Paper AS06

A Practical Approach to Causal Inference

Laura Watson Smith, SAS Institute, Cary, United States
Sherrine Eid, SAS Institute, Cary, United States

ABSTRACT

Any general introductory statistics course aims to leave students with one of the most important and foundational takeaways in the field: correlation does not imply causation. This leaves us with a rather glaring and consequential question: how can we infer causation? Such a question may be easiest to understand through a worked example.

This paper outlines a practical framework for performing causal analysis on observational data, guiding users through confounder identification, model specification, and effect estimation. As a motivating example, we leverage this framework against real-world data from the U.S. National Health and Nutrition Examination Survey (NHANES) to estimate the causal effect of smoking cessation on long-term weight change.

INTRODUCTION

PUBLISHED EXAMPLE

Consider the question, “Does smoking cessation cause weight change?” Framed as an associational question, this becomes, “Are smoking cessation and weight change associated?” In contrast, a causal inference approach asks, “Would the same weight change occur in the absence of smoking cessation as was observed when individuals quit smoking?” Smoking is a major public health concern, contributing to many chronic diseases and long-term changes in body weight.[1-4] Being overweight or obese is linked to increased chronic health risks and higher all-cause mortality.[5-7] Understanding how smoking cessation affects weight change is therefore important for public health research and for informing clinical guidelines. This study uses observational, cross-sectional data from the US National Health and Nutrition Examination Survey (NHANES) to estimate the causal effect of smoking cessation on weight change while adjusting for confounding factors, following the framework outlined in the textbook *What If* by Hernán and Robins.

ORIENTATION FOR CAUSAL INFERENCE

Causality describes a relationship wherein one event, the cause, directly influences another, the outcome, such that the outcome would not have occurred without the cause.[8] Causal inference is the process of determining whether such a relationship exists between variables in the general population and quantifying its strength.[8]

WHAT IS THE DIFFERENCE BETWEEN CAUSAL INFERENCE AND STATISTICAL ANALYSIS?

QUESTIONS

Statistical analysis and causal inference address different questions. Statistical analysis focuses on observed associations, examining the joint distribution of variables to answer associational questions.[9,10] In contrast, causal inference tackles counterfactual “what-if” scenarios, asking how outcomes would change under alternative conditions or interventions.[9] Because the same data distribution can align with multiple causal structures, answering causal questions requires ***knowledge of the underlying data-generating process, formalized through constructing a causal model and asserting causal assumptions.***[9] The key distinction is that while statistical analysis describes observed associations, causal inference explains what would have occurred under different, hypothetical circumstances.

POPULATIONS

Figure 1 taken from *What If* by Hernán and Robins, helps to illustrate the difference in considered populations between causal inference and statistical analysis in a general cross-sectional example.[11] Consider the question: *“Is the risk of outcome the same for treated and untreated individuals?”,* where Y denotes outcome and A defines treatment assignment.

Statistical analysis and causal inference differ also in the populations they conceptually compare. Statistical analysis divides the observed sample into treated and untreated groups, then compares average outcomes between these real individuals—Person X versus Person Z—using the data we have.[11,12] In contrast, causal inference imagines the same population under two hypothetical scenarios: everyone treated and everyone untreated.[11] Conceptually, it

compares Person X under treatment to that same Person X had they not been treated, which requires defining potential outcomes for each individual for each exposure condition.[11] Because only one outcome is observed for each person, one potential outcome is counterfactual and unobservable, making causal inference fundamentally a problem of missing data. While statistical analysis could, in principle, observe more data through additional sampling, causal inference uses observed data and assumptions to learn about outcomes that can never be directly observed.[12]

Figure 1. Illustration of the Populations Considered in Causal vs. Associational Statistics

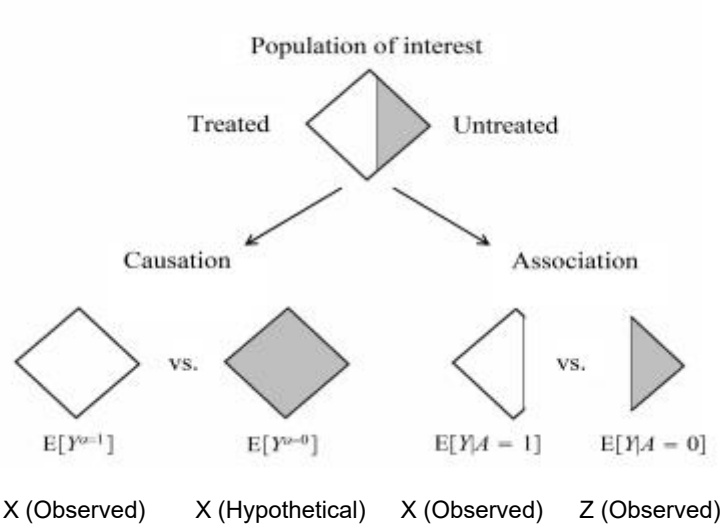

Figure 1 was captured from What If by Hernán and Robins, see reference 11.

ASSUMPTIONS

The final key distinction lies in the assumptions each approach makes. Associational assumptions are testable in principle, given a sufficiently large sample size.[9,13] These include assumptions about distributions or model structures, such as normality, linearity, and independence. Causal inference, on the other hand, requires additional assumptions beyond those of standard statistical analysis. These causal assumptions allow an observational study to be conceptualized as a *conditionally randomized experiment*.[11]

WHY DO WE DO CAUSAL INFERENCE INSTEAD OF STATISTICAL ANALYSIS?

We use causal inference when we want to understand the effect of an intervention, not just whether variables are related. Statistical analysis only identifies associations, while causal inference asks what would happen if conditions changed—making it more actionable for research, policy implementation, and medicine than statistical analysis.[14,15] For example, “Smoking cessation is associated with weight change” describes correlation, but a causal claim like “Smoking cessation causes a 3-pound weight increase compared to not quitting” explains the impact of intervention and supports decision-making.

WHY IS CAUSAL INFERENCE DIFFICULT?

UNOBSERVABLE DATA

Causal inference is often described as a problem of missing data. As previously noted, it conceptually compares the same population under two hypothetical exposure conditions. This framework yields two potential outcomes for each individual: one observed, and the other—the counterfactual—unobservable. This is the fundamental problem of causal inference: we can never observe both potential outcomes for the same individual at the same time.[16]

ASSOCIATION ^= CAUSATION

Mentioned earlier, traditional statistical analysis focuses on associational questions, whereas causal inference addresses counterfactual questions. This distinction leads directly to a familiar problem: association is not the same as causation. The presence of an association alone is insufficient to establish a causal relationship. Association and causation are defined with respect to different underlying populations, and an observed association may reflect phenomena other than a causal relationship, such as bias or confounding.[16-21]

TO IMPLY CAUSATION

Establishing causality requires conditions beyond an observed association between two variables.[17,20,21] Satisfying these conditions amounts to making appropriate assumptions, including identifying assumptions that imply the causal effect is measurable from observed data. This process of identification will be elaborated on in the next section.

COMMON CAUSAL TASKS

DIRECTED ACYCLIC GRAPHS (DAGS)

Directed acyclic graphs (DAGs) visualize assumed relationships among variables included in the system as nodes linked through one-way directed arrows which form no loops or cycles.[11,22-31] They clarify causal pathways in complex studies, informing study design by making causal assumptions explicit.[11,22-31] By explicitly showing pathways, DAGs help identify confounders, mediators, and colliders, guiding the selection of covariates to control confounding and improve causal estimates.[11,22-31] DAGs also provide a means of visually assessing potential biases in causal inference.[24-26,30] Using a DAG helps to bridge gaps in bias assessment by identifying biases from all sources, including data selection, study design, and analysis method.[32]

IDENTIFICATION

Identification determines whether the causal estimand—such as the average treatment effect (ATE) in this case—can be expressed using observed data.[11,33] When the identifying assumptions underlying causal analysis hold, it implies that the observed effect is the only possible effect that could have arisen given the data; if identification fails, the analysis remains associational.[11,34] This process is complicated by the fundamental problem of causal inference: we cannot observe both potential outcomes for the same individual at the same time.[34] As a result, causal inference operates at the population level, comparing averages rather than individual outcomes.[11] In causal identification, the goal is to prove that it is possible to express the causal estimand as a function of the observed data, a task complicated by the very nature of this fundamental problem. To justify that the causal estimand is identifiable from our data, certain causal identification assumptions must be asserted.[34]

- ***Positivity:*** There is a nonzero probability that individuals in the observed data receive every level of exposure for each combination of values for included covariates.[35-38]
- ***Consistency:*** The exposure is defined well enough that on the individual level, there are no multiple versions of the exposure leading to different potential outcomes.[39-41]
- ***No Interference:*** An individual's treatment status has no impact on the potential outcomes for any other individual.[37,41]
- ***Exchangeability:*** For all levels of exposure, the potential outcomes are independent of the exposure assignment conditional on the observed covariates.[37,42]

Taken together, these assumptions allow us to assert that the exposed and unexposed groups are identical in all respects, whether observable or unobservable, except for their treatment assignments.[37,41-43] This allows us to assert that, under these assumptions, our observed data identifies the causal estimand, which may be interpreted as reflecting a true causal relationship.[11,33,34]

EFFECT ESTIMATION

Effect estimation quantifies the impact of an exposure on an outcome across a population. After consulting a casual diagram, and once the causal parameter is identified, statistical methods are used to estimate its magnitude, providing insight into how changes in treatment would influence outcomes at the population level.

EXAMPLE

This example demonstrates effect estimation using SAS Viya to assess whether quitting smoking causes weight change, while appropriately adjusting for confounders. It is implemented as an interactive dashboard, leveraging native SAS procedure code through embedded SAS Jobs.

CONSTRUCT A CAUSAL DIAGRAM

Causal analyses begin by constructing a DAG to visualize the relationships among variables, particularly how they relate to the exposure and outcome. This helps identify confounders that must be controlled to obtain unbiased effect estimates. A DAG can be created manually or generated using SAS's PROC CAUSALDISCOVERY. In the Effect Estimation dashboard (**Figure 2**), an embedded SAS Job runs PROC CAUSALDISCOVERY, allowing users to specify exposure and outcome variables and producing outputs—such as the Estimated Adjacency Matrix—that describe the DAG's parent–child structure (

Figure *3*).

Figure *2* Effect Estimation Dashboard: Proc CausalDiscovery Page

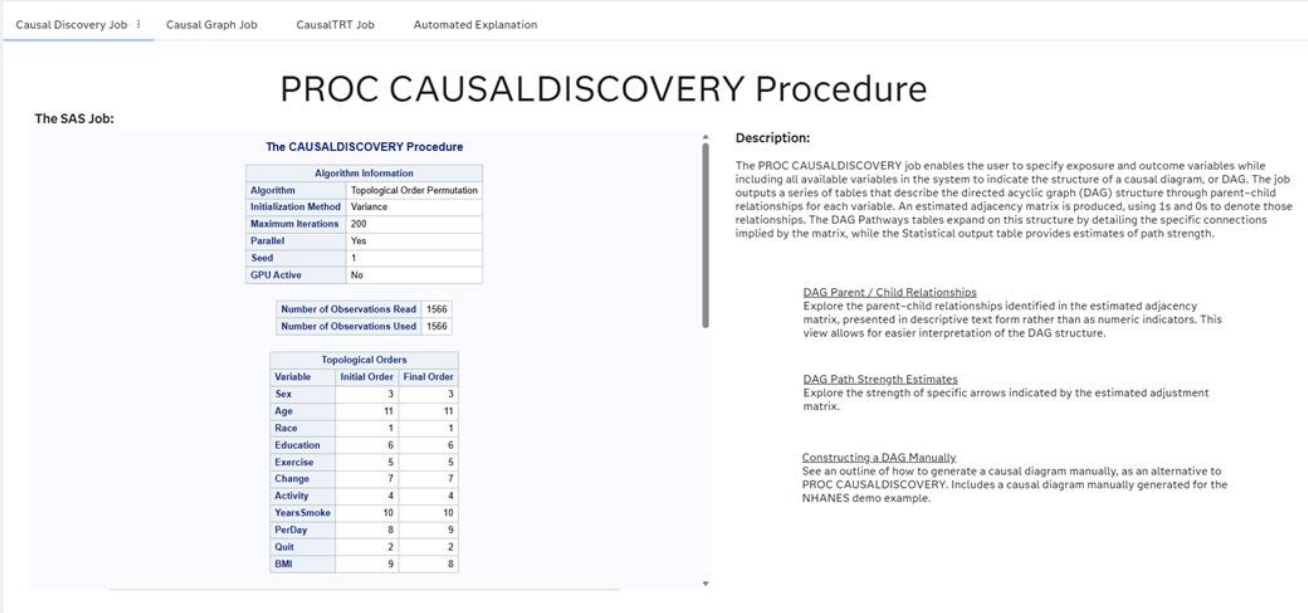

Figure 3 PROC CAUSALDISCOVERY Output: Estimated Adjacency Matrix

Estimated Adjacency Matrices

Alpha	Variable	Sex	Age	Race	Education	Exercise	Change	Activity	YearsSmoke	PerDay	Quit	BMI
1E-6	Sex	0	1	0	1	1	0	1	0	1	0	1
	Age	0	0	0	0	0	0	0	0	0	0	0
	Race	1	0	0	0	1	0	0	0	1	1	1
	Education	0	1	0	0	0	1	0	1	1	0	1
	Exercise	0	0	0	1	0	0	0	1	0	0	1
	Change	0	0	0	0	0	0	0	0	0	0	0
	Activity	0	0	0	1	1	0	0	0	0	0	1
	YearsSmoke	0	1	0	0	0	0	0	0	0	0	0
	PerDay	0	0	0	0	0	0	0	1	0	0	0
	Quit	1	1	0	1	1	1	1	0	1	0	1
	BMI	0	1	0	0	0	0	0	1	1	0	0

PROC CAUSALDISCOVERY helps identify the structure of a causal diagram but does not generate a visual DAG, so this analysis is supplemented with a manual construction method. A detailed description of this manual approach is provided in the linked GitHub page.

The completed causal diagram (**Figure 4**) summarizes the relationships among variables in the data and can reveal structures not evident from statistical analysis alone. In the smoking-cessation example, the DAG confirmed Age, Sex, and Years Smoked as confounders and additionally identified Exercise as a confounder because it influences both quitting status and BMI, which affects weight change. A subsequent statistical review supported this finding, illustrating how DAGs help uncover complex dependencies that may be missed by traditional analyses alone.

Figure 4 Manually Constructed Causal Diagram of NHANES Example

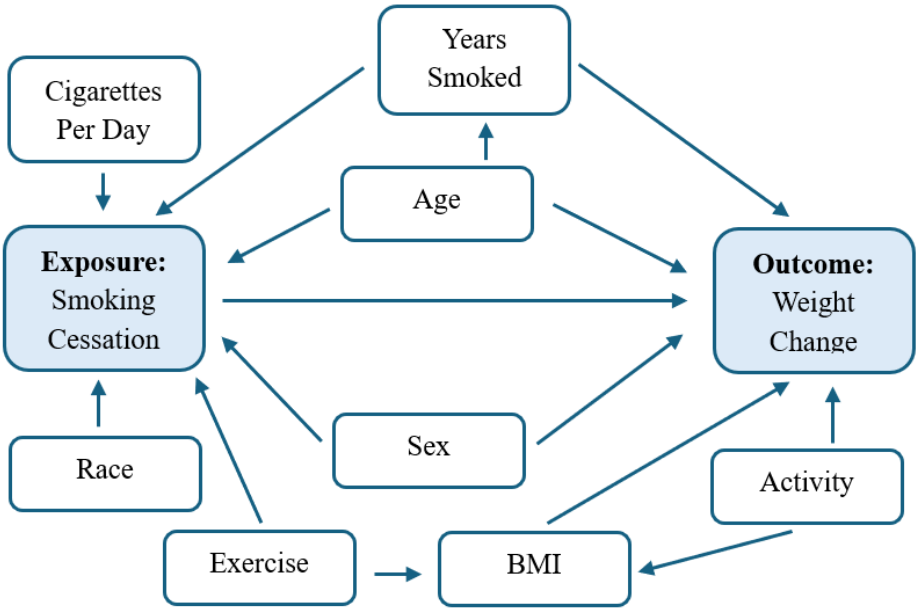

CAUSAL IDENTIFICATION

The next step, Identification, verifies whether the average treatment effect can be estimated from the observed data, moving the analysis beyond traditional associational methods into causal inference.

Figure 5 Effect Estimation Dashboard: Causal Graph Page

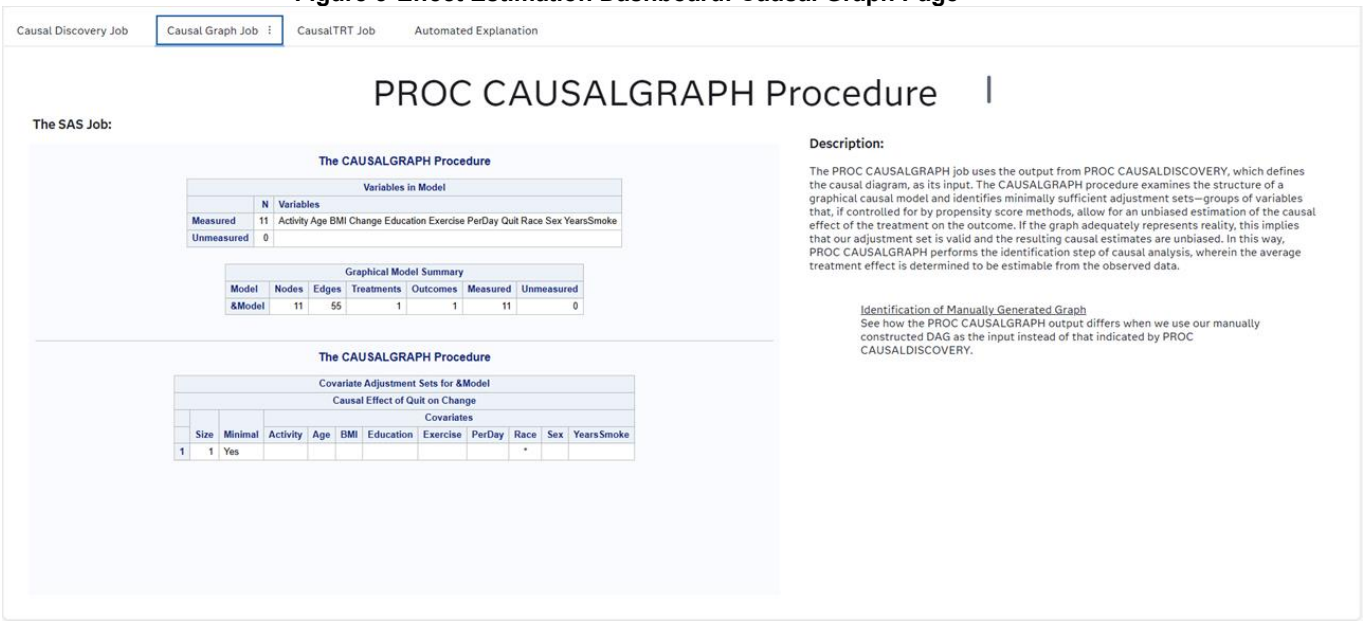

This step uses PROC CAUSALGRAPH on the second page of the Effect Estimation dashboard (**Figure 5**). The procedure takes the defined causal diagram from the PROC CAUSALDISCOVERY output, then evaluates the graph—given the user-specified exposure and outcome—to identify minimally sufficient adjustment sets for unbiased effect estimation.[44] The key output is the covariate adjustment set table (**Figure 6);** finding at least one minimally sufficient adjustment set indicates that the average treatment effect is identifiable when those variables are controlled for as confounders.

Figure 6 PROC CAUSALGRAPH Output: Covariate Adjustment Sets

Covariate Adjustment Sets for TESTID Demo										
Causal Effect of SmokingCessation on WeightChange										
			Covariates							
	Size	Minimal	Activity	Age	BMI	Exercise	PerDay	Race	Sex	YearsSmoke
1	4	Yes		*	*				*	*
2	4	Yes		*		*			*	*
3	5	No	*	*	*				*	*
4	5	No	*	*		*			*	*
5	5	No		*	*	*			*	*
6	5	No		*	*		*		*	*
7	5	No		*	*			*	*	*
8	5	No		*		*	*		*	*
9	5	No		*		*		*	*	*
10	6	No	*	*	*	*			*	*
11	6	No	*	*	*		*		*	*
12	6	No	*	*	*			*	*	*
13	6	No	*	*		*	*		*	*

In the smoking cessation example, two minimally sufficient adjustment sets are reported. The second set aligns exactly with the manually selected confounders—Age, Exercise, Sex, and Years Smoked—providing additional support for that confounder selection. This establishes identification, confirming that the average treatment effect is estimable from our data.

EFFECT ESTIMATION

In the final step, the treatment effect is estimated as the difference in predicted outcomes across the population under each treatment level. This is done using the PROC CAUSALTRT procedure, which integrates a propensity score model and an outcome regression model into a pair of estimating equations to compute the average treatment effect (Error! Reference source not found.).

In the dashboard, users specify the exposure, outcome, propensity score model, and outcome model (Error! Reference source not found.**).** PROC CAUSALTRT then runs through all supported causal inference methods, producing predicted outcomes by treatment level, average treatment effect estimates, propensity scores, and related statistical results. The main output is the Analysis of Causal Effects table (default method: AIPW), which reports predicted outcomes for each treatment level, the average treatment effect, and confidence measures such as p-values and confidence intervals (**Figure 8**). A summary table in the Effect Estimation dashboard also compiles average treatment effect estimates and confidence metrics across all methods (**Figure 9**).

In the smoking-cessation example, **Figure 8** and **Figure 9** show an AIPW-estimated average treatment effect of 3.37 kg, indicating that quitting smoking leads to roughly a 3.4 kg long-term weight gain after adjusting for confounders. This provides evidence that smoking cessation causes a meaningful increase in weight change. The causal effect analysis is complete, though additional visualizations could still be explored.

Figure 7 Effect Estimation Dashboard: CausalTrt Page

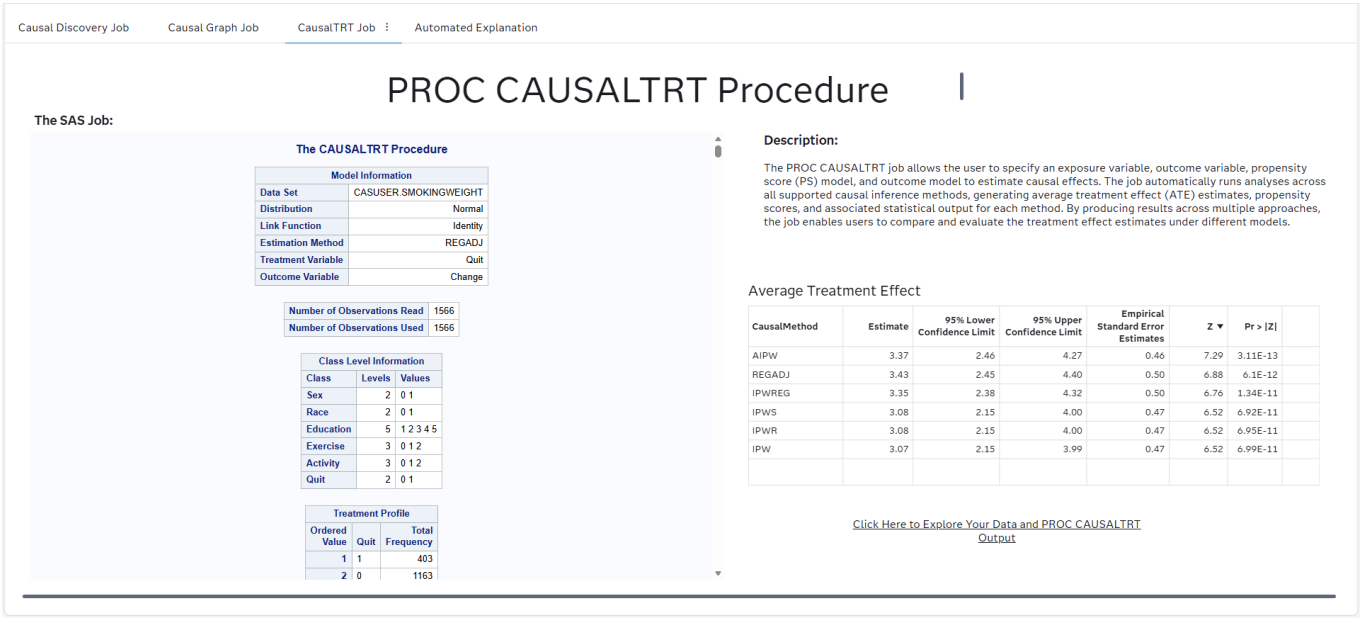

Figure 8 PROC CAUSALTRT Output: Analysis of Causal Effects AIPW

The CAUSALTRT Procedure

Analysis of Causal Effect

Parameter	Treatment Level	Estimate	Robust Std Err	Wald 95% Confidence Limits		Z	Pr > |Z|
POM	1	5.1467	0.4177	4.3282	5.9653	12.32	<.0001
POM	0	1.7782	0.2152	1.3564	2.2001	8.26	<.0001
ATE		3.3685	0.4621	2.4628	4.2742	7.29	<.0001

Figure 9 PROC CAUSALTRT Output: Analysis of Causal Effects Summary

CausalMethod	Estimate	95% Lower Confidence Limit	95% Upper Confidence Limit	Empirical Standard Error Estimates	Z ▼	Pr > |Z|
AIPW	3.37	2.46	4.27	0.46	7.29	3.11E-13
REGADJ	3.43	2.45	4.40	0.50	6.88	6.1E-12
IPWREG	3.35	2.38	4.32	0.50	6.76	1.34E-11
IPWS	3.08	2.15	4.00	0.47	6.52	6.92E-11
IPWR	3.08	2.15	4.00	0.47	6.52	6.95E-11
IPW	3.07	2.15	3.99	0.47	6.52	6.99E-11

CONCLUSION

In summary, we have introduced and explored the principles of causal inference, as well as demonstrated how causal analyses can be implemented using SAS Viya. Through the stepwise process of generating causal diagrams, establishing identification, and estimating treatment effects, we have shown how observational data can be leveraged to draw meaningful conclusions about causal relationships.

REFERENCES

1. Driva S, Korkontzelou A, Tonstad S, Tentolouris N, Katsaounou P. The Effect of Smoking Cessation on Body Weight and Other Metabolic Parameters with Focus on People with Type 2 Diabetes Mellitus. Int J Environ Res Public Health. 2022 Oct 14;19(20):13222. doi: 10.3390/ijerph192013222. PMID: 36293800; PMCID: PMC9603007.
2. Mishra A, Chaturvedi P, Datta S, Sinukumar S, Joshi P, Garg A. Harmful effects of nicotine. Indian J Med Paediatr Oncol. 2015 Jan-Mar;36(1):24-31. doi: 10.4103/0971-5851.151771. PMID: 25810571; PMCID: PMC4363846. https://pmc.ncbi.nlm.nih.gov/articles/PMC4363846/
3. Audrain-McGovern J, Benowitz NL. Cigarette smoking, nicotine, and body weight. Clin Pharmacol Ther. 2011 Jul;90(1):164-8. doi: 10.1038/clpt.2011.105. Epub 2011 Jun 1. PMID: 21633341; PMCID: PMC3195407. https://pmc.ncbi.nlm.nih.gov/articles/PMC3195407/#S3
4. Hu, Y., Zong, G., Liu, G., Wang, M., Rosner, B., Pan, A., Willett, W. C., Manson, J. E., Hu, F. B., & Sun, Q. (2018). Smoking cessation, weight change, type 2 diabetes, and mortality. New England Journal of Medicine, 379(7), 623–632. https://doi.org/10.1056/nejmoa1803626
5. National Institute of Diabetes and Digestive and Kidney Diseases. (September 2021). Overweight & Obesity Statistics. US National Institute of Health. https://www.niddk.nih.gov/health-information/health-statistics/overweight-obesity
6. National Institute of Diabetes and Digestive and Kidney Diseases. (May 2023). Health Risks of Overweight & Obesity. US National Institute of Health. https://www.niddk.nih.gov/health-information/weight-management/adult-overweight-obesity/health-risks
7. Kim HJ, Kim BS, Lee JH, Shin JH. Impact of underweight on 3-year all-cause mortality in patients with acute severe hypertension: a retrospective cohort study. Sci Rep. 2022 Mar 21;12(1):4798. doi: 10.1038/s41598-022-08892-9. PMID: 35314748; PMCID: PMC8938442. https://pmc.ncbi.nlm.nih.gov/articles/PMC8938442/
8. Garcia-Huidobro D, Michael Oakes J. Squeezing observational data for better causal inference: Methods and examples for prevention research. Int J Psychol. 2017 Apr;52(2):96-105. doi: 10.1002/ijop.12275. Epub 2016 Apr 20. PMID: 27094382; PMCID: PMC5549466.
9. Pearl, J. (2009). Caual inference in statistics: An overview. Statistical Surveys. Vol. 3 (2009): 96-146. ISSN:1935-7516. DOI: 10.1214/09-SS057.
10. Reizinger, P. (Nov 6, 2021). Pearls of Causality #4: Causal Queries. Casual Causality. Retrieved from https://rpatrik96.github.io/posts/2021/11/poc4-causal-queries/
11. Hernán, MA., Robins, JM. (2020). Causal Inference: What If. Boca Raton: Chapman & Hall/CRC.
12. Masten, M. (Oct 30, 2015). Statistical vs Causal Inference: Causal Inference Bootcamp. Mod-U: Powerful Concepts in Social Science. Youtube. Retrieved from https://www.youtube.com/watch?v=jfkysJxSifg
13. Reizinger, P. (Nov 6, 2021). Pearls of Causality #5: Statistical vs Causal Inference. Casual Causality. Retrieved from https://rpatrik96.github.io/posts/2021/11/poc5-stats-vs-causality/
14. Krohn, J., Taylor, S. (Oct 16, 2022). Causal Modeling: Why and when is it helpful? Super Data Science: ML & AI Podcast with Jon Krohn. Youtube. Retrieved from https://www.youtube.com/watch?v=F159s-Ar4vc
15. Hershaff, J. (Sept 4, 2025). Why Causal Inference Matters: From Marketing to Policy. Udacity. Youtube. Retrieved from https://www.youtube.com/watch?v=h0rHctkP7fQ
16. Lee, H., Aaronson, J., Nunan, D. (2019) Association or causation? How do we know? Catalogue of Bias, Nuffield Department of Primary Care. Retrieved from https://catalogofbias.org/2019/03/05/association-or-causation-how-do-we-ever-know/
17. Aandahl, EM. (). The difference between association, correlation, and causation. Ledidi. Retrieved from https://ledidi.com/academy/the-difference-between-association-correlation-and-causation
18. Hernán MA. A definition of causal effect for epidemiological research. J Epidemiol Community Health. 2004 Apr;58(4):265-71. doi: 10.1136/jech.2002.006361. PMID: 15026432; PMCID: PMC1732737.
19. McCormick, K. (March 18 2022). Machine Learning and AI Foundations: Prediction, Causation, and Statistical Inference. LinkedIn Career Hub. Retrieved from https://www.linkedin.com/learning/machine-learning-and-ai-foundations-prediction-causation-and-statistical-inference/prediction-causation-and-statistical-inference?dApp=218089040&leis=LAA&u=2101305
20. Irizarry, R. () Introduction to Data Science: Linear Models: Association is Not Causaltion. Github. Retrieved from https://rafalab.dfci.harvard.edu/dsbook-part-2/linear-models/association-not-causation.html
21. Barrat, H., Kirwan, M., Shantikuman, S. (2018). Association and Causation. Faculty of Public Health. Retrieved from https://www.healthknowledge.org.uk/public-health-textbook/research-methods/1a-epidemiology/association-causation
22. Faries, D., Obenchain, R., Haro, J., & Leon, A. (2010). Analysis of Observational Health Care Data Using SAS. Sas Institute.
23. Faries, D. E., Zhang, X., Kadziola, Z., Siebert, U., Kuehne, F., Obenchain, R. L., & Haro, J. M. (2020). Real World Health Care Data Analysis: Causal methods and implementation using SAS. SAS Institute.
24. Primbs, M., Bijlstra, G., Holland, R., & Thoemmes, F. (2024, August 1). Causal Inference for Dummies: A Tutorial on Directed Acyclic Graphs and Balancing Weights.
25. Darren Dahly. (September 25, 2019). Intro to causal inference and directed acyclic graphs.
26. CodeEmporium. (January 3, 2022). Causal Inference – EXPLAINED!
27. 21. Epidemiology Stuff. (March 1, 2022). Directed Acyclic Graphs (DAGs).
28. Pearl, J. (2010). An Introduction to Causal Inference. The International Journal of Biostatistics.

29. Wang SV, Schneeweiss S. (January 2022). Assessing and Interpreting Real-World Evidence Studies: Introductory Points for New Reviewers. Clin Pharmacol Ther.
30. Schneeweiss S, Rassen JA, Brown JS, Rothman KJ, Happe L, Arlett P, Dal Pan G, Goettsch W, Murk W, Wang SV. (2019 Mar 19). Graphical Depiction of Longitudinal Study Designs in Health Care Databases. Ann Intern Med.
31. Emmott, N., Nafie, M., Kim, D., & Hendricks-Sturrup, R. (June 18, 2024). Real-World Evidence to Support Causal Inference: Methodological Considerations for Non-Interventional Studies. Duke-Margolis Institute for Health Policy.
32. SAS Help Center: Causal Inference Overview
33. Kiciman, E., Sharma, A. (March 29, 2021). Getting Started With Causal Inference; Chapter 3: Identification. Causalinference.gitlab.io. Retrieved from https://causalinference.gitlab.io/causal-reasoning-book-chapter3/
34. Unal, M. (February 22, 2023). Identification: The Key to Credible Causal Inference. Towards Data Science. Retrieved from https://towardsdatascience.com/identification-the-key-to-credible-causal-inference-c3023143349e/
35. Cole, Stephen R.a; Frangakis, Constantine E.b. The Consistency Statement in Causal Inference: A Definition or an Assumption?. Epidemiology 20(1):p 3-5, January 2009. | DOI: 10.1097/EDE.0b013e31818ef366
36. Murray, E. (January 26, 2019). Positivity: What it is and why it matters for data science. Medium. Retrieved from https://medium.com/@EpiEllie/positivity-what-it-is-and-why-it-matters-for-data-science-d5e9c0bc1fcb
37. Didelez, V., Evans, R. (November 26, 2025). Causal Inference; Chapter 9: Causal Assumptions. Stats.ox.ac.uk. Retrieved from https://www.stats.ox.ac.uk/~evans/APTS/causassmp.html
38. Neal, B. (September 7, 2020). Positivity/Overlap and Extrapolation. Youtube. Retrieved from https://www.youtube.com/watch?v=4xc8VkrF98w&list=PLoazKTcS0Rzb6bb9L508cyJ1z-U9iWkA0&index=15
39. Neal, B. (September 7, 2020). No Interference and Consistency. Youtube. Retrieved from https://www.youtube.com/watch?v=3t4HPI8Gmto
40. Rehkopf DH, Glymour MM, Osypuk TL. The Consistency Assumption for Causal Inference in Social Epidemiology: When a Rose is Not a Rose. Curr Epidemiol Rep. 2016 Mar;3(1):63-71. doi: 10.1007/s40471-016-0069-5. Epub 2016 Feb 16. PMID: 27326386; PMCID: PMC4912021.
41. Perry, R. (January 19, 2020). SUTVA vs. Exchangeability. Github. Retrieved from https://rflperry.github.io/posts/sutva_vs_exchangeability/
42. Neal, B. (September 7, 2020). Conditional Exchangeability and the Adjustment Formula. Youtube. Retrieved from https://www.youtube.com/watch?v=FFjL5Hkeap4&list=PLoazKTcS0Rzb6bb9L508cyJ1z-U9iWkA0&index=14
43. Unal, M. (March 7, 2023). Why are Randomized Experiments the Gold Standard in Causal Inference? Towards Data Science. Retrieved from https://towardsdatascience.com/why-are-randomized-experiments-the-gold-standard-in-causal-inference-f3fa240a1d02/
44. SAS Institute Inc. 2018. SAS/STAT® 15.1 User's Guide: The CAUSALGRAPH Procedure. Cary, NC: SAS Institute Inc. https://support.sas.com/documentation/onlinedoc/stat/151/causalgraph.pdf

ACKNOWLEDGMENTS

The authors would like to thank Clay Thompson and Joe Zhao, subject matter experts who were consulted in that capacity and provided guidance to ensure the technical correctness and clarity of this paper. The authors also thank Jim Box for editorial support and mentorship throughout the development of the paper.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Author Name: Laura Watson
Company: SAS Institute, Inc
Address: 100 SAS Campus Drive, Cary, NC27513
Work Phone: 919.531.4894
Email: Laura.Watson@sas.com
Website: www.sas.com