



Published in final edited form as:

Epidemiology. 2019 November ; 30(6): 835–837. doi:10.1097/EDE.0000000000001064.

Mediational E-values: Approximate sensitivity analysis for unmeasured mediator–outcome confounding

Louisa H. Smith^a, Tyler J. VanderWeele^{a,b}

^aDepartment of Epidemiology, Harvard T.H. Chan School of Public Health

^bDepartment of Biostatistics, Harvard T.H. Chan School of Public Health

Abstract

Background: Mediation analysis is a powerful tool for understanding mechanisms, but conclusions about direct and indirect effects will be invalid if there is unmeasured confounding of the mediator–outcome relationship. Sensitivity analysis methods allow researchers to assess the extent of this bias but are not always used. One particularly straightforward technique that requires minimal assumptions is nonetheless difficult to interpret, and so would benefit from a more intuitive parameterization.

Methods: We conducted an exhaustive numerical search over simulated mediation effects, calculating the proportion of scenarios in which a bound for unmeasured mediator–outcome confounding held under an alternative parameterization.

Results: In over 99% of cases, the bound for the bias held when we described the strength of confounding directly via the confounder–mediator relationship instead of via the conditional exposure–confounder relationship.

Conclusions: Researchers can conduct sensitivity analysis using a method that describes the strength of the confounder–outcome relationship and the approximate strength of the confounder–mediator relationship that, together, would be required to explain away a direct or indirect effect.

Keywords

bias; mediation; unmeasured confounding; sensitivity analysis

Introduction

It is now widely recognized that valid mediation analyses require adjustment for mediator–outcome confounding.^{1,2} Failing to adjust for such confounders does not bias the total effect of the exposure, but the extent of mediation (i.e., the size of the indirect effect through the mediator relative to the direct effect through other pathways) may be over- or underestimated.^{1–3} However, it is not always the case that researchers collect data on these

Corresponding author: Louisa H. Smith, Department of Epidemiology, 677 Huntington Ave., Boston, MA 02115, Louisa_h_smith@g.harvard.edu, 207-653-4318.

Conflicts of interest: none declared

Data and computing code: Complete computing code is available in the supplementary material

confounders, particularly when mediation is secondary to the primary exposure–outcome analysis for which a study has been designed and for which confounders have been measured. Here we discuss a straightforward, although approximate, approach to sensitivity analysis for direct and indirect effects that essentially corresponds to the mediational equivalent to the “E-value”⁴ – a metric that was introduced as an approach to quantify robustness to confounding for total effects.

Recently a sensitivity analysis technique was proposed to examine the extent to which unmeasured confounding of the mediator–outcome relationship could explain observed mediation effects.^{1,2} Specifically, a bound for the bias can be used to assess the possible influence of an unmeasured confounder U on the observed direct or indirect effects (more formally, the natural direct and indirect effects; see eAppendix for counterfactual definitions). The bound describes the maximum extent to which unmeasured confounding by U could have resulted in the overestimation of the indirect effect and the corresponding underestimation of the direct effect, or vice versa.

Under the assumption that U is marginally independent of the exposure, A (see Figure), conditional on measured covariates C (omitted for simplicity from a probability statements that follow), the bound depends on the strength of the relationship of the unmeasured U with the outcome Y , specifically, the following risk ratio (RR):

$$RR_{UY | (A=1, M)} = \max_m \frac{\max_u \Pr(Y=1 | A=1, M=m, U=u)}{\min_u \Pr(Y=1 | A=1, M=m, U=u)},$$

as well as on the strength of the relationship between U and A that is induced within levels of mediator M , specifically,

$$RR_{AU | M} = \max_m \max_u \frac{\Pr(U=u | A=1, M=m)}{\Pr(U=u | A=0, M=m)}.$$

With these two parameters specified, the maximum ratio by which the estimate of the indirect effect (and correspondingly direct effect), on the RR scale, could differ from the true value was shown to be given by:

$$\frac{RR_{AU | M} \times RR_{UY | (A=1, M)}}{RR_{AU | M} + RR_{UY | (A=1, M)} - 1}.$$

The interpretation of the parameter $RR_{UY(A=1, M)}$ is relatively straightforward as the maximum direct effect, among the exposed, of U on Y not through M . However, the interpretation of the second parameter $RR_{AU | M}$ is more complex. That the $RR_{AU | M}$ parameter differs from 1 is due to so-called collider bias, which occurs when conditioning on the common effect (M) of two variables (A and U). However, the fact that A and U are not related except within levels of M makes the magnitude of this relationship difficult to intuit or speculate about.

Citing previous work by Greenland,³ the authors proposing this bound claimed that in most but not all situations, the relationship between U and M would be at least as strong as the conditional A - U relationship and could therefore potentially be used as its proxy in sensitivity analysis.¹ Specifically, instead of the parameter RR_{AUM} , defined as the maximum RR of A on some value of U within a level of M , the authors proposed using the maximum RR comparing two values of U on M within a level of A :

$$RR_{UM=1|A} = \max_a \frac{\max_u \Pr(M=1 | A=a, U=u)}{\min_u \Pr(M=1 | A=a, U=u)}.$$

This is a more intuitive parameter to specify as it directly reflects the strength of the confounder–mediator relationship, which is of course necessary for confounding to be present.

The purpose of the present study was to investigate to what extent this alternative parameter could be used to perform approximate instead of exact sensitivity analysis for natural direct and indirect effects.

METHODS

To evaluate the extent to which the bound for unmeasured confounding would hold if $RR_{UM=1|A}$ were used instead of RR_{AUM} , we conducted an exhaustive numerical search via simulation. Our main analysis searched over all possible probability combinations of binary exposure, mediator, outcome, and unmeasured confounder, without making functional form assumptions. We did this to ensure no restrictions on the size or direction of any interactions or on the prevalence of any of the variables.

Conditional probabilities for M , U , and Y were randomly drawn from a uniform (0,1) distribution to ensure that all possible probabilities were encountered. However, because allowing the probabilities to vary over any range results in many unrealistic situations, we also considered several possibly more plausible restrictions when generating the data, including following a log-linear model for the mediator and restricting interactions so that all effects were in the same direction.

After randomly generating the necessary probabilities to define the joint distribution of exposure, mediator, outcome, and confounder, we calculated the true direct effect on both the risk ratio and risk difference (RD) scales, as well as the effects that would be observed if the confounder were unmeasured. From these values we computed the bias as the ratio of observed to true direct effect on the RR scale and as the difference on the RD scale. We also computed the exact bound as defined above.

Next we computed the bound using $RR_{UM=1|A}$ instead of RR_{AUM} . We considered this alternative bound to be successful when it, like the true bound, was greater than the bias. We also “corrected” the observed effects using both the true and alternative bounds by dividing (for RRs) or subtracting (for RDs) the bound from the observed effect. When the alternative

bound failed to bound the bias, we computed the magnitude of the bias that would remain if we used it to correct the mediated effects anyway.

Because not only the bias parameters but also the bias itself is conditional on measured confounders that have been adjusted for in the estimation of the observed effects, we did not explicitly include C as one of the random variables in the numerical search. Instead, we assumed, as we do throughout the text, that we are working within strata of C .

More details can be found in section 2 of the eAppendix. R code can also be found in the eAppendix.

RESULTS

As expected, the bound computed with the appropriate parameters was always greater than or equal to the bias. We found that in 99.3% of simulation settings, the value calculated using the $RR_{UM=1|A}$ parameter did indeed also bound the bias (eTable 1). However, we did not discover any complete characterization that distinguished the failed bounds from the successes, although they occurred more frequently when the bias was larger (eTable 2). On average, the bound constructed with the alternative parameter was weaker than the true bound (eFigure 1), but overall effects corrected by both bounds had relatively similar distributions (eFigure 2). When the alternative bound did fail, the residual bias after correcting with it was generally small. Over 80% of the time, the corrected effect was less than 1.2 times as great as the true effect (eFigure 3). On the risk-difference scale, the remaining bias was less than 0.02 around 50% of the time that the bound failed (eTable 4).

We found similar results when we varied the data-generating distribution. While the distributions we assumed for the variables are unlikely to directly reflect reality, the few situations in which the bound failed were unremarkable and did not appear more likely to occur in practice than those in which the bound held. In section 4 of the eAppendix we consider continuous M and U as well as derive conditions under which other alternative bounds will hold.

DISCUSSION

These results support the claim that this sensitivity analysis technique can, at least roughly, be based on the proposed strength of the relationship between an unmeasured confounder and the mediator, which better matches intuition about the source and size of confounding. When unmeasured confounding is suspected, bounds can be computed by researchers or readers with values for $RR_{UM=1|A}$ that seem reasonable for a given situation in order to see how bias of that magnitude would affect the direct and indirect effects.

It is also possible to calculate the minimum size of the two parameters that would be necessary to completely explain away a direct or indirect effect. This is essentially a mediational analogue to the E-value to assess robustness to unmeasured confounding for total effects.⁴ For an observed natural direct or indirect effect risk ratio of magnitude RR^{obs} , this is given by:^{1,4}

$$\text{Mediational E-value} = RR^{obs} + \sqrt{RR^{obs} \times (RR^{obs} - 1)}$$

If both of the parameters are at least as large as the mediational E-value, conditional on the measured confounders, it is possible that unmeasured confounding is entirely responsible for a direct or indirect effect, and that effect is truly null. The mediational E-value expression applies exactly to the sensitivity analysis parameters $RR_{UY(A=1,M)}$ and RR_{AUM} and approximately, as above, to the parameters $RR_{UY(A=1,M)}$ and $RR_{UM=1|A}$.

As an example, Oberg et al.⁵ examined the indirect effect of fertility treatment on preterm birth mediated through multiple gestations (see Figure). They estimated an indirect effect risk ratio of 1.55 (95% CI: 1.52, 1.59). The mediational E-value for the estimate is 2.47, and for the limit of the CI closest to the null it is 2.41. We are then able to make statements of the form, “To completely explain away the observed indirect effect, an unmeasured confounder associated with both multiple gestations and preterm birth with approximate risk ratios of 2.47-fold each, above and beyond the measured covariates, could suffice, but weaker confounding could not. To shift the confidence interval to the null, an unmeasured confounder associated with both multiple gestations and preterm birth with approximate risk ratios of 2.41-fold each, above and beyond the measured covariates, could suffice, but weaker confounding could not.”

We hope that by simplifying an already straightforward sensitivity analysis method with a more easily specified parameter, researchers will increasingly assess the robustness of their mediation analyses to unmeasured confounding.

Supplementary Material


Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Sources of funding: This work was supported by grant R01CA222147 from the NIH (Tyler VanderWeele)

REFERENCES

1. Ding P, VanderWeele TJ. Sharp sensitivity bounds for mediation under unmeasured mediator–outcome confounding. *Biometrika*. 2016;103:483–490. [PubMed: 27279672]
2. VanderWeele TJ. *Explanation in causal inference: Methods for mediation and interaction*. 2015.
3. Greenland S Quantifying biases in causal models: Classical confounding vs collider-stratification bias. *Epidemiology*. 2003;14:300–306. [PubMed: 12859030]
4. VanderWeele TJ, Ding P. Sensitivity analysis in observational research: Introducing the E-value. *Ann Intern Med*. 2017;167:268–275. [PubMed: 28693043]
5. Oberg AS, VanderWeele TJ, Almqvist C, et al. Pregnancy complications following fertility treatment—disentangling the role of multiple gestation. *Int J Epidemiol*. 2018;47:1333–1342. [PubMed: 29939263]

**Figure.**

Directed acyclic graph representing a mediated effect subject to unmeasured mediator–outcome confounding. A represents the exposure of interest, M the mediator, and Y the outcome. C includes all measured confounders of both the exposure–outcome and mediator–outcome relationships, and U is an unmeasured confounder of the latter. In the example in the text, A is fertility treatment, M multiple gestations, and Y preterm birth. C includes all confounders the authors were able to measure and adjust for, and U the mediator–outcome confounders they were not able to measure.