

A Crash Course in Good and Bad Controls

Carlos Cinelli, Andrew Forney and Judea Pearl

March 21, 2020

Abstract

Many students of economics express frustration with the way a problem known as “bad control” is evaded, if not mishandled in the standard econometric literature. This technical report updates a previous note, which describes graphical tools for understanding and resolving the problem.¹ By making this report more accessible we hope to make those tools available to a broader community of scientists concerned with the interpretation of data.

Introduction

If you were trained in traditional regression pedagogy, chances are that you have heard about the problem of “bad controls” (Angrist and Pischke, 2009, 2014). The problem arises when one needs to decide whether the addition of a variable to a regression equation helps getting estimates closer to the parameter of interest. Analysts have long known that some variables, when added to the regression equation, can produce unintended discrepancies between the regression coefficient and the effect that the coefficient is expected to represent. Such variables have become known as “bad controls,” to be distinguished from “good controls” (also known as “confounders” or “deconfounders”) which are variables that must be added to the regression equation to eliminate what came to be known as “omitted variable bias.”

Recent advances in graphical models have produced a simple criterion to distinguish “good” from “bad” controls, and the purpose of this note is to provide practicing analysts a concise and visible summary of this criterion through illustrative examples. We will assume that readers are familiar with the notions of “path-blocking” (or d-separation) and back-door paths. For a gentle introduction, see Pearl (2009a, Sec. 11.1.2).

In the following set of models, the target of the analysis is the average causal effect (ACE) of a treatment X on an outcome Y , which stands for the

¹First posted in <https://ucla.in/2ZcRpRq>, August 14, 2019.

expected increase of Y per unit of a *controlled* increase in X . Observed variables will be designated by black dots and unobserved variables by white empty circles. Variable Z , highlighted in red, will represent the variable whose inclusion in the regression equation is to be decided, with “good control” standing for bias reduction, “bad control” standing for bias increase, and “neutral control” when the addition of Z does not increase nor reduce asymptotic bias. For this last case, we will also make a brief remark about how Z could affect the precision of the ACE estimate.

Models 1, 2 and 3 – Good Controls




Figure 1: Models 1, 2 , and 3.

In Model 1, Z stands for a common cause of both X and Y . Once we control for Z , we block the back-door path from X to Y , producing an unbiased estimate of the ACE. In Models 2 and 3, Z is not a common cause of both X and Y , and therefore, not a traditional “confounder” as in Model 1. Nevertheless, controlling for Z blocks the back-door path from X to Y due to the unobserved confounder U , and again, produces an unbiased estimate of the ACE.

Models 4, 5 and 6 – Good Controls




Figure 2: Models 4, 5 and 6.

When thinking about possible threats of confounding, one needs to keep in mind that common causes of X and any mediator (between X and Y) also confound the effect of X on Y . Therefore, Models 4, 5 and 6 are analogous

to Models 1, 2 and 3—controlling for Z blocks the backdoor path from X to Y and produces an unbiased estimate of the ACE.

Model 7 – Bad Control




Figure 3: Model 7

We now encounter our first “bad control.” Here Z is correlated with the treatment and the outcome and it is also a “pre-treatment” variable. Traditional econometrics textbooks would deem Z a “good control” (Angrist and Pischke, 2009, 2014; Imbens and Rubin, 2015). The backdoor criterion, however, reveals that Z is a “bad control.” Controlling for Z will induce bias by opening the backdoor path $X \leftarrow U_1 \rightarrow Z \leftarrow U_2 \rightarrow Y$, thus spoiling a previously unbiased estimate of the ACE. This structure is known as the “M-bias”, and has spurred several discussions. Readers can find further discussion in Pearl (2009a, p. 186), and Shrier (2009); Pearl (2009c,b); Sjölander (2009); Rubin (2009); Ding and Miratrix (2015); Pearl (2015).

Model 8 – Neutral Control (possibly good for precision)




Figure 4: Model 8

In Model 8, Z is not a confounder nor does it block any backdoor paths. Likewise, controlling for Z does not open any backdoor paths from X to Y . Thus, in terms of asymptotic bias, Z is a “neutral control.” Analysis shows, however, that controlling for Z reduces the variation of the outcome variable Y , and helps improving the precision of the ACE estimate in finite samples (Hahn, 2004; Henckel et al., 2019).

Model 9 – Neutral Control (possibly bad for precision)




Figure 5: Model 9

Similar to the previous case, here Z is “neutral” in terms of bias reduction. However, controlling for Z will reduce the variation of the treatment variable X and so may hurt the precision of the estimate of the ACE in finite samples (Henckel et al., 2019, Corollary 3.4). As a general rule of thumb, parents of X which are not necessary for identification are harmful for the asymptotic variance of the estimator; on the other hand, parents of Y which do not spoil identification are beneficial. See Henckel et al. (2019) for recent developments in graphical criteria for efficient estimation via adjustment in linear models.

Model 10 – Bad Control




Figure 6: Model 10

We now encounter our second “pre-treatment” “bad control,” due to a phenomenon called “bias amplification” (Bhattacharya and Vogt, 2007; Wooldridge, 2009; Pearl, 2011, 2010, 2013; Middleton et al., 2016; Steiner and Kim, 2016). Naive control for Z in this model will not only fail to deconfound the effect of X on Y , but, in linear models, will amplify any existing bias.

Models 11 and 12 – Bad Controls



Figure 7: Models 11 and 12

If our target quantity is the ACE, we want to leave all channels through which the causal effect flows “untouched.” In Model 11, Z is a mediator of the causal effect of X on Y . Controlling for Z will block the very effect we want to estimate, thus biasing our estimates. In Model 12, although Z is not itself a mediator of the causal effect of X on Y , controlling for Z is equivalent to partially controlling for the mediator M , and will thus bias our estimates. Models 11 and 12 violate the backdoor criterion (Pearl, 2009a), which excludes controls that are descendants of the treatment along paths to the outcome.

Model 13 – Neutral Control (possibly good for precision)






Figure 8: Model 13

At first look, Model 13 might seem similar to Model 12, and one may think that adjusting for Z would bias the effect estimate, by restricting variations of the mediator M . However, the key difference here is that Z is a cause, not an effect, of the mediator (and, consequently, also a cause of Y). Thus, Model 13 is analogous to Model 8, and so controlling for Z will be neutral in terms of bias and may increase precision of the ACE estimate in finite samples. Readers can find further discussion of this case in Pearl (2013).

Models 14 and 15 – Neutral Controls (possibly helpful in the case of selection bias)



(a) Model 14



(b) Model 15

Figure 9: Models 14 and 15

Contrary to econometrics folklore, not all “post-treatment” variables are inherently bad controls. In Models 14 and 15 controlling for Z does not open any confounding paths between X and Y . Thus, Z is neutral in terms of bias. However, controlling for Z does reduce the variation of the treatment variable X and so may hurt the precision of the ACE estimate in finite samples. Additionally, in Model 15, suppose one has only samples with $W = 1$ recorded (a case of selection bias). In this case, controlling for Z can help obtaining the W -specific effect of X on Y , by blocking the colliding path due to W .

Model 16 – Bad Control




Figure 10: Model 16

Contrary to Models 14 and 15, here controlling for Z is no longer harmless. Adjusting for Z opens the backdoor path $X \rightarrow Z \leftarrow U \rightarrow Y$ and so biases the ACE.

Model 17 – Bad Control




Figure 11: Model 17

In our last example, Z is not a mediator, and one might surmise that, as in Model 14, controlling for Z is harmless. However, controlling for the effects of the outcome Y will induce bias in the estimate of the ACE, making Z a “bad control.” A visual explanation of this phenomenon using “virtual colliders” can be found in (Pearl, 2009a, Sec. 11.3). Model 17 is usually known as a “case-control bias” or “selection bias.” Finally, although controlling for Z will generally bias numerical estimates of the ACE, it does have an exception when X has no causal effect on Y . In this scenario, X is still d-separated from Y even after conditioning on Z . Thus, adjusting for Z is valid for testing whether the effect of X on Y is *zero*.

Final remarks

In this note, we demonstrated through illustrative examples how simple graphical criteria can be used to decide when a variable should (or should not) be included in a regression equation—and thus whether it can be deemed a “good” or “bad” control. Many of these examples act as cautionary notes against prevailing practices in traditional econometrics: for instance, Models 7 to 10 reveal that one should be wary of the propensity score “mantra” of conditioning on all “pre-treatment predictors of the treatment assignment;” whereas Models 14 and 15 demonstrate that not all “post-treatment” variables are “bad-controls”, and some may even help with identification. In all cases, structural knowledge is indispensable for deciding whether a variable is a good or bad control. Graphical models provide a natural language for articulating such knowledge, as well as efficient tools for examining its logical ramifications.

References

- Angrist, J. and Pischke, J.-S. (2009). *Mostly harmless econometrics: an empiricists guide*. Princeton: Princeton University Press.
- Angrist, J. D. and Pischke, J.-S. (2014). *Mastering 'metrics: The path from cause to effect*. Princeton University Press.

- Bhattacharya, J. and Vogt, W. B. (2007). Do instrumental variables belong in propensity scores? Technical report, National Bureau of Economic Research.
- Ding, P. and Miratrix, L. W. (2015). To adjust or not to adjust? sensitivity analysis of m-bias and butterfly-bias. *Journal of Causal Inference*, 3(1):41–57.
- Hahn, J. (2004). Functional restriction and efficiency in causal inference. *Review of Economics and Statistics*, 86(1):73–76.
- Henckel, L., Perković, E., and Maathuis, M. H. (2019). Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *arXiv preprint arXiv:1907.02435*.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Middleton, J. A., Scott, M. A., Diakow, R., and Hill, J. L. (2016). Bias amplification and bias unmasking. *Political Analysis*, 24(3):307–323.
- Pearl, J. (2009a). *Causality*. Cambridge University Press.
- Pearl, J. (2009b). Letter to the editor: Remarks on the method of propensity score. *Statistics in Medicine*, 28:1420–1423. URL: <https://ucla.in/2Nbs14j>.
- Pearl, J. (2009c). Myth, confusion, and science in causal analysis. *UCLA Cognitive Systems Laboratory*, Technical Report (R-348). URL: <https://ucla.in/2EihVyD>.
- Pearl, J. (2010). On a class of bias-amplifying variables that endanger effect estimates. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 417–424. URL: <https://ucla.in/2N8mBMg>.
- Pearl, J. (2011). Invited commentary: understanding bias amplification. *American journal of epidemiology*, 174(11):1223–1227. URL: <https://ucla.in/2PORDX2>.
- Pearl, J. (2013). Linear models: A useful “microscope” for causal analysis. *Journal of Causal Inference*, 1(1):155–170. URL: <https://ucla.in/2LcpmHz>.
- Pearl, J. (2015). Comment on ding and miratrix: “to adjust or not to adjust?”. *Journal of Causal Inference*, 3(1):59–60. URL: <https://ucla.in/2PgOWNd>.
- Rubin, D. B. (2009). Should observational studies be designed to allow lack of balance in covariate distributions across treatment groups? *Statistics in Medicine*, 28(9):1420–1423.

- Shrier, I. (2009). Propensity scores. *Statistics in Medicine*, 28(8):1317–1318.
- Sjölander, A. (2009). Propensity scores and m-structures. *Statistics in medicine*, 28(9):1416–1420.
- Steiner, P. M. and Kim, Y. (2016). The mechanics of omitted variable bias: Bias amplification and cancellation of offsetting biases. *Journal of causal inference*, 4(2).
- Wooldridge, J. (2009). Should instrumental variables be used as matching variables. Technical report, Citeseer.