Monte Carlo Sensitivity Analysis for Unmeasured Confounding

Lara Maleyeff

June 9, 2016

INTRODUCTION

The conclusions of a regression analysis rely on a number of fundamental assumptions. One such assumption is that the model is correctly specified; the validity of all inferences from any analysis is dependent upon the extent to which these assumptions are met. Any violation of the necessary assumptions can result in bias caused by systematic error, i.e. error that cannot be reduced by increasing the number of observation units. Most epidemiological studies do not quantify bias in their results as many current methods are too complicated, expensive (because, for example, they require collection of additional data), or hard to understand. However, statisticians can use sensitivity analysis techniques to quantify the effect of a hypothesized unmeasured confounder and may then adjust the conclusions of the analysis accordingly. To do this the researcher must make assumptions about certain characteristics regarding the hypothesized unmeasured confounder. Sensitivity analysis techniques can be *local* (using a cost function and point parameter estimation) or *global* (exploring the design

space using a representative set of samples and varying parameter estimates). The more valid assumptions a statistician can make in their analysis, the more precise the results will be.

This paper will propose and evaluate a method for implementing a sensitivity analysis for the violation of the assumption of no unmeasured confounders. It will explore a global Monte Carlo based approach, varying inputs and parameter estimates, to quantify the robustness of the original analysis to the violation of the assumption of no unmeasured confounders. The first section will describe the existing sensitivity analysis approaches for the violation of the assumption of no unmeasured confounders. These methods range from simple point estimates based on sensitivity parameters, to Bayesian and Monte Carlo approaches. The second section describes the proposed method and evaluates its performance via a simple simulation study.

A BRIEF SURVEY OF THE LITERATURE

Background

Publications were reviewed that used sensitivity analysis methods to calculate bias due to unmeasured confounders in a study. The key publications are described below and summarized in Table 1.

Local Analysis

The early papers on sensitivity analysis of unmeasured confounding utilize a local sensitivity analysis method, yielding a point estimate. These papers estimated bias

using informed assumptions about potential confounders combined with the study's data. Most of the early studies assume a single exposure and a single confounder. Later publications generalize this method[5-8].

In 1978 Schelessman conducted a study[1] that utilizes local sensitivity analysis techniques. This method assumes only three factors (exposure, disease, and one confounding variable) and uses odds ratios and external information about a suspected confounder to calculate a point estimate of bias caused by unmeasured confounding. Similarly, Gail et al.[9] utilize a crude relative risk term and Flanders et al.[10] use a confounding risk ratio to quantify bias (See Table 1).

Schneeweiss et al.[2] describe a more complex implementation of this method. Here, the authors used informed assumptions about the hypothesized unmeasured confounder(s), combined with an analysis of the residual confounding that would have been necessary to explain the study results, to quantify the amount of bias present in the original analysis. This method uses a confounded relative risk (RR) term, which is also referred to as an "apparent" RR. This term depends on certain sensitivity parameters, which can be estimated/adjusted to explore effects. Using this method, a researcher can create an educated point estimate of bias in a fairly simple manner; that is, without an unreasonable amount of technical understanding. However, this method does not create a confidence interval for the bias, which researchers may find useful. Breslow et al.[3] also write about a confounded relative risk term in chapter 3 of their 1980 publication. This confounding risk ratio is given as the ratio of the crude odds ratio

to the post stratification odds ratio. If the confounding risk ratio is one, then there is no confounding present.

VanderWheele et al.[4] describe another method that uses the potential outcomes framework to assess the magnitude of bias caused by unmeasured confounding. The potential outcomes framework assumes a counterfactual world where the exposure is not present to determine its outcome. Here, the authors derive a general class of formulas for sensitivity analysis of uncontrolled confounding that can be simplified with further assumptions. This again presents a method to calculate a point estimate of bias in observational studies.

Bayesian and Monte Carlo Sensitivity Analysis

More recent publications on the subject describe a more automated approach - in the form of a Bayesian analysis or a Monte Carlo sensitivity analysis. A Bayesian analysis involves the declaration of a prior distribution or a class of prior distributions to represent uncertainties about one or more unmeasured confounders. This information is then combined with a likelihood model and Bayes' Theorem is applied to give a posterior distribution. A Monte Carlo sensitivity analysis, the focus of this paper, compares a number (in the thousands) of randomly sampled confounding scenarios to repeatedly estimate the bias. The scenarios can be based on a prior distribution, as in a Bayesian analysis. Both approaches offer the significant advantage over the local methods described in the previous section by providing a natural means of computing measures of variability such as standard errors or confidence intervals.

A 2006 publication by McCandless et al.[5] considers a Bayesian sensitivity analysis for unmeasured confounding. A Bayesian analysis involves the definition of a prior distribution and a "sampling" from that prior distribution, which in this case is done via a Markov Chain Monte Carlo approach. Here, they assume that the association between a binary exposure, a binary response, measured confounders, and a single binary unmeasured confounder can be formulated via logistic regression models. Since the model for the unmeasured confounder cannot be identified, the authors elect to use Markov Chain Monte Carlo approaches to investigate the effect of different priors.

In a study conducted by Steenland et al.[6], the authors provide a method for discussing the *range* of bias using either Monte Carlo sensitivity analysis or Bayesian sensitivity analysis (rather than the traditional point estimates). In conventional analyses, bias may cause the confidence intervals to be too narrow, shifted upward or shifted downward. The Monte Carlo analysis is analogous to the Bayesian analysis, yet is much simpler to implement because it can be accomplished within a more familiar, frequentist framework. In ordinary sensitivity analysis, external information (or hypothesized information) is used to estimate the effect of the confounder on the observed responses. Monte Carlo sensitivity analysis is simply an expanded version of this method. In the Monte Carlo approach, the authors used 5000 randomly sampled confounding scenarios to repeatedly estimate the bias factor. In the Bayesian sensitivity analysis, the authors combined the observed data entered into a data model with prior distributions for the parameters to derive a posterior distribution for the parameters. This latter approach is more complicated and difficult to understand compared to the Monte

Carlo version for the classically trained data analyst; thus analysts with frequentist training may prefer the Monte Carlo approach.

A paper by Phillips et al.[7] seeks to quantify the uncertainty introduced by systematic error. Their approach is to estimate a probability distribution for a bias-corrected effect based on external information in the form of externally derived distributions of data. Here, the authors used Monte Carlo simulation to combine multiple correction for bias. The bias calculations are similar to those in normal sensitivity analysis, while being a more complete reporting method.

A very simple method of bias calculation would be to calculate a point estimate of bias necessary for the study to have produced the observed data. This paper refers to this approach as target-adjustment sensitivity analysis (TASA). An alternative approach is bias-level sensitivity analysis, which specifies bias parameters and calculates the resulting adjusted effect. This approach differs from TASA because it calls for forming beliefs about the probabilities of various states of the world, rather than simply plausibility.

A study that uses Bayesian Additive Regression Trees (BART) was conducted by Dorie et al.[8] BART is a Markov Chain Monte Carlo method that draws from a regression function, f, assuming a prior on f and specifying a Markov Chain whose stationary distribution is the posterior distribution of f. In this study, the sensitivity parameters are the regression coefficients. The authors assume that a confounder exists and determine the level of confounding present to drive the naive treatment effect to zero/insignificance. The authors make no assumptions about the parameters of the

population distribution(s) from which the data were drawn. This approach removes the error that may result from misspecification of a distribution, e.g. the errors that result from parametric tests.

Summary

The conclusions of a regression analysis rely on many assumptions, one of which is that there are no unmeasured confounders. Literature from 1955 through today describes many sensitivity analysis methods to calculate bias due to unmeasured confounders in a study (see Table 1). The older studies focus on point estimates of bias by varying one or two parameters, usually using odds ratios to calculate these estimates. These methods certainly cost less (computationally and perhaps also monetarily) and are easier to understand than a Monte Carlo or Bayesian approach, but may be less accurate and cannot provide a range, e.g. in the form of a confidence interval. More modern studies focus on Bayesian and Monte Carlo sensitivity analyses. A Bayesian analysis requires more distributional assumptions and may be unfamiliar to analysts following the more traditional frequentists training. An automated Monte Carlo analysis does not require the same assumptions, but comes with the usual caveats of interpretation that all frequentist analyses face. This paper will explore a global Monte Carlo based approach, varying inputs and parameter estimates, to quantify the robustness of the original analysis to the violation of the assumption of no unmeasured confounders.

Table 1. Summary of Relevant Literature

Author	Target of Estimation	Parameters Required	Example	Conclusions
Schlesselman, 1978[1]	Relative Risk	The disease, study agent, and confounder (where each is considered to be binary).	Oral contraceptive (OC) use and smoking status among controls and cases of nonfatal myocardial infarction (MI) in women under 45 years	This method only deals with dichotomous variables. In order to extend these methods, more information than was required in this analysis would be necessary.
Gail et al., 1988[9]	Crude Relative Risk	External information on on the joint distribution of confounder and exposure, together with external information on the relative risk of disease due to the confounding factor among unexposed individuals.	Lung cancer rates due to age confounded by smoking status	Here the authors assume a joint model for exposures and confounders which must be explored sufficiently for reliable results. The authors also assume a constant effect of exposure and confounder. It may be impractical to get such thorough and complete information about the confounder.
Flanders et al., 1990[10]	Confounding Risk Ratio	Prevalence of the covariate in the population; the association between the exposure and the covariate; and the effect of the covariate on disease	Lung cancer rates due to occupation confounded by smoking status	Practical implementation: it may not be feasible to acquire data on the confounder; This approach only requires specification of one or two parameters; results can be extended to incidence rate ratio or to the odds ratio
Schneeweiss, 2006[2]	Confounded Relative Risk	External adjustment in a validation study of the unmeasured confounder; or in the absence of such information we can vary parameters in a sensitivity analysis	Associations between newer sedative hypnotics and hip fractures, statin use and cancer, selective COX-2 inhibitors and cardiovascular events, and anti-TNFa therapy and lymphatic malignancies	Quantitative sensitivity analysis are easy to perform (using spreadsheets); if external information is available it may be included to increase accuracy; Propensity Score Calibration can be used with more than one confounder present; must interpret results cautiously

Author	Target of Estimation	Parameters Required	Examples	Conclusion
Breslow et al., 1980[3]	Risk Ratio	Cross tabulate disease against exposure, based on pooling data over levels of the confounder	The disease could be lung cancer, the exposure some occupation primarily of blue-collar workers, and the confounder cigarette smoking	This method gives results for many levels of stratification of the confounder; the joint confounding risk ratio will always be more extreme than any one singly
VanderWeele, 2011[4]	General Bias Formulas; additive, risk-ratio and odds-ratio scales	Potential outcomes framework: for each treatment a, the potential outcome for an individual if the treatment had not been set to a	A study comparing coronary artery bypass surgery to medical therapy in the treatment of coronary artery disease. The outcome is binary: symptomatic relief after 6 months. The data available are 64 covariates.	These formulas have allowed for binary, ordinal, or continuous outcomes; categorical or continuous treatment; and categorical or continuous measured and unmeasured confounding variables. This method uses some simplifying assumptions: the confounder has a constant effect on the treatment across all covariate levels; does not assume the absence of interactions among treatments and covariates.
McCandless et al., 2006[5]	Bias parameters using a Bayesian analysis	External information about unmeasured confounding is incorporated into the analysis as prior distributions on bias parameters. The posterior distribution of the exposure effect summarizes uncertainty due to unmeasured confounding in addition to random error	Consider a retrospective cohort study designed to estimate the effectiveness of beta blocker therapy for treatment of heart failure using linked records of hospital episodes, prescription claims data, and death certificates.	One concern with this approach is the impact of the choice of prior distribution on the result (impact is usually unknown). We cannot guarantee confidence intervals as in standard large sample theory.

Author	Target of Estimation	Parameters Required	Examples	Conclusion
Steenland et al., 2004[6]	Likely range of bias using Monte Carlo Analysis and Bayesian Sensitivity Analysis	A distribution for the smoking habits of workers and referents, a distribution of rate ratios for the effect of smoking on lung cancer, and a model for the bias parameter	Data from a US silica-lung cancer study in which results were potentially confounded by smoking.	The Monte Carlo analysis improves upon traditional sensitivity analyses as it provides a range rather than a point estimate. In situations where the prior information is less precise, the bias distribution will lead to substantially wider Monte Carlo and Bayesian intervals relative to the conventional confidence intervals.
Phillips, 2003[7]	Probability distribution for a bias-corrected effect measure; corrected odds ratio using a Monte Carlo Analysis	Externally-derived distributions of bias levels	The report of the Hemorrhagic Stroke Project case-control study which linked the decongestant and diet aid, phenylpropanolamine (PPA), to hemorrhagic stroke	This method may be improved upon - especially in terms of the generation of input distributions. Compared to the overall cost of a study, this method is inexpensive and feasible
Dorie et al., 2015[8]	Coefficient on an unmeasured confounder	Assume a confounder exists and determine the level of confounding required to drive the native treatment effect to zero using Bayesian Additive Regression Trees. No assumptions about population distribution from which data is drawn.	The effectiveness of anti-hypertensive drugs on the level of blood pressure	This method makes no structural assumptions - which is much more conducive to a real life dataset. The sensitivity parameters should be based on real life assumptions and should be easily interpretable.

SIMULATION STUDY

Sensitivity analysis is the study of how uncertainty in the output of a model can be attributed to uncertainty in its inputs. In this section, a sensitivity analysis was used to examine the impact of a Monte Carlo based sensitivity analysis for no unmeasured confounders. The simulation will generate two correlated covariates along with an outcome which depends on both. The analysis of the data will then ignore one of the generated covariates, mimicking a situation in which there is an unmeasured confounder present.

Uncorrected Bias Estimation

Data were generated (x, z, and y) assuming there is one unmeasured confounder (x), which affects both the exposure (z) and the outcome (y). Then, the resulting bias was calculated (if the confounder is ignored or unavailable) and the outcome model regresses only on the exposure. In order to generate the data the following assumptions are made:

$$x \sim N(0,1)$$

$$z \sim N(f(x),\ 1)$$

$$y \sim \beta_0 + \beta_1 x + \beta_2 z + \varepsilon,\ where\ \varepsilon \sim N(0,1)$$

Note that the normal distribution can be replaced with any distribution (uniform, poisson, exponential, etc.). In the first analysis, the parameters of the following model are estimated: $y = \widehat{\beta}_0 + \widehat{\beta}_1 z + \varepsilon$. Finally, the bias in the estimate of β_1 is explored (as the

intercept is typically of less interest the an exposure effect). The R code below indicates how the data are generated and the bias is calculated. Note that, in the simulation, all three parameters (β_0, β_1, and β_2) are set equal to 1.

```
       N <- c(20,50,100,500)
       n.sim <- 1000
       all.ests.b0 <- NULL
       all.ests.b1 <- NULL
       true.b0 <- 0
       true.b1 <- 1

       for (n in N) {
              n.ests.b0 <- NULL
              n.ests.b1 <- NULL
              n.real.b0 <- NULL
              n.real.b1 <- NULL
        for (s in 1:n.sim)
        {
              x <- rnorm(n)
              z <- rnorm(n, 0.3*x)
              y <- rnorm(n, 1+z+x)

              n.ests.b0 <- c(n.ests.b0, coef(lm(y~z))[1])
              n.ests.b1 <- c(n.ests.b1, coef(lm(y~z))[2])

        }

              all.ests.b0 <- cbind(all.ests.b0, n.ests.b0)
              all.ests.b1 <- cbind(all.ests.b1, n.ests.b1)

       }

       bias.b0 <- all.ests.b0 - true.b0 ## Bias computed relative to data generation
       bias.b1 <- all.ests.b1 - true.b1 ## parameters

       ests.sd.b0 = c(sd(all.ests.b0[,1]), sd(all.ests.b0[,2]), sd(all.ests.b0[,3]),
sd(all.ests.b0[,4]))
       ests.sd.b1 = c(sd(all.ests.b1[,1]), sd(all.ests.b1[,2]), sd(all.ests.b1[,3]),
sd(all.ests.b1[,4]))
```

Table 2 includes summary statistics for the parameter estimates of β_1 resulting from 1000 simulated data sets. Figure 1 shows the distribution of the bias of β_1. The bias is is the difference between the estimated value and the true value of 1.

Table 2: Means and standard deviations of β_1 .

n	Mean of β_1	Standard Deviation
20	1.2748	0.3272
50	1.2584	0.1870
100	1.2731	0.1370
500	1.2751	0.0590

Figure 1: Bias of the uncorrected estimates of β_1 with varying n.

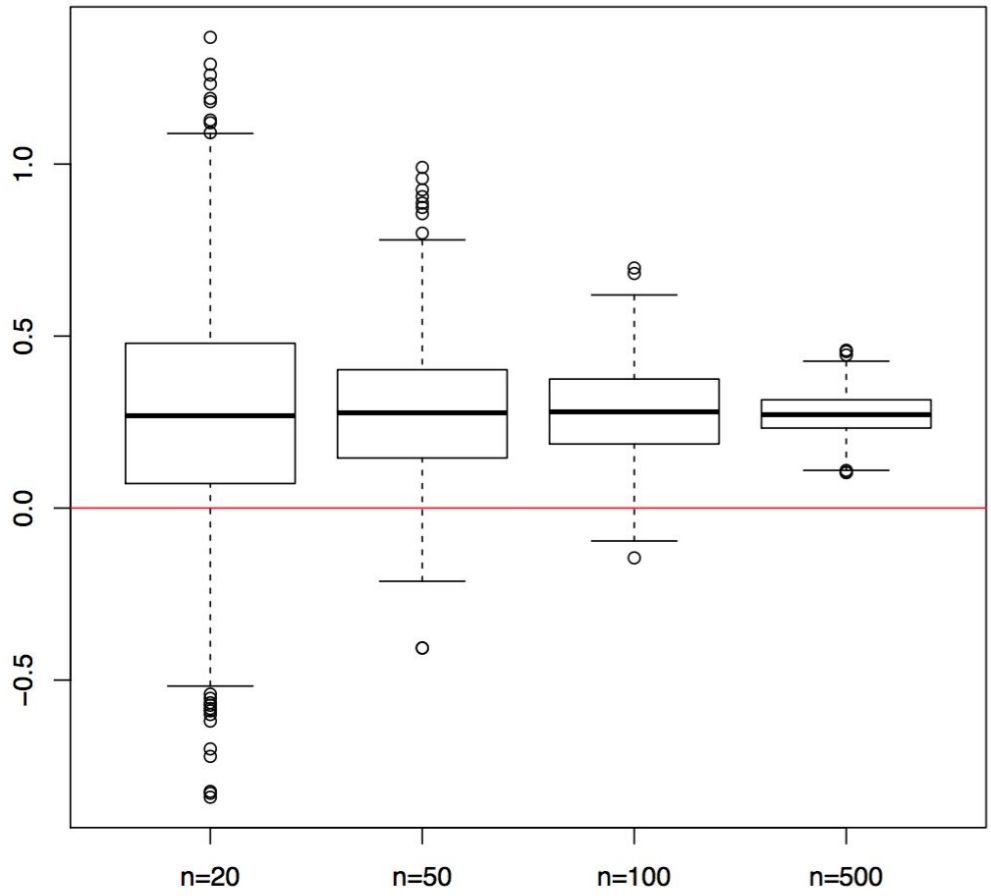

Proposed Bias-Corrected Estimation

Next, the algorithm is modified to correct for bias. The following additional notation is used. Let x be the unmeasured confounder, as stated earlier. Then x is of the form:

$$x = \alpha_0 + \alpha_1 \times z + error.$$

In a real data analysis study, this external information will likely be unknown or incomplete. In other words, we may have unknown or incomplete information about $\alpha_j ; j = 0, 1$. This (possibly incomplete or inaccurate) information would then be used to postulate plausible values for $\widehat{\alpha_j} ; j = 0, 1$. In this simulation, three postulated distributions for the confounding are considered: In the first scenario (which will be referred to as Scenario A) $\widehat{\alpha_j} \sim N(\alpha_j, 0.25) ; j = 0, 1$, i.e. correct and specific information about the confounder is available. In the next scenario, Scenario B, $\widehat{\alpha_j} \sim Unif[\alpha_j - 0.25, \alpha_j + 0.25] ; j = 0, 1$, i.e. correct, and more specific (variance is approximately one third the variance of Scenario A), information about the confounder is available. In Scenario C, $\widehat{\alpha_j} \sim Unif[\alpha_j - 0.75, \alpha_j + 0.75] ; j = 0, 1$. Here again the distribution is correctly centered, but the variance is nearly three times larger than in Scenario A. Thus, this is quite clearly a weaker case than Scenarios A and B, but it still may perform better than a naive analysis. In the final scenario, Scenario D, incorrect, somewhat nonspecific information is assumed: $\widehat{\alpha_j} \sim Unif[-1 - 0.1, -1 + 0.4] ; j = 0, 1$ so that variance is the same as in Scenario B but the distribution is not correctly centered.

The algorithm for generating the adjusted β_j's is as follows:

For each simulated dataset (Z, Y) available to the analyst:

Step 1: Sample from postulated distribution of α_j's.

Step 2: Generate x* according to the sampled $\widehat{\alpha_j}$'s.

Step 3: Estimate parameters from the model y ~ z + x* and record.

Step 4: Repeat until desired number of simulations completed (1000).

Figure 2 shows the bias of the estimator of β_1 for various sample sizes across the 4 scenarios, as well as the uncorrected case (reproducing the pattern observed in Figure 1). Table 3 includes summary statistics for the parameter estimates.

Figure 2. Bias of β_1 for various sample sizes across Scenarios A-D and the Uncorrected case (indicated with a U).

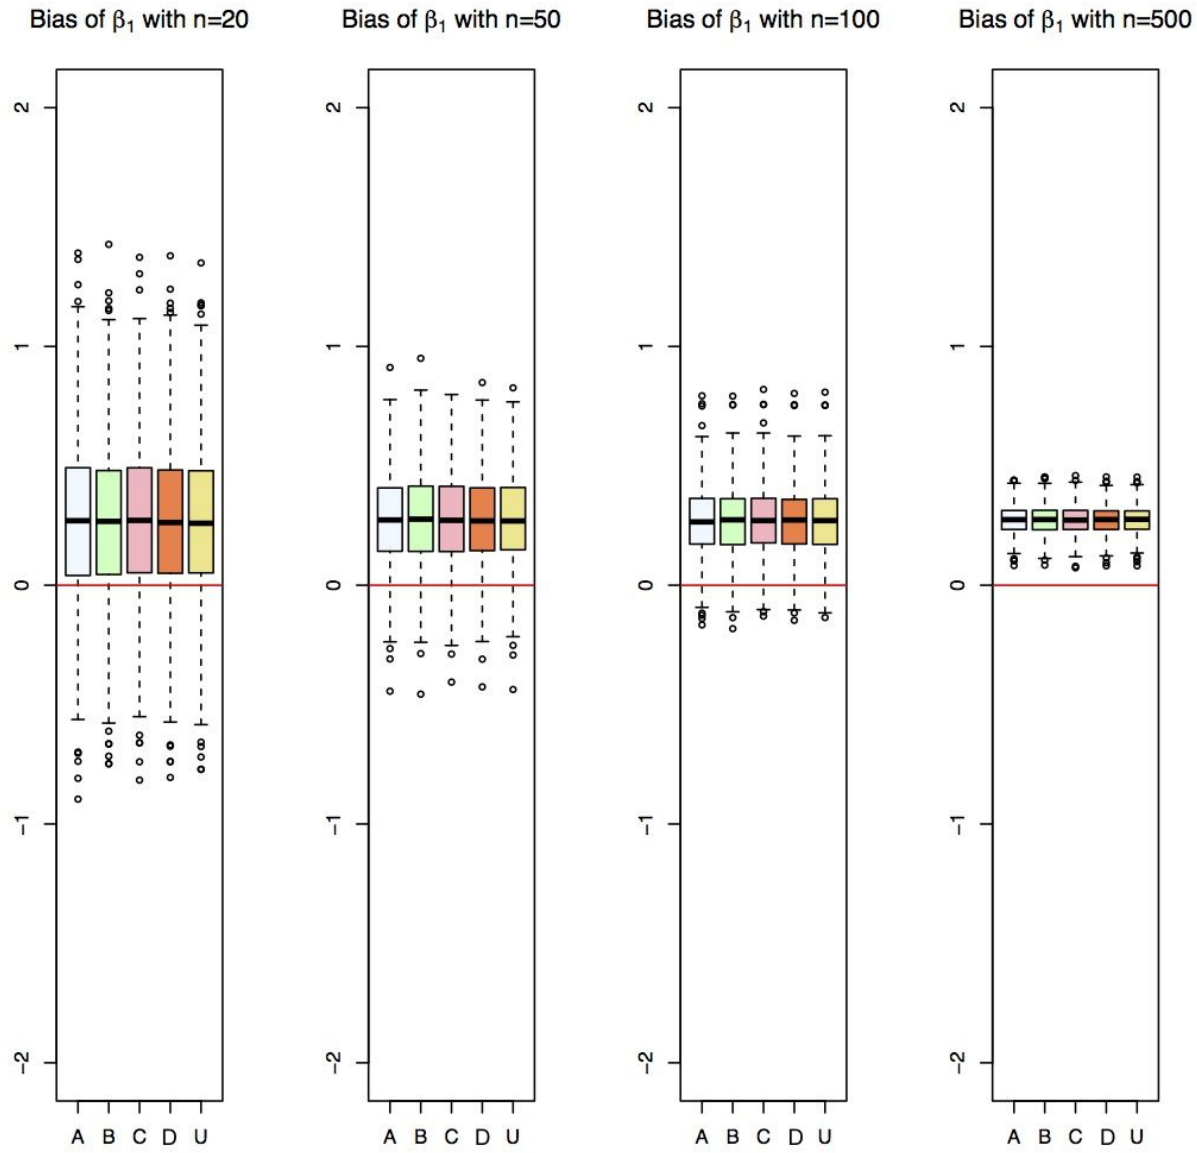

Table 3: Parameter estimates of β_1 for Scenarios A-D.

	Scenario A		Scenario B		Scenario C		Scenario D	
n	Mean	SD	Mean	SD	Mean	SD	Mean	SD
20	1.2812	0.3483	1.2720	0.3451	1.2773	0.3439	1.2756	0.3318
50	1.2596	0.1989	1.2610	0.1952	1.2614	0.1929	1.2590	0.1872
100	1.2716	0.1417	1.2713	0.1462	1.2701	0.1430	1.2728	0.1379
500	1.2744	0.0626	1.2748	0.0628	1.2754	0.0614	1.2751	0.0593

As observed in the tables, the proposed approach does not succeed in reducing bias due to unmeasured confounding whether using an informed distribution for the x-z relationship parameters or not. This finding can be explained by analogy with measurement error: in the Monte Carlo approach, the method attempted to “recover” the missing confounder, x, but generating a new variable that was a proxy for x given information on z. However, this variable was not sufficient to control for confounding because (i) it did not acknowledge the relationship between x and y, and (ii) it was a noisy (mis-measured) version of x -- more closely correlated with z than x itself, and therefore adjusting for this error-prone version of the confounder was insufficient to control for the confounding by x.

CONCLUSION

To conclude, accounting for confounding is an important aspect of any regression analysis. The results of an analysis can be severely biased if unmeasured confounding has not been accounted for. This can and will result in real life implications, especially and most notably in an epidemiological or health research context. Current literature on bias measurement for unmeasured confounding addresses these concerns, adding that most epidemiological studies lack a sufficient bias analysis. The literature on the subject includes a wide range of processes for approximating bias in a regression model, ranging from simple point estimates to a complete Bayesian analysis. In this simulation, a simple and intuitive approach to dealing with unmeasured confounding had unexpected effects, and served only to increase variability without any reduction in bias.

Thus, the research underscored the difficulty in accounting for unmeasured covariates even a simple epidemiological regression analysis.

References

1. Schlesselman, J. J. (1978). Assessing effects of confounding variables. *American Journal of Epidemiology*, *108*(1), 3–8. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/685974
2. Schneeweiss, S. (2006). Sensitivity analysis and external adjustment for unmeasured confounders in epidemiologic database studies of therapeutics {, (October 2005), 291–303.
3. Breslow, N.E. & Day, N.E. (1980). General Considerations for the Analysis of Case-Control Studies. (n.d.).
4. Vanderweele, T. J., & Arah, O. A. (2011). Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. *Epidemiology (Cambridge, Mass.)*, *22*(1), 42–52. http://doi.org/10.1097/EDE.0b013e3181f74493
5. Gran, J. M., Wasmuth, L., Amundsen, E. J., Lindqvist, B. H., & Aalen, O. O. (2008). Growth rates in epidemic models: application to a model for HIV/AIDS progression. *Statistics in Medicine*, *27(23)*(August 2006), 4817–4834. http://doi.org/10.1002/sim
6. Steenland, K., & Greenland, S. (2004). Monte Carlo sensitivity analysis and Bayesian analysis of smoking as an unmeasured confounder in a study of silica and lung cancer. *American Journal of Epidemiology*, *160*(4), 384–392. http://doi.org/10.1093/aje/kwh211
7. Phillips, C. V. (2003). Quantifying and reporting uncertainty from systematic errors. *Epidemiology*, *14*(4), 459–466. http://doi.org/10.1097/01.ede.0000072106.65262.ae
8. Dorie, V., Harada, M., Carnegie, N. B., & Hill, J. (2016). A flexible, interpretable framework for assessing sensitivity to unmeasured confounding. *Statistics in Medicine*, (October 2015). http://doi.org/10.1002/sim.6973

9. Gail, M. H., Wacholder, S., & Lubin, J. H. (1988). Indirect corrections for confounding under multiplicative and additive risk models. *Am J Ind Med*, *13*(1), 119–130. http://doi.org/10.1002/ajim.4700130108

10. Flanders, W., & Khoury, M. (1990). Indirect Assessment of Confounding: Graphic Description and Limits on Effect of Adjusting for Covariates. *Epidemiology*, *1*(3), 239–246. Retrieved from http://www.jstor.org/stable/25759804