.1111/j.1468-0262.2006.00668.x>.
- Banack, Hailey R., and Jay S. Kaufman. 2014. “The Obesity Paradox: Understanding the Effect of Obesity on Mortality Among Individuals with Cardiovascular Disease.” *Preventive Medicine* 62: 96–102. <https://doi.org/10.1016/j.ypmed.2014.02.001>.
- Callaway, Brantly. 2024. *CiC: Change in Changes Estimator for Quantile Treatment Effects*. <https://bcallaway11.github.io/qte/reference/CiC.html>.
- Catalog of Bias Collaboration. 2024a. *Catalog of Bias*. <https://catalogofbias.org/>.
- Catalog of Bias Collaboration. 2024b. *Collider Bias*. <https://catalogofbias.org/biases/collider-bias/>.
- “Continuation of Annual Screening Mammography and Breast Cancer Mortality in Women Older Than 70 Years.” 2020. *Annals of Internal Medicine* 172 (6): 381–89. <https://doi.org/10.7326/M18-1199>.
- Fox, Matthew P., Richard F. MacLehose, and Timothy L. Lash. 2022. *Applying Quantitative Bias Analysis to Epidemiologic Data*. Second edition. Statistics for Biology and Health. Springer. <https://doi.org/10.1007/978-3-030-82673-4>.
- Hernán, MA, and J Robins. 2020. “Causal Inference: What If. Boca Raton: Chapman & Hill/Crc.(2020).” *Publisher’s Note Springer Nature Remains Neutral with Regard to Jurisdictional Claims in Published Maps and Institutional Affiliations*. <https://miguelhernan.org/whatifbook>.
- Holland, Paul W. 1986. “Statistics and Causal Inference.” *Journal of the American Statistical Association* 81 (396): 945–60. <https://doi.org/10.1080/01621459.1986.10478354>.
- Judea Pearl, Nicholas P. Jewell, Madelyn Glymour. 2016. *Causal Inference in Statistics: A Primer*. 1st ed. Wiley.
- Lash, Timothy L., Matthew P. Fox, and Aliza K. Fink. 2009. *Applying Quantitative Bias Analysis to Epidemiologic Data*. 1st ed. Statistics for Biology and Health. Springer. <https://doi.org/10.1007/978-0-387-87959-8>.
- Sackett, D. L. 1979. “Bias in Analytic Research.” *Journal of Chronic Diseases* 32 (1-2): 51–63. [https://doi.org/10.1016/0021-9681\(79\)90012-2](https://doi.org/10.1016/0021-9681(79)90012-2).
- Schneeweiss, Sebastian. 2006. “Sensitivity Analysis and External Adjustment for Unmea-

sured Confounders in Epidemiologic Database Studies of Therapeutics.” *Pharmacoepidemiology and Drug Safety* 15 (5): 291–303. <https://doi.org/https://doi.org/10.1002/pds.1200>.