Assessing The Effectiveness of Empirical Calibration Under Different Bias Scenarios

Hon Hwang

University of New South Wales

Juan C Quiroz

University of New South Wales

Blanca Gallego (✉ b.gallego@unsw.edu.au)

University of New South Wales

Research Article

Keywords: Bias, confounding, empirical calibration, negative controls, positive controls, simulations

Posted Date: November 15th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1058822/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

1 Assessing the effectiveness of empirical calibration 2 under different bias scenarios

3 Hon Hwang¹, Juan C Quiroz¹, Blanca Gallego^{1*}

4 ¹ Centre for Big Data Research in Health (CBDRH), University of New South Wales, Sydney,
5 2052, NSW, Australia

6 Corresponding author:

7 Blanca Gallego

8 Centre for Big Data Research in Health

9 Level 2, AGSM Building

10 G27, Botany St, Kensington NSW 2052

11 b.gallego@unsw.edu.au

12

13

14 **Key words:** Bias, confounding, empirical calibration, negative controls, positive controls, simulations

15

16

17

Abstract

Background: Estimations of causal effects from observational data are subject to various sources of bias. These biases can be adjusted by using negative control outcomes not affected by the treatment. The empirical calibration procedure uses negative controls to calibrate p-values and both negative and positive controls to calibrate coverage of the 95% confidence interval of the outcome of interest. Although empirical calibration has been used in several large observational studies, there is no systematic examination of its effect under different bias scenarios.

Methods: The effect of empirical calibration of confidence intervals was analyzed using simulated datasets with known treatment effects. The simulations were for binary treatment and binary outcome, with simulated biases resulting from unmeasured confounder, model misspecification, measurement error, and lack of positivity. The performance of empirical calibration was evaluated by determining the change of the confidence interval coverage and bias of the outcome of interest.

Results: Empirical calibration increased coverage of the outcome of interest by the 95% confidence interval under most settings but was inconsistent in adjusting the bias of the outcome of interest. Empirical calibration was most effective when adjusting for unmeasured confounding bias. Suitable negative controls had a large impact on the adjustment made by empirical calibration, but small improvements in the coverage of the outcome of interest was also observable when using unsuitable negative controls.

Conclusions: This work adds evidence to the efficacy of empirical calibration on calibrating the confidence intervals of treatment effects in observational studies. We recommend empirical calibration of confidence intervals, especially when there is a risk of unmeasured confounding.

1 Introduction

Observational studies are often used when a randomised controlled trial design is unethical, too costly, or time-consuming (1). The trade-off in doing so is the loss of randomisation of treatment assignment,

1 which is not guaranteed in observational studies. The lack of randomisation introduces confounding,
2 where there is a common cause for both the treatment and outcome (2, 3). Thus, treatment effect
3 estimation in observational studies should include adjustment for confounders to take into account
4 their biases, i.e. using inverse probability score weighting.

5 Confounding can persist even after adjustment for confounding, referred to as residual confounding,
6 occurring when the confounding variable is not measured, measured incorrectly, or when the
7 relationship between the confounder and outcome of interest is incorrectly modelled (2, 3). In these
8 situations, it is not possible to adjust for them directly. One way to account for residual confounding
9 is with *negative controls*, exposure-outcome pairs where it is believed the exposure does not affect the
10 outcome. Treatment effect estimates on the negative controls can be used to adjust for bias in the
11 treatment effect estimate on the outcome of interest, provided that they share the same casual structure
12 (4).

13 A prior study proposed empirical calibration of the p-values of the treatment effect estimates using an
14 empirical null distribution derived from negative controls (5). This idea was extended to empirical
15 calibration of the confidence interval of the treatment effect estimate, using an empirical error model
16 derived from negative controls and *positive controls*—synthetically generated from the negative
17 controls by simulating a relationship between the treatment and control outcomes (6). Empirical
18 calibration has been used to detect and adjust for bias in several observational studies (7-11). However,
19 prior work has not explored under which circumstances empirical calibration corrects for different
20 sources of bias of the treatment effect estimate or the nature of the bias-variance trade-off.

21 In this paper, we systematically examined the effect of empirical calibration on the bias of treatment
22 effect estimates by constructing data simulations with different types of biases. The simulations were
23 carried out in the context of binary treatment and binary outcome with biases resulting from
24 unmeasured confounder, model misspecification, measurement error, and lack of positivity. The
25 simulations examined the effect of empirical calibration through the resulting bias-variance trade-off
26 for each type of bias. Our work has implications for observational studies estimating the comparative

effectiveness of treatment strategies that plan to use empirical calibration to address residual confounding.

2 Methods

2.1 Empirical calibration procedure

Empirical calibration adjusts the p-values and confidence interval of a treatment effect using a linear model (in the parameters) called the *systematic error model*. This error model is constructed from known and hypothetically generated treatment effects using *negative control outcomes* and *positive control outcomes*, respectively. Calibration involves adjusting the new estimate and its confidence interval as though they are samples from a distribution of the true estimate, taking into account the systematic error.

2.1.1 Negative Controls

A negative control outcome is an outcome not believed to be affected by the treatment of interest (4, 6, 12). For example, Jackson et al. used hospitalisation from injury or trauma as a negative control outcome when examining the effect of influenza vaccination, since injury or trauma hospitalisation is not considered to be plausibly linked to influenza vaccination (13). Similarly, Schuemie et al. used ingrown nail as a negative control outcome when replicating studies that compared the adverse effect of using either the drug dabigatran or warfarin on patients with atrial fibrillation (6). Negative controls should, to the extent possible, have the same potential confounding mechanism as the outcome of interest (bias-comparable). In an ideal scenario, the set of common causes of the treatment and negative control outcome are identical to those of the treatment and outcome interest. In practice, this case is unrealistic since covariates are likely to affect the various outcomes differently with varying magnitudes.

2.1.2 Positive Controls

Positive control outcomes are outcomes for which the treatment of interest has known effects. Unlike negative control outcomes, obtaining suitable positive controls is challenging. Even if there is a known positive control, an estimate of its effect size may be highly uncertain due to the study design (10). In

the case where the effect size of a positive control is obtained from a randomised controlled trial, factors such as inclusion/exclusion criteria may not match that of the study with the outcome of interest. A secondary issue is that there may not be a wide range of positive effects available.

To sidestep the challenges of obtaining a range of positive controls from existing literature, it is possible to derive synthetic positive control outcomes from the negative controls. Schuemie et al. modelled negative controls using penalised regression with incidence rate ratios as the unit for the treatment effect (6). Subjects with the highest predicted probability of experiencing the outcome were then re-sampled and added to the treated group (in a binary treatment), resulting in a positive control with the desired treatment effect.

In this study, following previous simulation studies of bias adjustment (14, 15), we chose the log odds ratio as the measure of treatment effect. We generated positive controls by reusing the estimated regression coefficients from negative controls and setting the treatment effect to an adjusted target log odds ratio. The adjustment took into account non-zero treatment effects from the negative controls indicative of potential biases (12). For simplicity, we assumed a linear relationship between confounders and the logit of the control outcomes.

2.2 Simulations

We conducted a set of simulation experiments to determine the effect of empirical calibration when bias is present due to four common sources (2): unmeasured confounding, model misspecification (due to missing quadratic or interaction term), lack of positivity, and measurement error. For each bias type, the performance of empirical calibration was assessed under three different data generation conditions: (1) the limit (“ideal”) case in which the negative controls share identical potential confounder effects as the outcome of interest; (2) confounders affect all outcomes (negative controls and outcome of interest) via beta parameters selected at random; (3) the other limit (“worse”) case scenario in which the negative controls do not share the same potential causal pathway with the outcome of interest. Conditions (1) and (2) emulate cases where suitable negative controls are selected, i.e. they are bias-comparable to the outcome of interest. Case (3) emulates cases where unsuitable negative controls are selected, i.e. the negative controls are not bias-comparable to the outcome of interest.

The simulation process consisted of three steps: (1) generating a set of potential confounders X ; (2) generating the binary treatment values Z ; and (3) generating the binary outcomes consisting of the outcome of interest Y^* and S negative controls Y^- . We assumed that the outcome of interest and the negative control outcomes are all measured as part of the same large observational dataset. For simplicity, and following previous literature (14, 16), we used logistic linear regression as our base treatment and outcome models.

2.2.1 Generating confounders

We included ten measured confounders ($m = 10$) sampled from a normalised Gaussian distribution $X_m \sim \mathcal{N}(\mu = 0, \sigma^2 = 1)$. Supplementary material 1 shows the directed acyclic graph of this simulation setup with one negative control outcome and two measured confounders.

2.2.2 Generating the treatment variable

The treatment assignment probability was modelled as a logistic function that depended on the linear combinations of the confounders parametrised by $\alpha = \{\alpha_0 + \alpha_1 + \dots + \alpha_m\}$. Treatment values were then sampled from a Bernoulli distribution with probabilities:

$$\Pr(Z = 1 | X) = \text{logistic}(\alpha_0 + \alpha_1 X_1 + \alpha_2 X_2 + \dots + \alpha_m X_m),$$

$$Z \sim \text{Bern}[\Pr(Z = 1 | X)],$$

where $X = \{X_1, X_2, \dots, X_m\}$.

2.2.3 Generating the outcome of interest and the negative controls

To generate the outcome of interest and negative controls, the probability of outcome $Y = 1$ was modelled as a logistic function that depends on confounders and treatment variable. The combination of confounders that the logistic function depends on were varied depending on the types of bias we wanted to introduce. As its base, the logistic function is a linear combination of independent variables:

$$\Pr(Y^* = 1 | X) = \text{logistic}(\beta_0^* + \beta_Z^* Z + \beta_1^* X_1 + \beta_2^* X_2 + \dots + \beta_m^* X_m) \quad (6)$$

$$Y^* \sim \text{Bern}[\Pr(Y^* = 1 | X, Z)],$$

where $\{\beta_0^*, \beta_1^*, \dots, \beta_m^*\}$ are the regression parameters. Different bias scenarios were simulated by modifying this reference outcome model to include non-linear or interactive terms in equation 6.

The negative control outcome probabilities were also generated from a logistic model, but without the dependency on the treatment:

$$\Pr(Y_s^- = 1 | X) = \text{logistic}(\beta_{s0}^- + \beta_{s1}^- X_1 + \beta_{s2}^- X_2 + \dots + \beta_{sm}^- X_m) \quad (7)$$

$$Y_s^- \sim \text{Bern}[\Pr(Y_s^- = 1 | X)],$$

where $\{\beta_{s0}^-, \beta_{s1}^-, \beta_{s2}^-, \dots, \beta_{sm}^-\}$ are the regression parameters for s -th negative control outcome.

For brevity, f_s^* denotes the data generating function for the outcome of interest and f_s^- denotes the data generating function (and outcome model) for the s -th negative control.

2.2.4 Simulating bias

The bias scenarios simulated in this study are summarised in Table 1: unmeasured confounding, model misspecification, lack of positivity, and measurement error. Each bias scenario consisted of 500 simulation iterations with $N = 50,000$ subjects, resulting in 500 comparisons between the calibrated and uncalibrated estimates of the treatment effect. The relatively large sample size of 50,000 was chosen in the reference scenario to avoid biases associated with sample size effects. Within each simulation iteration, the treatment, outcome, and control model parameters were randomly sampled from a uniform distribution extending from $\log(0.5) = -0.693$ to $\log(2) = 0.6931$.

Unmeasured confounding: Confounding not measured or not controlled for in the analysis (3). This was introduced in our experiments by adding an extra confounder U to the simulated outcome models f^* and f_s^- , without making the confounder available at the time of treatment effect estimation. In one experiment (2.1) the negative controls shared the same unmeasured confounder as the outcome of interest. In a first instance, we simulated an ideal case in which the parameter associated with the unmeasured confounding in the outcome of interest (f^*) is replicated in all negative control models f_s^- . We then relaxed this constraint and allowed all the confounder parameters to be independently

generated. In a second experiment (2.2), the extra confounder U was added to the outcome of interest (f^*), but not to the negative controls (f_s^-).

Model misspecification: Biases associated to various forms of misspecification in the outcome models.

In our experiments, we introduced a quadratic term X_1^2 (experiment group three) and an interaction term between confounders X_1X_2 (experiment group four). In the first group of experiments (3.1 and 4.1), the negative controls shared the same misspecification as the outcome of interest, first with identical parameters and then with parameters chosen at random. In the second group of experiments (3.2 and 4.2), model misspecification was only present in the outcome of interest.

Lack of positivity: Lack of overlap between the treatment groups (in terms of propensity score), also known as structural violation of the positivity assumption. This was introduced in experiments 5.1 and 5.2 by modifying regions of the propensity score distribution as described in (17).

Measurement error: Confounders contain a systematic error or are subject to noise. This was introduced in experiment 6 by adding an error term $E \sim \mathcal{N}(\mu_e, \sigma_e^2)$ to the confounder that has larger effect size. The mean (greater than zero) and standard deviations were chosen at random for each simulation iteration.

2.2.5 Applying the empirical calibration procedure to the simulated data

The *systematic error model* for empirical calibration can be derived using regression estimates from (1) all the negative and positive controls or (2) only the negative controls (referred to as the null error model). The null error model assumes the systematic error is the same for all true effect sizes, whereas using all negative and positive controls models the systematic error as a function of the true effect size. Treatment effect estimates for the negative and positive controls are calculated using a standard inverse propensity score weighted logistic regression (referred to as estimation function). The systematic error model is then constructed using these estimates. The estimation function is also applied to the outcome of interest to estimate the treatment effect and its model-robust ‘sandwich’ standard error. This estimate is then calibrated using the systematic error model resulting in calibrated estimates and their

corresponding standard errors. See Supplementary material 2 for additional details on modelling the systematic error in empirical calibration.

For brevity, we refer to empirical calibration with the systematic error model as “default empirical calibration.” Sensitivity analysis included comparing default empirical calibration with empirical calibration with the null model. Five negative controls were used for empirical calibration, but sensitivity analysis included calibration with 30 negative controls to examine the effect of increasing the number of negative controls.

We used version R v4.0.3 (18) on the x86_64 architecture. Stabilised weights in the estimation function were obtained using version 0.10.0 of the WeightIt package (19), and their sandwich standard errors were obtained using the v4.0 survey package (20). Calibration of the estimates were performed using version v2.1.0 of the EmpiricalCalibration package (6). Our simulation software is licensed under GPLv3 and is available at <https://github.com/clinical-ai/assess-empcalib/>.

Funnel plots were used to represent the biases in the estimates of the treatment effect on the outcome of interest with and without calibration across each confounding scenario (Table 1). For each experiment, the funnel plots were generated for default empirical calibration with five negative controls. The funnel plots also include the results from sensitivity analysis: (1) empirical calibration with the null model, and (2) empirical calibration with 30 negative controls.

3 Results

3.1 Calibration of treatment effect on the outcome of interest

Figure 1 to 5 show funnel plots representing the biases in the estimates of the treatment effect on the outcome of interest with and without calibration across the bias scenarios outlined in Table 1 (experiments 2 to 6). The corresponding differences in coverage by the 95% confidence interval and bias of the outcome of interest after applying empirical calibration are presented in Table 2. Supplementary material 3 contains coverage plots of the calibrated controls in this study (presented as a diagnostic of empirical calibration performance in Schuemie et al. (6)). Running an experiment without any introduced confounding (experiment 1) resulted in zero bias (see Supplementary material).

Empirical calibration increased the coverage of the confidence intervals of the outcome of interest in 8/10 cases when negative controls were suitable (bias-comparable to the outcome of interest), with coverage increasing in the range of +1% to +60%. Empirical calibration increased coverage when bias was due to unmeasured confounding, model misspecification (quadratic term), non-positivity, and measurement error, though the increase for non-positivity (+2%) and measurement error (+1%) was small. When bad negative controls (not bias-comparable) were used for empirical calibration, coverage increased in 4/4 cases, though the increase in coverage was modest by comparison (+1% - 8%). The small increase in coverage for non-positivity and measurement error are due to uncalibrated confidence intervals having coverage of the outcome of interest close to 95% (Table 3).

Empirical calibration decreased bias in the outcome of interest in 8/10 cases when negative controls were suitable, in the scenarios of unmeasured confounder, model misspecification (interaction term), and measurement error. For non-positivity and model misspecification (quadratic term), the effect in bias was inconsistent. When unsuitable negative controls were used, calibration decreased bias in 3/4 cases.

3.2 Empirical Calibration with the Null Model

Empirical calibration with the null model decreased coverage in 9/10 cases when compared with default empirical calibration when suitable negative controls were used (coverage decrease ranging from -1% to -52%). Calibration decreased coverage for model misspecification (quadratic and interaction term), non-positivity, and measurement error. For unmeasured confounding, the change in coverage was inconsistent (-6% and +9%). With unsuitable negative controls, coverage decreased in 4/4 cases when compared with default empirical calibration (-3% to -8%).

With suitable negative controls, empirical calibration with the null model increased bias in the outcome of interest in 8/10 cases when compared with default empirical calibration. Bias increased for model misspecification (quadratic and interaction term) and measurement error. For unmeasured confounder and non-positivity, the change in bias was inconsistent. With unsuitable negative controls, bias increased in 2/4 cases.

3.3 *Empirical Calibration with a Higher Number of Negative Controls*

Increasing the number of negative controls from five to 30 increased coverage in 7/10 cases when suitable negative controls were used, the increase ranging from +1% to +6%. Coverage increased for unmeasured confounder, model misspecification (interaction term), and non-positivity. When using unsuitable negative controls, coverage increased in 2/4 cases. Increasing the number of negative controls had inconsistent effect on the bias, decreasing bias in 5/10 cases when suitable negative controls were used and decreasing bias in 2/4 cases when unsuitable negative controls were used.

4 Discussion

4.1 *Main Findings*

This paper examined the impact of empirical calibration across different types of biases introduced in simulation scenarios. In the majority of the simulations, empirical calibration increased coverage of the 95% confidence interval and decreased bias of the outcome of interest. Of all the bias scenarios, empirical calibration performed best in bias due to unmeasured confounder. Both suitable and unsuitable negative controls resulted in an increase in coverage and a decrease in bias, though the increase in coverage was lower when unsuitable negative controls were used for calibration. This suggests that when suitable negative controls are chosen for calibration, it benefits both coverage and bias of the outcome of interest, but calibration with unsuitable negative controls provides small benefits.

While prior work discussed the assumption that the negative control outcome should share the same potential causal mechanism as the outcome of interest to be effective, our work demonstrates some gains even when less than ideal negative controls were used for empirical calibration (4, 12). While the performance of empirical calibration was mostly positive, its performance was at times inconsistent in our bias scenarios in how it affected coverage and bias, with cases where coverage was increased at the expense of increased bias. Given the recommendation for empirical calibration in observational studies, our work suggests further examination of empirical calibration in controlled simulated datasets and in real-world datasets.

Our results showed minimal gains from running empirical calibration with 30 negative controls instead of five negative controls. A previous study {Suchard, 2019 #37} used 75 negative controls, with each one obtained by examining literature. Depending on the domain, this may not be feasible. Our results confirm the results from prior work showing gains when the number of negative controls is increased (6), but if a few negative controls are chosen with proper evidence, they can be sufficient for robust performance of empirical calibration.

Empirical calibration with the null model (using only negative controls) yielded lower coverage gains than default empirical calibration (using negative and positive controls). That is, the assumptions made by null model calibration decreased the effectiveness of empirical calibration. However, the small gains in coverage of the 95% confidence interval should encourage practitioners to apply empirical calibration with the null model in cases where it is unfeasible or too complex to derive positive controls from the negative controls.

4.2 Limitations

Our simulations explored five common bias scenarios, but real-world data is likely to result in combinations of different types of biases. The data generation and estimation functions were linear models, with non-linear and non-parametric functions to be explored in future work. The confounders in our simulations were independent and the number of measured confounders was fixed, which limited the data simulations explored. The treatment effect in our analysis was the log of an odds ratio (regression coefficient of a logistic regression). Future work needs to analyze how other treatment effects, such as relative risk and the corresponding generation of positive controls, may affect the performance of empirical calibration.

5 Conclusions

Our work adds evidence that empirical calibration can increase coverage of the 95% confidence interval and decrease bias of the outcome of interest. This is relevant for observational studies estimating the comparative effectiveness of treatments. Empirical calibration performed best when adjusting bias due to unmeasured confounder. Caution must be taken to select suitable negative controls, as unsuitable negative controls lessen the efficacy of empirical calibration.

1

2

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and materials

The datasets generated and analysed during the current study are available in the assessing empirical calibration repository, <https://github.com/clinical-ai/assess-empcalib/>.

Competing interests

The authors declare that they have no competing interests.

Funding

This work was supported by National Health and Medical Research Council, project Grant No. 1125414.

Authors' contributions

HH and BG designed the data simulations and experiments. HH coded the data simulations and empirical calibration procedure. HH, BG, and JQ analyzed and interpreted the results. HH drafted the initial manuscript. All authors contributed to critical revisions of the manuscript. All authors approved the final draft.

Acknowledgements

The authors thank Luca Maestrini for helpful comments. This work includes results produced on the computational cluster Katana at UNSW Sydney.

References

1. Black N. Why we need observational studies to evaluate the effectiveness of health care. *BMJ*. 1996;312(7040):1215-8.
2. Hernán MA, Robins JM. *Causal Inference: What If*: Chapman & Hall/CRC; 2020.
3. Sterne JAC, Hernán MA, McAleenan A, Reeves BC, Higgins JPT. Chapter 25: Assessing risk of bias in a non-randomized study. In: Higgins JPT, Thomas J, Chandler J, Cumpston M, Li T, Page MJ, et al., editors. *Handbook for Systematic Reviews of Interventions* version 61: Cochrane; 2020.
4. Lipsitch M, Tchetgen ET, Cohen T. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology*. 2010;21(3):383-8.
5. Schuemie MJ, Ryan PB, DuMouchel W, Suchard MA, Madigan D. Interpreting observational studies: why empirical calibration is needed to correct p-values. *Statistics in Medicine*. 2014;33(2):209-18.
6. Schuemie MJ, Hripcsak G, Ryan PB, Madigan D, Suchard MA. Empirical confidence interval calibration for population-level effect estimation studies in observational healthcare data. *Proceedings of the National Academy of Sciences of the United States of America*. 2018;115(11):2571-7.
7. Hripcsak G, Suchard MA, Shea S, Chen R, You SC, Pratt N, et al. Comparison of Cardiovascular and Safety Outcomes of Chlorthalidone vs Hydrochlorothiazide to Treat Hypertension. *JAMA Internal Medicine*. 2020;180(4):542-51.
8. Kim Y, Tian Y, Yang J, Huser V, Jin P, Lambert CG, et al. Comparative safety and effectiveness of alendronate versus raloxifene in women with osteoporosis. *Scientific Reports*. 2020;10(1):11115.
9. McQuade ETR, Liu J, Kang G, Kosek MN, Lima AAM, Bessong PO, et al. Protection from natural immunity against enteric infections and etiology-specific diarrhea in a longitudinal birth cohort. *The Journal of Infectious Diseases*. 2020;222(11):1858-68.
10. Suchard MA, Schuemie MJ, Krumholz HM, You SC, Chen R, Pratt N, et al. Comprehensive comparative effectiveness and safety of first-line antihypertensive drug classes: a systematic, multinational, large-scale analysis. *The Lancet*. 2019;394(10211):1816-26.
11. Schuemie MJ, Ryan PB, Pratt N, Chen R, You SC, Krumholz HM, et al. Large-scale evidence generation and evaluation across a network of databases (LEGEND): assessing validity using hypertension as a case study. *Journal of the American Medical Informatics Association*. 2020;27(8):1268-77.
12. Shi X, Miao W, Tchetgen ET. A Selective Review of Negative Control Methods in Epidemiology. *Current Epidemiology Reports*. 2020;7(4):190-202.
13. Jackson LA, Jackson ML, Nelson JC, Neuzil KM, Weiss NS. Evidence of bias in estimates of influenza vaccine effectiveness in seniors. *International Journal of Epidemiology*. 2005;35(2):337-44.
14. Austin PC. A comparison of 12 algorithms for matching on the propensity score. *Statistics in Medicine*. 2014;33(6):1057-69.
15. Austin PC, Grootendorst P, Normand S-LT, Anderson GM. Conditioning on the propensity score can result in biased estimation of common measures of treatment effect: a Monte Carlo study. *Statistics in Medicine*. 2007;26(4):754-68.
16. Austin PC. The performance of different propensity score methods for estimating marginal hazard ratios. *Statistics in Medicine*. 2012;32(16):2837-49.
17. Kang J, Chan W, Kim M-O, Steiner PM. Practice of causal inference with the propensity of being zero or one: assessing the effect of arbitrary cutoffs of propensity scores. *CSAM*. 2016;23(1):1-20.
18. Team RC. *R: A Language and Environment for Statistical Computing*: R Foundation for Statistical Computing; 2020.
19. Greifer N. *WeightIt: Weighting for Covariate Balance in Observational Studies* 2020.
20. Lumley T. Analysis of Complex Survey Samples. *Journal of Statistical Software*. 2004;9(1):1-19.

Table 1. Bias scenarios with covariate dependencies in the data generation and estimation functions. For each bias scenario, two experiment groups explore the effect of suitable negative controls (having the same potential confounding mechanism as the outcome of interest) and unsuitable negative controls. The third column shows the covariate dependencies of the data generating functions: f^* for the outcome of interest and f^- for the negative controls. X is the set of measured confounders.

Bias Scenario	Experiment Group	Outcome generation function
Reference case	1	f^* and f^- depend on X
Unmeasured confounder (U)	2.1	f^* and f^- depend on X and U
	2.2	Only f^* depends on X and U
Model misspecification: Quadratic term (X_1^2)	3.1	f^* and f^- depend on X and X_1^2
	3.2	Only f^* depends on X and X_1^2
Model misspecification: Interaction between two confounders (X_1X_2)	4.1	f^* and f^- depend on X and X_1X_2
	4.2	Only f^* depends on X and X_1X_2
Lack of positivity	5.1	f^* and f^- depend on X .
	5.2	Only f^* depends on X .

Lack of overlap between treatment groups in terms of propensity score.		
Measurement error in confounder	6	f^* and f^- depend on X . One of the confounders with large effect have random measurement error.

Table 2. Differences in coverage of the outcome of interest by the 95% confidence interval and difference in mean of standardised absolute biases of the outcome of interest. Default empirical calibration was performed with five negative controls and using the all controls systematic error model.

Difference Coverage				Difference in Mean of Standardised Absolute Bias		
	Ideal— Suitable negative controls	Random— Suitable negative controls	Random— Unsuitable negative controls	Ideal—Suitable negative controls	Random—Suitable negative controls	Random—Unsuitable negative controls
No calibration vs default empirical calibration						
Unmeasured confounder	0.48	0.60	0.05	-0.46	-0.23	0.00
Model misspecification Quadratic term	0.24	0.27	0.08	-0.05	0.03	-0.01
Model misspecification Interaction term	-0.09	0.04	0.02	-0.19	-0.09	-0.02
Non-positivity	0.02	0.00	0.01	-0.02	0.07	-0.01
Measurement error	0.01	0.01	NA	-0.02	-0.02	NA
Default empirical calibration vs empirical calibration with the NULL systematic error model						
Unmeasured confounder	0.09	-0.06	-0.04	0.55	-1.56	0.00
Model misspecification Quadratic term	-0.09	-0.13	-0.08	0.20	0.35	0.02
Model misspecification Interaction term	-0.52	-0.06	-0.03	0.29	0.00	0.01
Non-positivity	-0.02	-0.01	-0.03	-0.01	0.09	-0.01

Measurement error	-0.04	-0.02	NA	0.02	0.01	NA
Default empirical calibration (5 negative controls) vs empirical calibration with 30 negative controls						
Unmeasured confounder	0.05	0.06	0.01	0.85	-8.31	-0.36
Model misspecification Quadratic term	0.00	0.00	-0.05	0.04	-0.11	0.08
Model misspecification Interaction term	0.02	0.02	0.02	0.06	0.06	-0.60
Non-positivity	0.02	0.03	0.00	-0.02	0.43	0.12
Measurement error	0.01	-0.01	NA	-0.11	-0.06	NA

Table 3. Coverage of the outcome of interest by the 95% confidence interval without calibration and with empirical calibration.

	Ideal—Suitable negative controls		Random—Suitable negative controls		Random—Unsuitable negative controls	
	Calibrated	Uncalibrated	Calibrated	Uncalibrated	Calibrated	Uncalibrated
Unmeasured Confounder	0.79	0.31	0.91	0.31	0.36	0.31
Model misspecification Quadratic term	0.99	0.75	0.99	0.72	0.82	0.74
Model misspecification Interaction term	0.87	0.96	0.98	0.94	0.96	0.94
Non-positivity	0.96	0.94	0.94	0.94	0.96	0.95
Measurement error	0.96	0.95	0.96	0.95	NA	NA

Figure 1. Bias of treatment estimates and the coverage of their standard errors by the 95% confidence interval in the unmeasured confounding scenario. In the “ideal” case the effect of confounders to the outcomes are the same, and in the “random coefficients” case these effects are randomised. Within each cell, the left funnel plot shows the estimates calibrated with negative controls only and the right funnel plot shows estimates calibrated with negative and positive controls.

Figure 2. Bias of treatment estimates and the coverage of their standard errors by the 95% confidence interval in the model misspecification – missing quadratic term scenario. In the “ideal” case the effect of confounders to the outcomes are the same, and in the “random coefficients” case these effects are randomised. Within each cell, the left funnel plot shows the estimates calibrated with negative controls only and the right funnel plot shows estimates calibrated with negative and positive controls.

Figure 3. Bias of treatment estimates and the coverage of their standard errors by the 95% confidence interval in the model misspecification – missing

interaction term scenario. In the “ideal” case the effect of confounders to the outcomes are the same, and in the “random coefficients” case these effects are randomised. Within each cell, the left funnel plot shows the estimates calibrated with negative controls only and the right funnel plot shows estimates calibrated with negative and positive controls.

Figure 4. Bias of treatment estimates and the coverage of their standard errors by the 95% confidence interval in the non-positivity scenario. In the “ideal” case the effect of confounders to the outcomes are the same, and in the “random coefficients” case these effects are randomised. Within each cell, the left funnel plot shows the estimates calibrated with negative controls only and the right funnel plot shows estimates calibrated with negative and positive controls.

Figure 5. Bias of treatment estimates and the coverage of their standard errors by the 95% confidence interval in the measurement error scenario. In the “ideal” case the effect of confounders to the outcomes are the same, and in the “random coefficients” case these effects are randomised. Within each cell, the left funnel plot shows the estimates calibrated with negative controls only and the right funnel plot shows estimates calibrated with negative and positive controls.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementarymaterial.pdf](#)