

Commentary

Selection Bias Requires Selection: The Case of Collider Stratification Bias

Haidong Lu, Gregg S. Gonsalves, and Daniel Westreich*

* Correspondence to Daniel Westreich, Department of Epidemiology, UNC Gillings School of Global Public Health, 105C McGavran Greenberg Hall CB #7435, Chapel Hill, NC 27599 (e-mail: westreich@email.unc.edu).

Initially submitted March 8, 2023; accepted for publication October 27, 2023.

In epidemiology, collider stratification bias, the bias resulting from conditioning on a common effect of two causes, is oftentimes considered a type of selection bias, regardless of the conditioning methods employed. In this commentary, we distinguish between two types of collider stratification bias: collider restriction bias due to restricting to one level of a collider (or a descendant of a collider) and collider adjustment bias through inclusion of a collider (or a descendant of a collider) in a regression model. We argue that categorizing collider adjustment bias as a form of selection bias may lead to semantic confusion, as adjustment for a collider in a regression model does not involve selecting a sample for analysis. Instead, we propose that collider adjustment bias can be better viewed as a type of overadjustment bias. We further provide two distinct causal diagram structures to distinguish collider restriction bias and collider adjustment bias. We hope that such a terminological distinction can facilitate easier and clearer communication.

causal diagrams; collider adjustment bias; collider stratification bias; epidemiologic research; overadjustment bias; selection bias

In epidemiology, “collider stratification bias” is a broad term describing the bias resulting from conditioning on a common effect of two causes, of which one is the exposure or a cause of the exposure, and the other is the outcome or a cause of the outcome (1–4). The term “conditioning” refers in common usage variously to adjustment (as in a regression analysis or tabular, Mantel-Haenszel-type analysis, both cases in which such an approach is also called “controlling” for variables) and to restriction (e.g., in a 2-level variable, only performing analysis in one level of that variable; or alternatively, performing analysis in each level of that variable, separately). Note that this latter approach, looking at the effect of an exposure on an outcome in each level of a covariate, is sometimes referred to as “stratification.” This can lead to semantic confusion since that same word, used in the term “collider stratification bias,” means something more general. Here, we think of stratification as a kind of restriction applied to each level of the variable in parallel.

Hernán et al. (2) considered collider stratification bias to be a type of selection bias, regardless of conditioning methods employed. Here we distinguish between two types of collider stratification bias: collider restriction bias due to restricting to one level of a collider (or a descendant of a

collider, including restricting simultaneously to all levels of a collider and reporting the results separately) and collider adjustment bias through inclusion of a collider (or a descendant of a collider) in a regression model. We contend that while collider restriction bias can be deemed to be a form of selection bias, categorizing collider adjustment bias as selection bias may lead to a definitional error. Instead, it may be more useful and intuitive to think of collider adjustment bias as a type of overadjustment bias (5, 6).

Selection bias is broadly defined as “distortions that result from procedures used to select subjects and from factors that influence participation in the study” in *A Dictionary of Epidemiology* (7, p. 225), and the word “select” is defined as “to choose (as by fitness or excellence) from a number or group: pick out” in the Merriam-Webster Online Dictionary (8). Both terms imply that when we select, we are taking some (research subjects, observations, etc.) and leaving others—and thus that selection relates to a change in the sample analyzed, whether that change comes from the selection of target population into study sample, or from study sample into analytical sample (4). This accords with common or colloquial usage: If a friend asks you to go to the store and select some apples, they (generally speaking) expect you to bring home 3 or 4 apples, to take some apples, and leave




Figure 1. The proposed relationship between selection bias, collider stratification bias, collider restriction bias, and collider adjustment bias. Note that collider restriction bias is also referred to as type 1 selection bias, whereas type 2 selection bias is due to restricting to one level of an effect measure modifier (4). In contrast, collider adjustment bias can be deemed to be a type of overadjustment bias.

the rest. Indeed, even the paper by Hernán et al. (which argues that all collider stratification bias, including collider adjustment bias, is selection bias) states that “the common consequence of selection bias is that the association between exposure and outcome among those selected for analysis differs from the association among those eligible” (2, p. 615), implying similarly that selection bias is intimately tied to selection into analysis, and that the sample analyzed changes from those eligible for the study to those included into the analytical sample due to selection.

Collider restriction bias, analogous to restriction as a method for minimizing confounding, involves bias due to the selection of a portion of a study sample, and is straightforwardly a type of selection bias. Collider adjustment bias, however, is less clear. Although simple inclusion of a collider in a regression model does indeed cause bias, it does not result in a change in the sample being analyzed. In collider adjustment, analogous to adjustment for confounding, the entire sample is used for analysis: There is no sample selection involved. Or to echo Hernán et al.’s language (2), adjustment for a collider in a regression model does not result in sample “select[ion] for analysis.” For this reason, we propose that collider adjustment bias might be better categorized as a form of overadjustment bias.

Typically, we think of adjusting for variables as reducing bias: When adjusting for a variable creates bias, we refer to that bias as overadjustment bias (5, 6). Generally, there are 4 types of overadjustment bias. Apart from 1) adjustment for a collider (or a descendant of a collider), overadjustment bias can also stem from 2) adjustment for a mediator (or a descendant of a mediator) when estimating a total effect, 3) adjustment for an instrumental variable when unmeasured confounding is present, and 4) adjustment for a descendant of the outcome when the exposure has a causal effect on the outcome (specifically, in the absence of a sharp causal null) and the effect measure is not the odds ratio (this type of overadjustment bias can also be considered to be a special type of collider bias, as detailed in the work of Lu et al. (5)).




Figure 2. Examples of two distinct causal diagrams for collider restriction bias and collider stratification bias. Consider a randomized trial, A is certain antiretroviral therapy, Y is the 5-year mortality, L is the presence of human immunodeficiency virus–related symptoms, U is the unmeasured common cause for variable L and outcome Y (e.g., levels of immunosuppression), and $S = 1$ means selection into the sample (i.e., being uncensored during the study follow-up in this context). A) An example of collider restriction bias (a type of selection bias) where only those who were selected (i.e., $S = 1$) are included in the sample being analyzed; B) an example of collider stratification bias through inclusion of a collider in the regression model.

Unless we are willing to consider (for example) adjustment for a mediator to be a kind of selection bias, we find ourselves at a definitional impasse. Specifically, selection bias requires, denotatively and connotatively, a restriction, but adjustment alone does not restrict. As such, the logical conclusion is that adjustment for a collider alone cannot be a type of selection bias.

To be clear, a regression model that adjusts for a collider typically yields a biased estimate of effect in the stratum-specific associations, which can be further compounded by pooling the estimates across levels of the collider. Hence, the effect estimate from regression with adjustment for a collider is not generally equal to the true causal effect we are interested in. But it is not a selection bias, because there is no selection—no taking of some observations, leaving others. Rather, adjusting for a collider in a regression without selecting the sample being analyzed can be better understood as a type of overadjustment bias (4, 6) or causal model misspecification bias (9), because a collider is the variable we do not wish to adjust for, and mistakenly adjusting for it introduces bias. The view described here on the relationship between selection bias, collider stratification bias, and collider restriction bias, and collider adjustment bias is depicted in Figure 1.

Why is this distinction important? Because to most people learning epidemiology in English, the plain meaning of “selection” is taking some observations and leaving others. In our experience teaching epidemiology, including collider adjustment bias in this category leads to semantic, and thus methodological, confusion.

To better distinguish collider restriction bias (as a form of selection bias) and collider adjustment bias (as a form of overadjustment bias), we advocate for two distinct structures in causal diagrams to illustrate these two biases (see Figure 2). Briefly, it is suggested that a selection node in causal diagrams be used to describe the action of “selection” (10, 11). For example, in Figure 2A as an example of collider restriction bias, the selection node S , is the descendant of the collider L . Restricting to one level of S (i.e., $S = 1$) does result in a change in the sample analyzed, and can lead to selection bias. Without the node $S = 1$, the causal diagram

represents the entire population (i.e., those eligible, as in Hernán et al.'s work (2)); by including the node $S = 1$, the sample being analyzed is now those who were selected (i.e., those selected for analysis, as in Hernán et al.'s work (2)). $S = 1$ does not necessarily indicate that one whole group at a certain level of L (e.g., $L = 1$) is selected, and the other groups at different levels of L (e.g., $L = 0$) are not selected; it can also reflect a real-world scenario, for example, that 70% of people with $L = 1$ and 20% of people with $L = 0$ are selected disproportionately, and thereby we say selection S is dependent on L (i.e., L is a cause of S). In contrast, Figure 2B represents an example of collider adjustment bias that might occur by mistakenly including the collider L in a regression model (as indicated by placing a box around L). This causal diagram neither implies any change of the sample under study nor indicates any action of "selection for analysis." Therefore, to illustrate selection bias, one could add the $S = 1$ node into causal diagrams, an approach advocated around missing data (10) and being taken in the emerging literature on generalizability. Indeed, as Greenland suggested (11), every causal diagram in epidemiologic research should include the $S = 1$ node to symbolize sample selection and better depict the data-generating process in realistic scenarios.

In conclusion, mistakenly adjusting for a collider in a regression analysis can introduce bias, just as can restricting on a collider in such an analysis. It should be noted that propensity score weighting and matching are often viewed as "adjustment" methods as well, and including a collider in the propensity score models can also result in collider adjustment bias. Our argument here is not about whether each introduces bias; instead, we are addressing whether we should restrict (pun intended) the use of the term "selection bias" to situations in which we take some observations into analysis and leave others out. Some might argue that every adjustment approach inherently "selects" distinct analytical samples from the study population, followed by pooling the data across those samples, or that choosing which variables to adjust for involves "picking out" from a set of options and thereby should be regarded as sources of selection bias. However, if this were true, then any type of overadjustment as mentioned (e.g., adjusting for a mediator) could be deemed to be a source of selection bias, consequently obfuscating the definition of selection bias and blurring the demarcation between selection bias and other biases. Hence we believe such a terminological distinction facilitates easier and clearer communication, as well as faster understanding of these ideas among learners. Thus, we propose that collider adjustment bias, in particular and as distinct from collider restriction bias, be henceforth considered a form of overadjustment bias.

ACKNOWLEDGMENTS

Author affiliations: Section of General Internal Medicine, Department of Internal Medicine, Yale School of

Medicine, New Haven, Connecticut, United States (Haidong Lu); Public Health Modeling Unit, Yale School of Public Health, New Haven, Connecticut, United States (Haidong Lu, Gregg S. Gonsalves); Program in Addiction Medicine, Yale School of Medicine, New Haven, Connecticut, United States (Haidong Lu, Gregg S. Gonsalves); Department of Epidemiology of Microbial Diseases, Yale School of Public Health, New Haven, Connecticut, United States (Gregg S. Gonsalves); and Department of Epidemiology, UNC Gillings School of Global Public Health, Chapel Hill, North Carolina, United States (Daniel Westreich).

This work was supported by grants from National Institutes of Health (award numbers DP2DA049282 and K99DA057487).

We are grateful to Dr. Stephen Cole and Dr. Chanelle Howe for helpful comments.

The content of this work was presented at the Society for Epidemiologic Research annual meeting, June 13–16, 2023, Portland, Oregon.

The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Conflict of interest: none declared.

REFERENCES

- Greenland S. Quantifying biases in causal models: classical confounding vs collider-stratification bias. *Epidemiology*. 2003;14(3):300–306.
- Hernán MA, Hernandez-Diaz S, Robins JM. A structural approach to selection bias. *Epidemiology*. 2004;15(5):615–625.
- Cole SR, Platt RW, Schisterman EF, et al. Illustrating bias due to conditioning on a collider. *Int J Epidemiol*. 2010;39(2):417–420.
- Lu H, Cole SR, Howe CJ, et al. Toward a clearer definition of selection bias when estimating causal effects. *Epidemiology*. 2022;33(5):699–706.
- Lu H, Cole SR, Platt RW, et al. Revisiting overadjustment bias. *Epidemiology*. 2021;32(5):e22–e23.
- Schisterman EF, Cole SR, Platt RW. Overadjustment bias and unnecessary adjustment in epidemiologic studies. *Epidemiology*. 2009;20(4):488–495.
- Porta M. *A Dictionary of Epidemiology*. 5th ed. Oxford, UK: Oxford University Press; 2008.
- Merriam-Webster Dictionary. Select. <https://www.merriam-webster.com/dictionary/select>. Accessed August 29, 2022.
- Keil AP, Mooney SJ, Jonsson Funk M, et al. Resolving an apparent paradox in doubly robust estimators. *Am J Epidemiol*. 2018;187(4):891–892.
- Westreich D. Berkson's bias, selection bias, and missing data. *Epidemiology*. 2012;23(1):159–164.
- Greenland S. The causal foundations of applied probability and statistics. In: Geffner H, Dechter R, Halpern JY, eds. *Probabilistic and Causal Inference: The Works of Judea Pearl*. New York, NY: Association for Computing Machinery; 2022:605–624.