

How likely is it that omitted variable bias will overturn your results?

Deepankar Basu*

August 20, 2024

Abstract

Building on a recently developed methodology for sensitivity analysis that parametrizes omitted variable bias in terms of partial R^2 measures, I propose a simple statistic to capture the severity of omitted variable bias in any observational study: the probability of omitted variable bias overturning the reported result. The central element of my proposal is formal covariate benchmarking, whereby researchers choose an observed regressor (or a group of observed regressors) to benchmark the relative strength of association of the omitted regressor with the outcome variable and with the treatment variable. These relative strengths of association function as the two sensitivity parameters of the analysis. By allowing these sensitivity parameters to take all permissible values, we get the most conservative estimate of the probability that omitted variable bias can overturn the reported results. By using absolute and relative limits on the maximum values of the sensitivity parameters based on institutional knowledge or other details of the particular study, a researcher can generate less conservative estimates of that probability. For empirical studies with relatively large number of regressors and sample sizes, I suggest bounds for the sensitivity parameters based on simulation studies. I illustrate the methodology using an empirical example that studies the effect of exposure to violence on attitudes towards peace.

Keywords: omitted variable bias; sensitivity analysis.

JEL Codes: C20.

1 Introduction

In many disciplines, like economics, epidemiology, political science, public health, sociology, etc., it is of utmost importance to estimate causal effects from observational data, e.g. the causal effect of years of schooling on wages (Card, 2001), the causal effect of class size on student scores (Angrist and Lavy, 1999), the causal effect of exposure to violence on attitudes towards peace (Cinelli and Hazlett, 2020), or the causal effect of breastfeeding on child outcomes (VanderWeele and Ding, 2017; Oster, 2019). In each of these cases, and in observational studies more generally, to distinguish between correlation and causation researchers need to take account of unmeasured confounders (omitted variables), i.e. unobserved variables that are correlated both with the treatment and outcome variables. When plausible instrumental variables are not available for treatment assignment, as is often the case in observational studies, or natural experiments

*Department of Economics, 310 Crotty Hall, University of Massachusetts Amherst, Amherst 01002, USA. Email: dbasu@econs.umass.edu

cannot be exploited to estimate causal effects, researchers can turn to sensitivity analysis to investigate the robustness of their reported results to omitted variable bias.

There is a large literature on sensitivity analysis that goes back to at least [Cornfield et al. \(1959\)](#), and has been continued in [Rosenbaum and Rubin \(1983\)](#); [Rosenbaum \(2002\)](#); [Imbens \(2003\)](#); [Imai et al. \(2010\)](#), and others. Among more recent proposals for sensitivity analysis ([Frank, 2000](#); [Krauth, 2016](#); [Ding and VanderWeele, 2016](#); [VanderWeele and Ding, 2017](#); [Oster, 2019](#); [Diegert et al., 2023](#)), a most innovative and promising approach has been presented by [Cinelli and Hazlett \(2020\)](#). The key novelty in this proposal involved re-parametrizing the traditional omitted variable bias expression using partial R^2 measures.¹ This re-parametrization has opened up a fruitful way to conduct sensitivity analysis about omitted variable bias.

The main contribution of this paper is to build on the methodology of [Cinelli and Hazlett \(2020\)](#) to develop a simple and intuitive statistic that helps answer the following question: how likely is it that omitted variable bias will overturn the baseline results reported for any observational study? Since a reported result for any study will be overturned if, once omitted variable bias has been taken into account, either the bias-adjusted estimate is statistically indistinguishable from zero or has reversed its sign (compared to the unadjusted estimate), I provide an estimate of this occurring. That is, my proposed methodology provides an estimate of the probability that omitted variable bias will overturn the baseline reported result.

In developing my methodology, I stress the importance of formal covariate benchmarking. In most observational studies the true effect of a treatment variable on an outcome variable is sought to be estimated. If the model has omitted variables, then the estimated (unadjusted) effect is different from the true effect—due to omitted variable bias. In such a context, formal covariate benchmarking involves the researcher choosing some observed regressor (or group of regressors) to use as a benchmark for the *relative* strength of association of the *residualized* omitted variable with both the treatment and outcome variables, where association is measured by total or partial R^2 measures.

By a residualized omitted variable, I refer to the part of the omitted variable that is not explained by a linear function of all the included regressors, i.e. it is the residual in a hypothetical regression of the omitted variable on all the included regressors (including a constant). Since this entity can always be defined, the method is very general. In the special case that the omitted variable is exogenous, it coincides with its residualized version.²

¹For a definition of the partial R^2 , see equation 1 below.

²In a recent contribution, [Diegert et al. \(2023\)](#) has highlighted some analytical problems of residualizing omitted variables for sensitivity analysis. The argument in [Diegert et al. \(2023\)](#) applies only to the methodology of [Oster \(2019\)](#). It does not apply to the methodology of [Cinelli and Hazlett \(2020\)](#) that I build on in this paper. The key difference between the approaches of [Cinelli and Hazlett \(2020\)](#) and [Diegert et al. \(2023\)](#) is that the former defines the sensitivity parameters symmetrically and the latter define it asymmetrically ([Diegert et al., 2023](#), appendix F.3). There is no obvious reason to choose the former over

Following Cinelli and Hazlett (2020), I use two parameters for sensitivity analysis that captures the relative strength of association of the omitted variable. The first is k_D , which measures the relative strength of association of the residualized omitted variable, compared to the benchmark covariate, with the treatment variable; the second is k_Y , which denotes the relative strength of association of the residualized omitted variable, compared to the benchmark covariate, with the outcome variable.

My proposed methodology works as follows. For any empirical study, the researcher chooses a benchmark covariate (or group of covariates) and a significance level, α , to test whether the true estimate is zero. Once the benchmark covariate has been chosen, we can compute maximum permissible values of the two sensitivity parameters, k_D and k_Y , and denote them as k_D^{\max} and k_Y^{\max} , respectively. These maximum permissible values are derived from the fact that some relevant R^2 measures must be bounded above by 1.

Next we define a bounded box on the (k_D, k_Y) plane, B, given by the Cartesian product of $0 \leq k_D \leq k_D^{\max}$ and $0 \leq k_Y \leq k_Y^{\max}$. At each point of B we compute the omitted variable bias and the standard error of the true estimate. We use these to compute a $100(1 - \alpha)\%$ bias-adjusted confidence interval for the true estimate.³

In the next step, we draw the following contour plots on the (k_D, k_Y) plane: (a) if the unadjusted estimate is positive, we draw the contour plot of the *lower boundary* of the $100(1 - \alpha)\%$ bias-adjusted confidence interval; (b) if the unadjusted estimate is negative, we draw the contour plot of the *upper boundary* of the $100(1 - \alpha)\%$ bias-adjusted confidence interval. In each contour plot, the contour line for 0 divides the contour plot area into two parts that are relevant for our analysis: one part of the contour plot area give combinations of k_D and k_Y where the lower boundary of the $100(1 - \alpha)\%$ bias-adjusted confidence interval is weakly negative, and another part where the lower boundary is strictly positive. If the lower boundary is strictly positive this implies that the whole confidence interval lies strictly to the right of zero; on the other hand, if the lower boundary is strictly negative this implies that the whole confidence interval lies strictly to the left of zero.

Now consider two cases. For the case where the unadjusted estimate is positive, the fraction of the contour plot area where the lower boundary is strictly positive represents the probability that the reported result cannot be overturned (because the $100(1 - \alpha)\%$ bias-adjusted confidence interval is wholly to the right of zero). Thus, 1 minus the fraction of the area where the lower boundary is strictly positive gives us the probability that omitted variable bias can overturn the reported result. For the case where the unadjusted

the latter.

³At any point k_D, k_Y , we need to compute two partial R^2 measures involving the omitted variable to compute the omitted variable bias. If either of these partial R^2 measures exceed 1, we discard that point as impermissible.

estimate is negative, by the same reasoning, 1 minus the fraction of the area where the upper boundary is strictly negative gives us the probability that omitted variable bias can overturn the reported result. When these probabilities are small, then researchers can conclude that omitted variable bias is unlikely to overturn reported results. If the probabilities are large, then researchers should draw the opposite conclusion.

In computing these probabilities, the upper bounds on k_D and k_Y play very important roles. If a researcher uses the full permissible range for k_D and k_Y , she would get the most conservative estimate of the probability that omitted variable bias can overturn reported results. In many cases, deeper knowledge of the institutional set up or context of the study can allow researchers to put some bounds on how large k_D and k_Y can be that are significantly lower than k_D^{\max} and k_Y^{\max} . Using these bounds, she can compute less conservative estimates of the probability that omitted variable bias will overturn reported results.

If the researcher has no basis to choose reasonable bounds on k_D and k_Y , then she can use the following simple rules of thumb (which I arrived at on the basis of simulation studies reported in section 6.5 below).

- For empirical studies with relatively large sample sizes ($N > 300$) and relatively large number of regressors ($k > 50$), reasonable bounds to use are $0.1k_D^{\max}$ and $0.1k_Y^{\max}$.
- For empirical studies with lower number of regressors, e.g. $k < 25$, the bounds depend on the relative magnitude of the variance of the error terms in the outcome and treatment equations.
 - If the variance of the error term in the treatment equation is significantly higher than the variance of error term in the outcome equation, then $0.07k_D^{\max}$ and $0.4k_Y^{\max}$ can be used as reasonable bounds.
 - If the variance of the error term in the treatment equation is significantly lower than the variance of error term in the outcome equation, then $0.4k_D^{\max}$ and $0.07k_Y^{\max}$ are reasonable bounds.

In implementing my methodology, I construct contour plots and compute the probability that omitted variable bias will overturn reported results using all three covariate benchmarking approaches discussed in the appendix of Cinelli and Hazlett (2020): total R^2 -based covariate benchmarking, and partial R^2 -based covariate benchmarking (with and without conditioning on the treatment variable).⁴ I illustrate my proposed methodology using the running example in Cinelli and Hazlett (2020) where exposure to violence is used to explain attitudes towards peace in the context of the civil war in Darfur.⁵

⁴The details are explained in section 5 below.

⁵The details can be found in section 6 below.

Some recent work on sensitivity analysis that are close to my paper are [Krauth \(2016\)](#), [VanderWeele and Ding \(2017\)](#), [Ding and VanderWeele \(2016\)](#) and [Oster \(2019\)](#). In [Krauth \(2016\)](#), the key sensitivity parameter, λ , does not have natural bounds. This is because λ is the *ratio* of two correlation coefficients, and it is difficult to restrict this ratio to a finite interval on the real line, as [Krauth \(2016, assumption 2\)](#) does, without additional assumptions. On the other hand, the methodology of [Oster \(2019\)](#) suffers from two problems: (a) that one of the key sensitivity parameters, δ , cannot be interpreted in a way that is useful for sensitivity analysis ([Cinelli and Hazlett, 2020](#), section 6.3) and (b) that an untenable assumption is required to construct the identified set ([De Luca et al., 2019b](#), p. 219). On the other hand, [Ding and VanderWeele \(2016\)](#) and [VanderWeele and Ding \(2017\)](#) do not use formal covariate benchmarking for sensitivity analysis.

The rest of the paper is organized as follows: in section 2, I present the basic set up and the key expression for bias, and discuss the rationale for formal covariate benchmarking; in section 3, I present some results about the total and partial R^2 that will be useful later in the paper; in section 4, I discuss some theoretical results about formal covariate benchmarking; in section 5, I discuss details of my proposal for sensitivity analysis; in section 6, I illustrate the methodology proposed by working through an empirical example; finally, I conclude in section 7. Proofs are collected in appendix A.

2 The Setup

2.1 What is partial R^2 ?

The concept of partial R^2 is not very widely used in econometrics.⁶ Since this concept is central to the analysis in this paper, I begin by discussing it briefly. The partial R^2 of the random variables Y and Z , conditional on a set of covariates, X , can be computed as follows ([Greene, 2012](#), section 3.4): (a) collect the vector of residuals from a regression of Y on X ; (b) collect the vector of residuals from a regression of Z on X ; (c) take the square of the correlation coefficient between the two vectors of residuals. This is the partial R^2 of the random variables Y and Z , conditional on a set of covariates, X . While this computation clarifies the *partialling out* involved in defining the partial R^2 , there are two different, more useful, ways to define it.⁷

The partial R^2 between the random variables Y and Z , conditional on a set of covariates, X , denoted by $R_{Y \sim Z|X}^2$, can also be defined in terms of the more familiar total R^2 (the coefficient of determination in a

⁶I could not find this concept being discussed in any one of the popular graduate-level textbooks on econometrics. The closest one comes is the discussion of the partial correlation coefficient in [Greene \(2012, section 3.4\)](#). The classic treatment of partial R^2 can be found in [Yule \(1911, chapter 12\)](#), or in later editions of the book, e.g. [Yule and Kendall \(1948, chapter 14\)](#).

⁷These definitions are given in [Cinelli and Hazlett \(2020, page 12\)](#).

regression), as follows:

$$R_{Y \sim Z|X}^2 = \frac{R_{Y \sim Z+X}^2 - R_{Y \sim X}^2}{1 - R_{Y \sim X}^2}, \quad (1)$$

where $R_{Y \sim Z+X}^2$ denote the total R^2 in the regression of Y on Z and X , and $R_{Y \sim X}^2$ denotes the total R^2 in a regression of Y on X . From this definition in (1) we get some more intuition about what the partial R^2 measures: it is the ratio of (a) the increment in the total R^2 when Z is added as a covariate to the regression of Y on X , and (b) the difference of the total R^2 of the regression of Y on X from 1. Since the total R^2 can be at most 1 and since the total R^2 always increases weakly with the addition of a regressor, the numerator can at most be as large as the denominator, with both the numerator and denominators being positive. Hence, it is immediately clear that the partial R^2 , like the total R^2 , must lie between 0 and 1.

There is yet another definition of the partial R^2 that is motivated by another consideration: to re-express the partial R^2 of two random variable by removing one random variable, or several random variables, from the conditioning set. Suppose, for concreteness, that we wish to express the partial R^2 of Y and Z , conditional on X and D , in terms of partial R^2 measures conditional only on X , i.e. we remove D from the conditioning set. This can be accomplished with the recursive definition of partial R^2 :

$$R_{Y \sim Z|D,X} = \frac{R_{Y \sim Z|X} - R_{Y \sim D|X} R_{D \sim Z|X}}{\sqrt{1 - R_{Y \sim D|X}^2} \sqrt{1 - R_{D \sim Z|X}^2}}. \quad (2)$$

Both (1) and (2) are useful for parametrizing omitted variable bias in terms of partial R^2 measures.

2.2 Omitted variable bias

Consider the linear regression of an outcome on a treatment, controlling for a set of covariates given by X and Z ,

$$Y = \hat{\tau}D + X\hat{\beta} + \hat{\gamma}Z + \hat{\varepsilon}_{\text{full}} \quad (3)$$

where Y is the $n \times 1$ vector of the outcome (dependent variable), X is the $n \times k$ matrix of observed covariates, including a constant, Z is the $n \times 1$ (unobserved) confounder vector, and all hat-quantities denote estimated (sample, and not population) quantities. Since Z is unobserved, the researcher cannot estimate (3) but is forced to estimate the following restricted regression

$$Y = \hat{\tau}_{\text{res}}D + X\hat{\beta}_{\text{res}} + \hat{\varepsilon}_{\text{res}} \quad (4)$$

Letting

$$\widehat{\text{bias}} = \hat{\tau}_{\text{res}} - \hat{\tau} \quad (5)$$

denote the bias of the treatment effect arising from the restricted model, [Cinelli and Hazlett \(2020, page 48\)](#) show, by combining the Frisch-Waugh-Lovell theorem and definitions of partial R^2 , that

$$|\widehat{\text{bias}}| = \widehat{\text{se}}(\hat{\tau}_{\text{res}}) \sqrt{\frac{\text{df} \times R_{Y \sim Z|D,X}^2 \times R_{D \sim Z|X}^2}{1 - R_{D \sim Z|X}^2}} \quad (6)$$

and the standard error of the true estimate, $\hat{\tau}$, is given by

$$\widehat{\text{se}}(\hat{\tau}) = \widehat{\text{se}}(\hat{\tau}_{\text{res}}) \sqrt{\frac{1 - R_{Y \sim Z|D,X}^2}{1 - R_{D \sim Z|X}^2} \times \frac{\text{df}}{\text{df} - 1}} \quad (7)$$

where ‘se’ denotes standard error, ‘df’ denotes the degrees of freedom of the restricted regression in (4), $\widehat{\text{se}}(\hat{\tau}_{\text{res}})$ denotes the standard error of the treatment effect in the restricted regression, $R_{Y \sim D,X}^2$ denotes the total R^2 from a regression of Y on D and X , $R_{D \sim Z|X}^2$ refers to the partial R^2 from a regression of D on Z conditioning on X and we assume that $0 \leq R_{D \sim Z|X}^2 < 1$ (to make sure we do not attempt to divide by zero).

2.3 The rationale for formal covariate benchmarking

If a researcher knew the values of $R_{Y \sim Z|D,X}^2$ and $R_{D \sim Z|X}^2$, then she could compute the bias and standard error using (6) and (7). She could then construct confidence intervals for the true parameter estimate, τ , with the correct coverage. But she cannot do so directly because $R_{Y \sim Z|D,X}^2$ and $R_{D \sim Z|X}^2$ involve the omitted (unobserved) regressor, Z . Hence, the researcher must compute these partial R^2 measures indirectly by using information about observed covariates. This is where formal covariate benchmarking comes in.

Let $Z^{\perp X}$ denote the part of the omitted variable that is not explained linearly by the set of included regressors. Suppose the researcher has chosen a benchmark covariate (or set of covariates). Now, following [Cinelli and Hazlett \(2020\)](#) let us introduce two *sensitivity parameters*, k_D and k_Y . The first parameter, k_D , captures the relative strength, in terms of total or partial R^2 , of $Z^{\perp X}$ in explaining variation in the treatment variable, as compared to the benchmark covariate (or set of covariates); the second parameter, k_Y , captures the relative strength, in terms of total or partial R^2 , of $Z^{\perp X}$ in explaining variation in the outcome variable. These two parameters, defined more precisely in section 4 below, capture the judgment of the researcher

based on her knowledge of the context of the research.

Before proceeding, let us note that omitted variable bias should be deemed problematic only when it reduces the absolute magnitude of the unadjusted estimate, i.e. if the unadjusted estimate is positive, taking account of the omitted variables reduces it, and if the unadjusted estimate is negative, taking account of the omitted variables increases it. This has the potential to make the ‘true’ estimate statistically indistinguishable from zero or even in reversing its sign. It is this situation that sensitivity analysis should be designed to address.⁸

Under the assumption that omitted variable bias reduces the absolute magnitude of the parameter estimate, once a benchmark covariate is chosen, we can compute the exact values of $R_{Y \sim Z|D,X}^2$ and $R_{D \sim Z|X}^2$ as functions of k_D and k_Y . Using these, we can then compute the omitted variable bias and standard error of the true estimate given in (6) and in (7), respectively. These can, in turn, be used to compute bias-adjusted confidence intervals for the ‘true’ estimate. Contour plots of the boundaries of the bias-adjusted confidence intervals can be used, as I explain in section 5 below, to generate an estimate of the probability that omitted variable bias can overturn reported results. Before presenting details of the computations of $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|D,X}^2$ in section 4, I need some initial results on total and partial R^2 .

3 A result about total and partial R^2

For any $n \times r$ matrix, W , let $P_W = W(W'W)^{-1}W'$, denote the $n \times n$ projection matrix that projects onto the column space of W ; let M^0 denote the $n \times n$ matrix that generates deviations from means when pre-multiplied to a n vector (Greene, 2012, page 978–79), i.e.,

$$M^0 = \left[I - \frac{1}{n} ii' \right],$$

where I is the identity matrix of dimension n and i denotes a column vector of 1s. Note that P_W and M^0 are symmetric and idempotent matrices (Greene, 2012, page 32, 979).

For a $n \times 1$ vector, Z , and a $n \times k$ matrix X , let $Z^{\perp X} = Z - P_X Z = (I - P_X) Z$ denote the $n \times 1$ vector of ordinary least squares (OLS) residuals obtained from a regression of Z on X , and consider the following

⁸When omitted variable bias increases the absolute magnitude of the unadjusted estimate, then that estimate can be used to bound the ‘true’ effect away from zero. In such cases, there is no need for any sensitivity analysis.

four regressions estimated by OLS:

$$Y \text{ on } X, Z \tag{8}$$

$$Y \text{ on } X \tag{9}$$

$$Y \text{ on } Z \tag{10}$$

$$Y \text{ on } Z^{\perp X} \tag{11}$$

Let $R_{Y \sim X+Z}^2$, $R_{Y \sim X}^2$, $R_{Y \sim Z}^2$, and $R_{Y \sim Z^{\perp X}}^2$, denote the total R-squared (coefficient of determination) for the regressions in (8), (9), (10), and (11), respectively; and let $W = (X : Z)$ denote the $n \times (k + 1)$ augmented matrix obtained by appending Z as an additional column to the matrix X . Using the definition of the R-squared (Greene, 2012, page 41), we have

$$R_{Y \sim X+Z}^2 = \frac{(P_W Y)' M^0 (P_W Y)}{Y' M^0 Y} = \frac{Y' P_W M^0 P_W Y}{Y' M^0 Y} \tag{12}$$

$$R_{Y \sim X}^2 = \frac{(P_X Y)' M^0 (P_X Y)}{Y' M^0 Y} = \frac{Y' P_X M^0 P_X Y}{Y' M^0 Y} \tag{13}$$

$$R_{Y \sim Z}^2 = \frac{(P_Z Y)' M^0 (P_Z Y)}{Y' M^0 Y} = \frac{Y' P_Z M^0 P_Z Y}{Y' M^0 Y} \tag{14}$$

$$R_{Y \sim Z^{\perp X}}^2 = \frac{(P_Z^{\perp X} Y)' M^0 (P_Z^{\perp X} Y)}{Y' M^0 Y} = \frac{Y' P_Z^{\perp X} M^0 P_Z^{\perp X} Y}{Y' M^0 Y} \tag{15}$$

where P_W, P_X, P_Z , and $P_{Z^{\perp X}}$ denote $n \times n$ projection matrices onto the column spaces of W, X, Z , and $Z^{\perp X}$ respectively. I will need a result on the decomposition of projection matrices that is given in Rao et al. (2008, page 323).

Lemma 1. *The projection matrix of W can be decomposed into two orthogonal projection matrices as:*

$$P_W = P_X + P_{Z^{\perp X}}, \text{ and } P_X P_{Z^{\perp X}} = 0. \tag{16}$$

Using lemma 1, we can prove the following result about the decomposition of the total R^2 .

Theorem 1. *For the regressions in (8), (9), (10), and (11), we have:*

$$R_{Y \sim X+Z}^2 = R_{Y \sim X}^2 + R_{Y \sim Z^{\perp X}}^2 = R_{Y \sim X+Z^{\perp X}}^2, \tag{17}$$

and

$$R_{Y \sim X+Z}^2 - R_{Y \sim X}^2 - R_{Y \sim Z}^2 = \eta_{Y,Z,X}, \quad (18)$$

where

$$\eta_{Y,Z,X} = R_{Y \sim Z^\perp X}^2 - R_{Y \sim Z}^2. \quad (19)$$

Theorem 1 shows that the total R^2 from a regression of Y on X and Z can be decomposed in two ways. In (17), it is decomposed into the total R^2 from a regression of Y on X , and the total R^2 from a regression of Y on $Z^\perp X$ (the part of Z that is orthogonal to X). In (18), it is decomposed into three terms: the total R^2 from a regression of Y on X , the total R^2 from a regression of Y and Z , and a remainder, $\eta_{Y,Z,X}$.

Corollary 1. *For the regressions in (8), (9), and (11), we have*

$$R_{Y \sim Z|X}^2 = \frac{R_{Y \sim X+Z}^2 - R_{Y \sim X}^2}{1 - R_{Y \sim X}^2} = \frac{R_{Y \sim X+Z^\perp X}^2 - R_{Y \sim X}^2}{1 - R_{Y \sim X}^2} = R_{Y \sim Z^\perp X|X}^2 \quad (20)$$

Proof. From (17), we have $R_{Y \sim X+Z}^2 = R_{Y \sim X}^2 + R_{Y \sim Z^\perp X}^2$. The right hand side is equal to the total R^2 from a regression of Y on X and $Z^\perp X$ because $Z^\perp X$ is orthogonal to X by construction. Hence, $R_{Y \sim X+Z}^2 = R_{Y \sim X+Z^\perp X}^2$. Using the definition of partial R^2 of Y and Z conditional on X , we have

$$R_{Y \sim Z|X}^2 = \frac{R_{Y \sim X+Z}^2 - R_{Y \sim X}^2}{1 - R_{Y \sim X}^2}$$

Now replacing $R_{Y \sim X+Z}^2$ with $R_{Y \sim X+Z^\perp X}^2$, we have the desired result. \square

4 Some theoretical results about formal covariate benchmarking

In this section, we see how to compute $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|D,X}^2$ once a benchmark covariate and specific values of the sensitivity parameters, k_D and k_Y are chosen.

4.1 Total R^2 -based covariate benchmarking

Suppose there are j observed covariates, $\{X_1, X_2, \dots, X_j\}$, and the researcher wishes to use the j -th one, X_j , for covariate benchmarking.

Proposition 1. *Let*

$$k_D := \frac{R_{D \sim Z^\perp X}^2}{R_{D \sim X_j}^2}, \quad k_Y := \frac{R_{Y \sim Z^\perp X}^2}{R_{Y \sim X_j}^2}; \quad (21)$$

then, we have

$$R_{D \sim Z|X}^2 = \frac{k_D R_{D \sim X_j}^2}{1 - R_{D \sim X}^2}, \quad R_{Y \sim Z|X}^2 = \frac{k_Y R_{Y \sim X_j}^2}{1 - R_{Y \sim X}^2}. \quad (22)$$

Proof. We first note, using corollary 1, that $R_{D \sim Z|X}^2 = R_{D \sim Z^{\perp X}|X}^2$. But

$$R_{D \sim Z^{\perp X}|X}^2 = \frac{R_{D \sim X + Z^{\perp X}}^2 - R_{D \sim X}^2}{1 - R_{D \sim X}^2} = \frac{R_{D \sim Z^{\perp X}}^2}{1 - R_{D \sim X}^2}.$$

Now using k_D defined in (21), we get the first part of (22). A similar argument establishes the second part of (22). \square

Assuming that omitted variable bias reduces the absolute magnitude of the unadjusted estimate, Cinelli and Hazlett (2020, appendix, page 34) show that

$$R_{Y \sim Z|D,X}^2 = \frac{(|R_{Y \sim Z|X}| - |R_{Y \sim D|X}| |R_{D \sim Z|X}|)^2}{(1 - R_{Y \sim D|X}^2)(1 - R_{D \sim Z|X}^2)}. \quad (23)$$

4.2 Partial R^2 -based covariate benchmarking without conditioning on the treatment variable

We would like, as in the total R^2 case, to generate exact expressions (or upper bounds) for $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|X}^2$. Suppose, like before, there are j covariates, $\{X_1, X_2, \dots, X_j\}$, and the researcher wishes to use the j -th observed covariate, X_j , for benchmarking. Let X_{-j} refer to the set of observed covariates that is *not* used for benchmarking.

Proposition 2. *Let*

$$k_D := \frac{R_{D \sim Z^{\perp X}|X_{-j}}^2}{R_{D \sim X_j|X_{-j}}^2}, \quad k_Y := \frac{R_{Y \sim Z^{\perp X}|X_{-j}}^2}{R_{Y \sim X_j|X_{-j}}^2}; \quad (24)$$

then, we have

$$R_{D \sim Z|X}^2 = \frac{k_D R_{D \sim X_j|X_{-j}}^2}{1 - R_{D \sim X_j|X_{-j}}^2}, \quad R_{Y \sim Z|X}^2 = \frac{k_Y R_{Y \sim X_j|X_{-j}}^2}{1 - R_{Y \sim X_j|X_{-j}}^2} \quad (25)$$

Proof. If we insert $Z^{\perp X}$ in place of Z in the argument in Cinelli and Hazlett (2020, appendix B.2, p. 35), and use corollary 1, we get

$$R_{D \sim Z|X}^2 = R_{D \sim Z^{\perp X}|X}^2 = \frac{k_D R_{D \sim X_j|X_{-j}}^2}{1 - R_{D \sim X_j|X_{-j}}^2},$$

where k_D is defined in (24). This establishes the first part of (25). A similar argument establishes the second part. \square

Just like in the case of total R^2 -based covariate benchmarking, if taking account of the omitted variables reduces the absolute magnitude of the coefficient estimate, then we have

$$R_{Y \sim Z|D,X}^2 = \frac{(|R_{Y \sim Z|X}| - |R_{Y \sim D|X}| |R_{D \sim Z|X}|)^2}{(1 - R_{Y \sim D|X}^2)(1 - R_{D \sim Z|X}^2)}. \quad (26)$$

4.3 Partial R^2 -based covariate benchmarking by conditioning on the treatment variable

This variant of partial R^2 -based covariate benchmarking differs from the previous case only in how it defines k_Y (the definition of k_D remains unchanged).

Proposition 3. *Let*

$$k_D := \frac{R_{D \sim Z^{\perp X}|X_{-j}}^2}{R_{D \sim X_j|X_{-j}}^2}, \quad k_Y := \frac{R_{Y \sim Z^{\perp X}|X_{-j},D}^2}{R_{Y \sim X_j|X_{-j},D}^2}; \quad (27)$$

then, we have

$$R_{D \sim Z|X}^2 = \frac{k_D R_{D \sim X_j|X_{-j}}^2}{1 - R_{D \sim X_j|X_{-j}}^2}, \quad R_{Y \sim Z|D,X}^2 = \eta^2 f_{Y \sim X_j|X_{-j},D}^2 \quad (28)$$

where

$$\eta = \frac{\sqrt{k_Y} + |f_{k_D} \times f_{D \sim X_j|X_{-j}}|}{\sqrt{1 - f_{k_D}^2 \times f_{D \sim X_j|X_{-j}}^2}} \quad (29)$$

and

$$f_{k_D} = \frac{\sqrt{k_D R_{D \sim X_j|X_{-j}}^2}}{\sqrt{1 - k_D R_{D \sim X_j|X_{-j}}^2}} \quad (30)$$

Proof. The proof of the first part of (28) is the same as in proposition 2; for the second part, see Cinelli and Hazlett (2020, appendix B.2, p. 37–38). \square

5 Main proposal for sensitivity analysis

After computing the crucial partial R^2 measures, $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|D,X}^2$, researchers can use them to report three sets of results about sensitivity analysis of omitted variable bias: (a) bias-adjusted confidence intervals for *specific values* of the sensitivity parameters; (b) contour plot of the *relevant* boundary of bias-adjusted confidence intervals in the *whole permissible region* of the k_D and k_Y space; and (c) probability that reported results will be overturned once omitted variable bias is taken into account.

5.1 Bias-adjusted confidence interval for specific k_D, k_Y values

The first set of results that researchers could report are bias-adjusted estimates and bias-adjusted confidence intervals for some specific values of the sensitivity parameters, e.g. $k_D = k_Y = 1$, $k_D = k_Y = 2$, and $k_D = k_Y = 3$. These results would show whether a residualized omitted variable that is equally important in determining variation in the outcome and treatment measured with total or partial R^2 measures contains zero if the strength of association between the residualized omitted variable and the outcome/treatment is *exactly* 1, 2 or 3 times as high as the corresponding association between the benchmark covariate and the outcome/treatment.

Bias-adjusted estimates of the true treatment effect and bias-adjusted confidence intervals for the true treatment effect are easy to compute once a benchmark covariate and specific values of k_D and k_Y are chosen. For instance, for the case of total R^2 -based covariate benchmarking, specific values of k_D and k_Y determine magnitudes of $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|X}^2$ according to (22) and these, in turn, are used to determine $R_{Y \sim Z|D,X}^2$ according to (23). These partial R^2 measures can then be used to compute the absolute bias according to (6) and the standard error of the true estimate according to (7). These quantities can be then used to construct the bias-adjusted estimate of the treatment effect and the $100 \times (1 - \alpha)\%$ bias-adjusted confidence interval for the true estimate as follows:

- if the unadjusted estimate is positive, then an adverse omitted variable bias would reduce the unadjusted estimate; hence, the bias-adjusted estimate of the treatment effect is given by $\hat{\tau}_{\text{res}} - \widehat{|\text{bias}|}$, and the $100 \times (1 - \alpha)\%$ bias-adjusted confidence interval for the ‘true’ estimate is

$$(\hat{\tau}_{\text{res}} - \widehat{|\text{bias}|} - |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau}), \hat{\tau}_{\text{res}} - \widehat{|\text{bias}|} + |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau})), \quad (31)$$

- if the unadjusted estimate is negative, then an adverse omitted variable bias would increase the unadjusted estimate; hence, the bias-adjusted estimate of the treatment effect is given by $\hat{\tau}_{\text{res}} + \widehat{|\text{bias}|}$, and the $100 \times (1 - \alpha)\%$ bias-adjusted confidence interval for the ‘true’ estimate is

$$(\hat{\tau}_{\text{res}} + \widehat{|\text{bias}|} - |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau}), \hat{\tau}_{\text{res}} + \widehat{|\text{bias}|} + |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau})), \quad (32)$$

where $|t_{\alpha/2}^*|$ is the absolute magnitude of the critical value. To see this note that since $\hat{\tau}/\widehat{\text{se}}(\hat{\tau})$ is distributed

as a t random variable with df degrees of freedom,

$$P\left(\hat{\tau} - |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau}), \hat{\tau} + |t_{\alpha/2}^*| \widehat{\text{se}}(\hat{\tau})\right) = 1 - \alpha.$$

But $\hat{\tau} = \widehat{\tau}_{\text{res}} - \widehat{\text{bias}}$. Hence, under the assumption of adverse omitted variable bias, if $\widehat{\tau}_{\text{res}} > 0$, $\hat{\tau} = \widehat{\tau}_{\text{res}} - |\widehat{\text{bias}}|$; and if $\widehat{\tau}_{\text{res}} < 0$, $\hat{\tau} = \widehat{\tau}_{\text{res}} + |\widehat{\text{bias}}|$. This delivers the above confidence intervals.⁹

Construction of bias-adjusted confidence intervals rely on researchers choosing *exact* values of the sensitivity parameters k_D and k_Y . It is unlikely that researchers will know these values exactly for a given empirical study. Rather it might be more reasonable, and less demanding, for a researcher to be able to know about *ranges* for k_D and k_Y than for specific values. Moreover, using ranges rather than exact values of the sensitivity parameters make the results more robust and reliable. The next two set of results work with ranges of values for k_D and k_Y .

5.2 Contour plots in (k_D, k_Y) space

One important limitation of computing bias-adjusted confidence intervals for a few specific values of k_D and k_Y is that this does not give us the complete picture. They do not tell us how the magnitude of omitted variable bias varies for *all possible* values of the sensitivity parameters, k_D and k_Y . One way to summarize the relevant information for the full range of sensitivity parameters is to report contour plots of the *relevant* boundary of bias-adjusted confidence intervals in k_D and k_Y space.

By relevant, I mean the following: when the unadjusted estimate is positive, they should report a contour plot of the lower boundary of the 95% confidence interval, i.e. a contour plots of $\widehat{\tau}_{\text{res}} - |\widehat{\text{bias}}| - |t_{\alpha/2}^*| \widehat{\text{se}}_{\widehat{\tau}_{\text{res}}}$; on the other hand, when the unadjusted estimate is negative, they should report a contour plot of the upper boundary of the 95% confidence interval, i.e. a contour plot of $\widehat{\tau}_{\text{res}} + |\widehat{\text{bias}}| + |t_{\alpha/2}^*| \widehat{\text{se}}_{\widehat{\tau}_{\text{res}}}$. The relevant level to look at, in both cases, is the contour curve at zero. This gives us the combinations of the sensitivity parameters k_D and k_Y which entail zero being included in the confidence interval.

The construction of the contour plots in (k_D, k_Y) space is helped by the fact that we can put bounds on both k_D and k_Y . The results in section 4 allow us to derive such bounds. In each case, the contour plots will be created on the bounded box created by the Cartesian product of $0 \leq k_D \leq k_D^{\max}$ and $0 \leq k_Y \leq k_Y^{\max}$.

⁹Similar arguments can be used to compute bias-adjusted confidence intervals for the two partial R^2 -based covariate benchmarking approaches. For partial R^2 -based covariate benchmarking without conditioning on treatment, $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|X}^2$ can be computed via proposition 2; for partial R^2 -based covariate benchmarking with conditioning on treatment, $R_{D \sim Z|X}^2$ and $R_{Y \sim Z|X}^2$ can be computed via proposition 3.

For total R^2 -based covariate benchmarking,

$$k_D^{\max} = \frac{1 - R_{D \sim X}^2}{R_{D \sim X_j}^2}, k_Y^{\max} = \frac{1 - R_{Y \sim X}^2}{R_{Y \sim X_j}^2}. \quad (33)$$

These bounds are arrived at from proposition 1 using the fact that $0 \leq R_{D \sim Z|X}^2 \leq 1$ and $0 \leq R_{Y \sim Z|X}^2 \leq 1$.

For partial R^2 -based covariate benchmarking without conditioning on the treatment variable,

$$k_D^{\max} = \frac{1 - R_{D \sim X_j|X_{-j}}^2}{R_{D \sim X_j|X_{-j}}^2}, k_Y^{\max} = \frac{1 - R_{Y \sim X_j|X_{-j}}^2}{R_{Y \sim X_j|X_{-j}}^2}. \quad (34)$$

These bounds are arrived at from proposition 2 using the fact, once again, that $0 \leq R_{D \sim Z|X}^2 \leq 1$ and $0 \leq R_{Y \sim Z|X}^2 \leq 1$. For partial R^2 -based covariate benchmarking with conditioning on the treatment variable,

$$k_D^{\max} = \frac{1 - R_{D \sim X_j|X_{-j}}^2}{R_{D \sim X_j|X_{-j}}^2} \quad (35)$$

and

$$k_Y^{\max} = \left(\frac{\sqrt{1 - f_{k_D}^2 \times f_{D \sim X_j|X_{-j}}^2}}{f_{Y \sim X_j|X_{-j}, D}} - |f_{k_D} \times f_{D \sim X_j|X_{-j}}| \right)^2. \quad (36)$$

These bounds are arrived at from proposition 3 and using the fact, once again, that $0 \leq R_{D \sim Z|X}^2 \leq 1$ and $0 \leq R_{Y \sim Z|D, X}^2 \leq 1$.

While the contour plots are useful in displaying how the lower (or upper) boundary of the bias-adjusted confidence interval varies across the whole possible space of the sensitivity parameters, k_D and k_Y , we can use them, as I discuss next, to compute an extremely interesting statistic: the probability that omitted variable bias can overturn the reported result in any specific study.

5.3 Probability of reported results being overturned

There are two cases to consider: (a) when the reported (or unadjusted) estimate is positive, and (b) when the reported (or unadjusted) estimate is negative.

5.3.1 Unadjusted estimate is positive

When the unadjusted estimate is positive and the researcher is using a significance level of α (for implicit hypothesis tests of the null hypothesis that the true estimate is zero), then omitted variable bias *cannot* overturn the result as long as the $100 \times (1 - \alpha)\%$ bias-adjusted confidence interval is strictly to the right of

zero. The $100*(1 - \alpha)\%$ bias-adjusted confidence interval being strictly to the right of zero is equivalent to the lower boundary of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval being strictly positive. The probability of the lower boundary of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval being strictly positive is the *fraction* of the area of the contour plot where the ‘level’ of the plot is strictly positive. This fraction is the relevant probability in a straightforward frequentist sense because the numerator counts all possible combinations of k_D and k_Y where the lower boundary of the bias-adjusted confidence interval is strictly to the right of zero and the denominator counts all possible combinations of k_D and k_Y that are permissible.

Note that we can compute this fraction easily because the contour plot lies within the *finite* bounded box given by the Cartesian product of $0 \leq k_D \leq k_D^{\max}$ and $0 \leq k_Y \leq k_Y^{\max}$. Hence the total area of the contour plot is always positive and finite. Thus, once the contour plots are created, we can compute this probability, and subtract it from 1 to get the probability of omitted variable bias overturning reported results, for the full range of permissible variation in k_D and k_Y or for subsets therein (as I do in the example in section 6 below; for instance, see Table 2).

The computation of the relevant fraction of the area of the contour plot is implemented by a grid search method. I choose a value of N , e.g. $N = 100$. Then, I choose N equally-spaced points in the interval $0 \leq k_D \leq k_D^{\max}$; similarly, I choose N equally-spaced points in the interval $0 \leq k_Y \leq k_Y^{\max}$. For each of the N^2 points (k_D, k_Y) formed by combinations of points from these two intervals, I test whether: (a) the value of $R_{Y \sim Z|D,X}^2$ is weakly less than 1; and (b) the lower boundary of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval is strictly positive. Suppose (a) is true at M_1 points (these are all the valid points on the contour plot) and (b) is true in M_2 points (these are the points where omitted variable bias cannot overturn the reported result). Then, the probability that omitted variable bias can overturn the reported results is given by

$$1 - \left(\frac{\text{area where lower boundary of confidence interval is strictly positive}}{\text{valid area of contour area}} \right) = 1 - \left(\frac{M_2}{M_1} \right).$$

5.3.2 Unadjusted estimate is negative

When the unadjusted estimate is negative and the researcher is using a significance level of α , then omitted variable bias cannot overturn the reported result as long as the $100*(1 - \alpha)\%$ bias-adjusted confidence interval is strictly to the left of zero. This is equivalent to the upper boundary of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval being strictly negative.

By the same argument as above, the probability of the lower bound of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval being strictly negative is the fraction of the area of the contour plot of the lower boundary

of $100*(1-\alpha)\%$ bias-adjusted confidence interval in k_D and k_Y space where the level of the function is strictly negative. Once the contour plots are created, we can compute this probability, and subtract it from 1 to get the probability of omitted variable bias overturning reported results, for the full range of permissible variation in k_D and k_Y or for subsets therein. The computation of the relevant fraction of the contour plot is implemented, as above, through a grid search method so that the probability that omitted variable bias can overturn the reported results is given by

$$1 - \left(\frac{\text{area where lower boundary of confidence interval is strictly positive}}{\text{valid area of contour area}} \right) = 1 - \left(\frac{M_2}{M_1} \right).$$

where, just like in the other case, on M_1 points on the grid $R_{Y \sim Z|D,X}^2 \leq 1$, and on M_2 points of the grid, the upper boundary of confidence interval is strictly negative.

6 Empirical example

6.1 Set up

I take up the running empirical example discussed in [Cinelli and Hazlett \(2020\)](#), where the effect of exposure to violence on attitudes towards peace is studied.¹⁰ To investigate this question, a researcher estimates the following regression model with individual-level data:

$$\text{PeaceIndex}_i = \beta_0 + \beta_1 \text{DirectHarm}_i + \text{Controls}_i + \varepsilon_i, \tag{37}$$

where ‘PeaceIndex’ is an index to measure attitudes towards peace efforts, ‘DirectHarm’ measures the exposure to violence. The control variables are: gender of the individual, age, whether the individual was a farmer, herder, merchant or trader, household size, whether or not the individual voted in the past, and village-level fixed effects.

It is suspected that there is at least one important unobserved confounder, Z , (e.g. wealth of the individual) that is correlated both with exposure to violence (treatment) and attitude towards peace (outcome). The researcher does not have data on wealth of individuals. Hence, the researcher wishes to conduct sensitivity analysis regarding the possible effect of this omitted variable. Suppose, finally, the researcher knows, on the bases of domain knowledge, that gender of the individual is one of the, if not the, most important

¹⁰For further details, see [Cinelli and Hazlett \(2020\)](#), section 2). I accessed the data set used for the analysis in the R package `sensemakr`.

variables determining treatment (exposure to violence measured by ‘DirectHarm’), and that gender is also correlated with the outcome (attitude towards peace measured by ‘PeaceIndex’). Hence, she chooses to use this variable for formal covariate benchmarking to conduct sensitivity analysis of her results to omitted variable bias.

6.2 Bias-adjusted estimates and confidence intervals

In table 1, I report the estimate of β_1 (the coefficient on ‘DirectHarm’ in equation 37), the bias-adjusted estimate of β_1 , the bias-adjusted standard error of β_1 and the lower and upper bounds of the bias-adjusted 95% confidence interval for β_1 based on total R^2 -based benchmarking, partial R^2 -based benchmarking (without conditioning on the treatment variable) and partial R^2 -based benchmarking (with conditioning on the treatment variable). I present results for two combinations of the sensitivity parameters: panel A uses $k_D = k_Y = 1$, and panel B uses $k_D = k_Y = 3$.

From Table 1, we see that the point estimate is 0.097. The bias adjusted estimate in panel A ($k_D = k_Y = 1$) is: 0.068, 0.059 and 0.058 for total R^2 -based benchmarking, partial R^2 -based benchmarking (without conditioning on the treatment variable) and partial R^2 -based benchmarking (with conditioning on the treatment variable), respectively. In each of these three cases, the 95% confidence interval does not contain zero even though the lower and upper bounds of the confidence intervals are slightly different (the main difference is between total and partial R^2 -based covariate benchmarking).

When we turn to panel B ($k_D = k_Y = 3$), we see the same pattern as in panel A though there are some important differences as well. The magnitude of the bias-adjusted estimate reduces to 0.046, 0.030 and 0.030 respectively, for the three cases. The 95% confidence intervals are now qualitatively different between total and partial R^2 -based covariate benchmarking (though there is not much difference between the two variants of partial R^2 -based covariate benchmarking). For total R^2 -based benchmarking, the confidence interval is (0.024, 0.067); for partial R^2 -based benchmarking (without conditioning on the treatment variable), the confidence interval is (−0.007, 0.066); and for the other variants of partial R^2 -based covariate benchmarking, the confidence interval is similar at (−0.006, 0.067). Thus, if a researcher used total R^2 -based benchmarking, she would conclude that taking account of omitted variable bias and estimation uncertainty would still keep the estimate away from zero; if she used either variant of partial R^2 -based benchmarking, she would reach the opposite conclusion that taking account of omitted variable bias along with estimation uncertainty would make the estimate statistically indistinguishable from zero.

Table 1: Omitted variable bias-adjusted estimate, standard error and 95% confidence intervals

	Panel A: $k_D = 1, k_Y = 1$			Panel B: $k_D = 3, k_Y = 3$		
	Total	Partial 1	Partial 2	Total	Partial 1	Partial 2
$R_{Y \sim D X}^2$	0.022	0.022	0.022	0.022	0.022	0.022
$R_{D \sim Z X}^2$	0.008	0.027	0.027	0.008	0.027	0.027
$R_{Y \sim Z D,X}^2$	0.255	0.120	0.126	0.781	0.381	0.374
Estimate	0.097	0.097	0.097	0.097	0.097	0.097
Bias-Adjusted Estimate	0.068	0.059	0.058	0.046	0.030	0.030
Bias-Adjusted Standard Error	0.020	0.022	0.022	0.011	0.019	0.019
Bias-Adjusted Lower Boundary of Conf. Int.	0.028	0.016	0.015	0.024	-0.007	-0.006
Bias-Adjusted Upper Boundary of Conf. Int.	0.107	0.103	0.102	0.067	0.066	0.067

Notes: This table presents omitted variable bias-adjusted estimate, standard error and 95% confidence intervals for two combinations of the sensitivity parameters: (a) panel A: $k_D = 1, k_Y = 1$, and (b) panel B: $k_D = 3, k_Y = 3$. Total: total R^2 -based benchmarking; Partial 1: partial R^2 -based benchmarking (without conditioning on treatment); Partial 2: partial R^2 -based benchmarking (with conditioning on treatment).

6.3 Contour plots in (k_D, k_Y) space

Figure 1, 2 and 3 present contour plots based on total R^2 -based benchmarking, partial R^2 -based benchmarking (without conditioning on the treatment variable) and partial R^2 -based benchmarking (with conditioning on the treatment variable), respectively. In each figure, I present the contour plot of the lower boundary of the 95% confidence interval (because the unadjusted parameter estimate is positive at 0.097). In these contour plots, the most relevant level to look at is zero. Each contour plot is constructed in (k_D, k_Y) space in the bounded box given by the Cartesian product of $0 \leq k_D \leq k_D^{\max}$ and $0 \leq k_Y \leq k_Y^{\max}$. Specific values of k_D^{\max} and k_Y^{\max} can be found as memo items in Table 2.

Looking at Figure 1, 2 and 3, we should read the yellow region as showing all combinations of k_D and k_Y , where the 95% bias-adjusted confidence interval for the ‘true’ estimate would exclude zero (because the lower boundary point of the interval is strictly positive). Thus, these are all the possible combinations of k_D and k_Y , i.e. all the possible strengths of the residualized omitted variable, that would still not make the bias-adjusted estimate zero even after we account for estimation uncertainty.

Why are the contour plots cut-off towards the north-east region? This is because the combinations of k_D and k_Y in the north-east region lead to $R_{Y \sim Z|D,X}^2 > 1$. Since R^2 cannot exceed unity, these combinations of k_D and k_Y are impermissible. Hence, they are dropped from the contour plots—giving rise to the empty north-east regions.

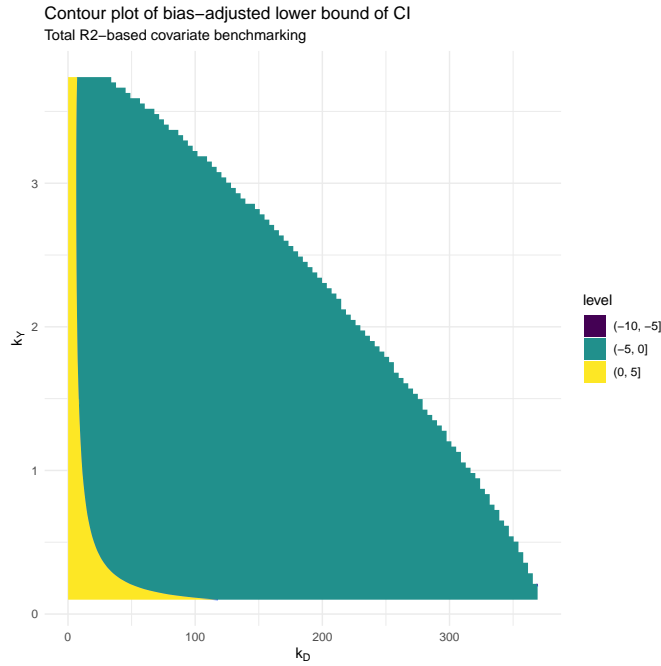


Figure 1: Contour plot for the the omitted variable bias-adjusted estimate and the lower bound of the bias-adjusted 95% confidence interval. Bounds have been estimated with total R^2 -based covariate benchmarking.

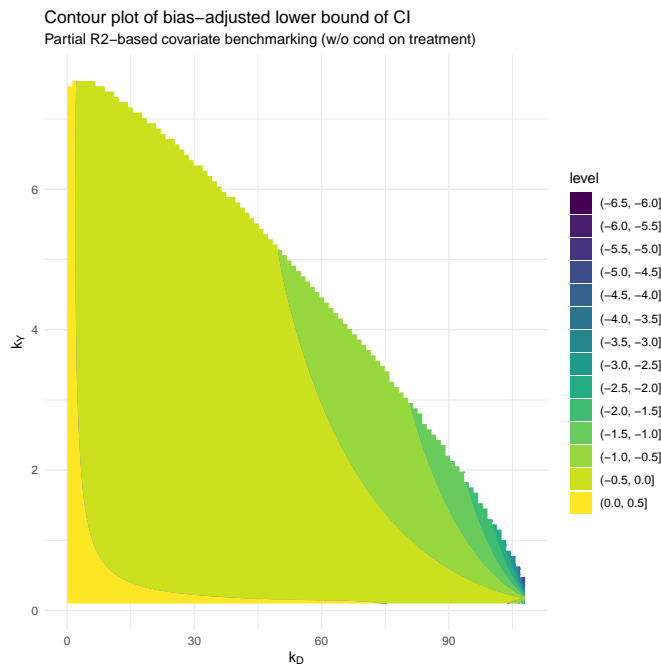


Figure 2: Contour plot for the the omitted variable bias-adjusted estimate and the lower bound of the bias-adjusted 95% confidence interval. Bounds have been estimated with partial R^2 -based covariate benchmarking without conditioning on the treatment.




Figure 3: Contour plot for the the omitted variable bias-adjusted estimate and the lower bound of the bias-adjusted 95% confidence interval. Bounds have been estimated with partial R^2 -based covariate benchmarking with conditioning on the treatment.

6.4 Probability that reported results can be overturned

In Table 2, I present results that is pertinent to answering the following question: what is the probability that omitted variable bias can overturn the reported results of this study that the estimate of exposure to violence on attitudes towards peace is positive (the coefficient estimate is 0.097)? The answer is found by computing the fraction of the area of the contour plot where the lower boundary of the 95% confidence interval is strictly positive. For total R^2 -based covariate benchmarking, the relevant contour plot is given in Figure 1; for partial R^2 -based covariate benchmarking without conditioning on the treatment variable, the relevant contour plot is given in Figure 2; and for partial R^2 -based covariate benchmarking with conditioning on the treatment variable, the relevant contour plot is given in Figure 3.

To compute the probabilities reported in Table 2, I use a grid search procedure with $N = 100$. Thus, in each case, I choose 100 equally-spaced points in the relevant intervals on the k_D and k_Y axes. For each different specification reported in Table 2, the relevant intervals are given in the first column of Table 2. For instance, for the most conservative estimates (panel A), the relevant intervals are: $0 \leq k_D \leq k_D^{\max}$ and $0 \leq k_Y \leq k_Y^{\max}$, with values of k_D^{\max} and k_Y^{\max} given as memo items in Table 2.

For each of the 10000 points (k_D, k_Y) on the grid, I test whether: (a) the value of $R_{Y \sim Z|D, X}^2 \leq 1$; and (b)

the lower boundary of the $100*(1 - \alpha)\%$ bias-adjusted confidence interval is strictly positive. Suppose the first condition is true at M_1 points and the second is true at M_2 points. Then, the probability that omitted variable bias can overturn the reported results is given by $1 - (M_2/M_1)$.

The results in panel panel A of Table 2 show that the probability of reported results being overturned is high and lies between 0.93 and 0.96, varying by the type of covariate benchmarking that is used. This means that if the researcher undertaking this study of the effect of violence on attitudes towards peace had no information to limit the variation in the sensitivity parameters k_D and k_Y , and had to perform use the *full range* of these sensitivity parameters for the analysis, then they would have to conclude that there is very high chance that their reported result (that the estimate is 0.097) can be overturned by omitted variable bias. This is the most conservative estimate of the probability that omitted variable bias can overturn their result.

In many cases, researchers might be able, on the basis of detailed knowledge of the context of the study, to limit the range of the sensitivity parameters k_D and k_Y . There are at least two ways to do this using either *absolute bounds* or *relative bounds* on the maximum value of the sensitivity parameters k_D and k_Y . When we use an absolute bound, we choose two positive real numbers, c_1, c_2 , and work with the intervals $0 \leq k_D \leq c_1$ and $0 \leq k_Y \leq c_2$; when we use a relative bound, we choose two positive fractions d_1, d_2 , and work with the intervals $0 \leq k_D \leq d_1 k_D^{\max}$ and $0 \leq k_Y \leq d_2 k_Y^{\max}$. To compute the probability that omitted variable bias can overturn reported results, I implement the same grid search method discussed above.

In panel B, I report results from using absolute bounds, i.e. using 1, 3 and 5 as the maximum values of the sensitivity parameters k_D and k_Y . For both 1 and 3, the probability of the reported result being overturned is zero. For a value of 5, the corresponding probability is 0 for total R^2 -based covariate benchmarking, but is about 0.34 for both variants of partial R^2 -based covariate benchmarking. These results can be interpreted as follows: if the residualized omitted variable is up to 3 times as strongly associated with the outcome and treatment variables as the benchmark covariate, the probability that omitted variable bias will overturn the reported result is zero.¹¹ If the relative strength of association of the residualized omitted variable increases to 5, then there is about 0.3 chance (using partial R^2 -based covariate benchmarking) of the reported result being overturned.

In panel C, I report results from using relative bounds, i.e. I use 0.1 and 0.25 of k_D^{\max} and k_Y^{\max} as the upper bounds for k_D and k_Y . This would be appropriate if a researcher were able to argue that k_D and k_Y

¹¹Here we can see an advantage of using ranges instead of specific values of sensitivity parameters. Even though the 95% bias-adjusted confidence interval contains 0 for $k_D = k_Y = 3$, as can be seen from Table 1, when we use the ranges $0 \leq k_D \leq 3$ and $0 \leq k_Y \leq 3$, the contribution of that single point is negligible. Hence, we arrive at the estimate that the probability of omitted variable bias overturning the results is zero.

can only vary up to 10% or 25% of its technical maximums k_D^{\max} and k_Y^{\max} . When we use the 10% bound, the probability of the reported result being overturned is small at about 0.04 for all the three types of covariate benchmarking. For the 25% bound, the probability of the reported result being overturned increases to about 0.7 for all the three types of covariate benchmarking.

Table 2: Probability that a 95% confidence interval for the true estimate includes zero.

	Total R^2 -based covariate benchmarking	Partial R^2 -based covariate benchmarking (w/o cond on treatment)	Partial R^2 -based covariate benchmarking (cond on treatment)
Panel A: Full range:			
$0 \leq k_D \leq k_D^{\max}; 0 \leq k_Y \leq k_Y^{\max}$	0.939	0.927	0.961
Panel B: Absolute upper bounds:			
$0 \leq k_D \leq 1; 0 \leq k_Y \leq 1$	0.000	0.000	0.000
$0 \leq k_D \leq 3; 0 \leq k_Y \leq 3$	0.000	0.000	0.000
$0 \leq k_D \leq 5; 0 \leq k_Y \leq 5$	0.000	0.342	0.338
Panel C: Relative upper bounds:			
$0 \leq k_D \leq 0.1k_D^{\max}; 0 \leq k_Y \leq 0.1k_Y^{\max}$	0.037	0.033	0.100
$0 \leq k_D \leq 0.25k_D^{\max}; 0 \leq k_Y \leq 0.25k_Y^{\max}$	0.678	0.645	0.717
Memo:			
k_D^{\max}	373.13	109.12	109.12
k_Y^{\max}	3.84	7.64	8.16

Notes: For total R^2 -based covariate benchmarking, k_D^{\max} and k_Y^{\max} are the upper bounds given in (33). For partial R^2 -based covariate benchmarking (without conditioning on treatment), k_D^{\max} and k_Y^{\max} are the upper bounds given in (34). For partial R^2 -based covariate benchmarking (with conditioning on treatment), k_D^{\max} is the upper bound in (35) and k_Y^{\max} is the upper bound given in (36).

Table 2 has presented results for various limits on the ranges of k_D and k_Y . But the following question remains unanswered: Which of the bounds should a researcher use in her study? It is quite obvious that using the full range of k_D and k_Y gives very conservative estimates of the probability that omitted variable bias can overturn reported results (panel A, Table 2). Thus, it seems preferable to put some upper bound on the maximum possible values of k_D and k_Y that are lower than k_D^{\max} and k_Y^{\max} . While the absolute bounds used in panel B are low, nonetheless they are arbitrary. Because the value of k_D and k_Y vary by data sets and specifications of the empirical model, it is difficult to justify absolute bounds like 1, 3 or 5 used in panel B of Table 2. Relative bounds seem more promising because, while they can potentially be low (and therefore generate less conservative estimates), they are also benchmarked by k_D^{\max} and k_Y^{\max} —quantities that are specific to each study. In the next sub-section, I use simulation studies to generate reasonable

relative bounds that can be used across different studies.

6.5 Simulation evidence on k_D/k_D^{\max} and k_Y/k_Y^{\max}

6.5.1 Simulated data set

I draw a sample of size N from a k -dimensional multivariate Gaussian distribution with mean zero (for each of the k variables) and a covariance matrix given by $A'A$, where A is a $k \times k$ matrix formed by drawing k^2 random numbers from a univariate Gaussian distribution with mean zero and variance 1. Let us call the resultant $N \times k$ matrix as X and the ‘standardized’ version of this matrix as X^s .¹²

I generate the scalar treatment variable, D , as

$$D = a_0 + X_1^s a_1 + X_2^s a_2 + \cdots + X_{k-1}^s a_{k-1} + X_k^s a_k + u_D, \quad (38)$$

where a is a $(k+1)$ -vector formed by drawing $k+1$ random numbers from a uniform distribution over $(-1, 1)$ and u_D is an N -vector formed by drawing N random numbers from an independent Gaussian distribution with mean zero and variance $\sigma_{u_D}^2$. I generate the scalar outcome variable, Y , as

$$Y = b_0 + X_1^s b_1 + X_2^s b_2 + \cdots + X_{k-1}^s b_{k-1} + X_k^s b_k + u_Y, \quad (39)$$

where b is a $(k+1)$ -vector formed by drawing $k+1$ random numbers from a uniform distribution over $(-1, 1)$ and u_Y is an N -vector formed by drawing N random numbers from an independent Gaussian distribution with mean zero and variance $\sigma_{u_Y}^2$.

6.5.2 Estimated model and distribution of k_D/k_D^{\max} and k_Y/k_Y^{\max}

Y, D, X^s comprise the simulated data set. Using this data set, I estimate the following model using ordinary least squares (OLS),

$$Y = \beta_0 + D\tau + X_1^s \beta_1 + X_2^s \beta_2 + \cdots + X_{k-1}^s \beta_{k-1} + \varepsilon, \quad (40)$$

where I treat the k -th column of X^s as the unobserved confounder (the omitted variable) and *leave it out of the estimated model*, i.e. $Z = X_k^s$. Note that the DGP for Y in (39) does not include D , which means that the true value of τ in the estimated model (40) is zero. But since I have left out X_k^s from the estimated model

¹²Each column of X^s is the standardized version of the corresponding column of X , i.e. from each element I subtract the column mean and divide the difference by the column standard deviation. Using standardized covariates in the simulation study ensures that the results are not impacted by the scales of the variables.

Table 3: Quantiles of the empirical distribution of k_D/k_D^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.05	0.24	0.44	0.03	0.10	0.22	0.02	0.06	0.14
N=200	0.05	0.20	0.45	0.02	0.10	0.24	0.01	0.05	0.13
N=300	0.06	0.19	0.39	0.02	0.09	0.22	0.01	0.05	0.12
N=400	0.05	0.19	0.42	0.02	0.07	0.20	0.01	0.04	0.10
N=500	0.05	0.19	0.39	0.02	0.09	0.19	0.01	0.04	0.10
N=600	0.04	0.19	0.39	0.03	0.09	0.22	0.01	0.04	0.12
N=700	0.05	0.19	0.41	0.02	0.09	0.20	0.01	0.05	0.12
N=800	0.06	0.20	0.42	0.02	0.08	0.20	0.01	0.04	0.11
N=900	0.06	0.21	0.44	0.02	0.08	0.20	0.01	0.04	0.11
N=1000	0.05	0.20	0.43	0.02	0.08	0.21	0.01	0.05	0.12

Notes: For this simulation, we use $\sigma_{u_D} = \sigma_{u_Y} = 1$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_D/k_D^{\max} .

as the omitted variable, OLS will estimate τ with bias (because X_k^s is correlated with D by equation 38).

Moreover, since the true value of τ is zero, the estimated $\hat{\tau}$ will be the omitted variable bias.

To generate the empirical distribution of the ratios k_D/k_D^{\max} and k_Y/k_Y^{\max} , I do the following:

- for total R^2 -based covariate benchmarking, I compute the true magnitude of k_D and k_Y using (21), and then I compute k_D^{\max} and k_Y^{\max} using (33);
- for partial R^2 -based covariate benchmarking without conditioning on the treatment variable, I compute the true magnitude of k_D and k_Y using (24), and then I compute k_D^{\max} and k_Y^{\max} using (34);
- for partial R^2 -based covariate benchmarking with conditioning on the treatment variable, I compute the true magnitude of k_D and k_Y using (27), and then I compute k_D^{\max} and k_Y^{\max} using (35) and (35), respectively.
- Finally, in each case, I compute the ratio k_D/k_D^{\max} and k_Y/k_Y^{\max} . I do this 1000 times and then report the 50-th (p50), 75-th (p75) and 90-th (p90) quantiles of the empirical distribution of these two ratios in Table 4 through 9.

For the results reported in Table 4 and 5, the error terms in the outcome and treatment equation have equal variances; for those reported in Table 6 and 7, the variance of the error term in the treatment equation is 9 times the variance of the error term in the outcome equation; finally, for the results in Table 6 and 7, the variance of the error term in the outcome equation is 9 times the variance of the error term in the treatment equation.

Table 4: Quantiles of the empirical distribution of k_Y/k_Y^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.06	0.21	0.41	0.03	0.10	0.23	0.02	0.06	0.14
N=200	0.05	0.19	0.40	0.03	0.09	0.21	0.01	0.04	0.11
N=300	0.05	0.20	0.42	0.02	0.09	0.22	0.01	0.04	0.12
N=400	0.06	0.19	0.41	0.02	0.09	0.22	0.01	0.05	0.12
N=500	0.05	0.19	0.39	0.02	0.08	0.20	0.01	0.05	0.13
N=600	0.05	0.20	0.43	0.02	0.08	0.20	0.01	0.04	0.11
N=700	0.05	0.20	0.43	0.02	0.08	0.17	0.01	0.04	0.10
N=800	0.05	0.20	0.45	0.02	0.09	0.22	0.01	0.04	0.12
N=900	0.05	0.18	0.40	0.02	0.08	0.21	0.01	0.04	0.11
N=1000	0.05	0.19	0.39	0.02	0.07	0.18	0.01	0.04	0.11

Notes: For this simulation, we use $\sigma_{u_D} = \sigma_{u_Y} = 1$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_Y/k_Y^{\max} .

6.5.3 Results

The most striking result that emerges from Table 4 through Table 9 is that for empirical studies with relatively large sample size, i.e. $N \geq 300$, and with relatively large number of regressors, i.e. $k = 50$ or more, the 90-th percentiles of both k_D and k_Y are close to 0.1. This result holds irrespective of the relative magnitudes of the variances of the error terms in the outcome and treatment equations (in fact, some times it is lower) and can be seen by looking down the last column of Table 4 through 9. This means that in studies with relatively large sample sizes and large number of regressors, $0.1k_D^{\max}$ and $0.1k_Y^{\max}$ are reasonable upper bounds to use for k_D and k_Y , respectively.

For empirical studies with relatively few regressors, e.g. $k = 10$ or $k = 25$, we see an interesting pattern. For situations when the variance of u_D (error term in the treatment equations) is higher than the variance of u_Y (error term in the outcome equations), the 90-th percentile of the empirical distribution of k_D is low at around 0.07 but the 90-th percentile of the empirical distribution of k_Y is quite high at around 0.4. For the case when the variance of u_D (error term in the treatment equations) is lower than the variance of u_Y (error term in the outcome equations), the empirical distributions are exactly reversed.

This means that if researchers have any basis to reason about the relative magnitudes of the variances of the error terms in the treatment and outcome equations, they can use either of the following bounds: (a) when the variance of the error term in the treatment equation is higher than the variance of error term in the outcome equation, then $0.07k_D^{\max}$ and $0.4k_Y^{\max}$ are reasonable upper bounds to use for k_D and k_Y , respectively; (b) when the variance of the error term in the treatment equation is lower than the variance of

error term in the outcome equation, then $0.4k_D^{\max}$ and $0.07k_Y^{\max}$ are reasonable upper bounds to use for k_D and k_Y , respectively

Table 5: Quantiles of the empirical distribution of k_D/k_D^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.01	0.04	0.10	0.01	0.03	0.06	0.01	0.03	0.07
N=200	0.01	0.03	0.08	0.01	0.02	0.04	0.01	0.01	0.03
N=300	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=400	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=500	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=600	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=700	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=800	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=900	0.01	0.03	0.07	0.00	0.01	0.04	0.00	0.01	0.02
N=1000	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.01

Notes: For this simulation, we use $\sigma_{u_D} = 3$ and $\sigma_{u_Y} = 1$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_D/k_D^{\max} .

Table 6: Quantiles of the empirical distribution of k_Y/k_Y^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.05	0.21	0.44	0.03	0.09	0.21	0.02	0.07	0.17
N=200	0.06	0.20	0.43	0.02	0.09	0.23	0.01	0.04	0.10
N=300	0.06	0.21	0.42	0.02	0.07	0.19	0.01	0.05	0.12
N=400	0.05	0.22	0.43	0.02	0.09	0.21	0.01	0.04	0.09
N=500	0.05	0.20	0.43	0.02	0.09	0.21	0.01	0.04	0.10
N=600	0.05	0.18	0.42	0.02	0.08	0.19	0.01	0.05	0.11
N=700	0.05	0.17	0.36	0.02	0.08	0.19	0.01	0.04	0.10
N=800	0.05	0.20	0.41	0.02	0.07	0.18	0.01	0.04	0.11
N=900	0.05	0.21	0.44	0.02	0.08	0.21	0.01	0.04	0.11
N=1000	0.05	0.19	0.40	0.02	0.08	0.21	0.01	0.04	0.11

Notes: For this simulation, we use $\sigma_{u_D} = 3$ and $\sigma_{u_Y} = 1$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_Y/k_Y^{\max} .

6.6 Back to the empirical example

When we return to Table 2 with the insights gained from the simulation study, we see that the first row of panel C gives us a reliable answer. This is because in the empirical example whose results are reported in Table 2, $N = 1276$ and $k = 493$. Thus, we can safely use $0.1k_D^{\max}$ and $0.1k_Y^{\max}$ as reasonable upper bounds

Table 7: Quantiles of the distribution of k_D/k_D^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.07	0.21	0.42	0.03	0.10	0.23	0.02	0.07	0.15
N=200	0.05	0.20	0.44	0.03	0.09	0.21	0.01	0.05	0.13
N=300	0.05	0.22	0.41	0.02	0.08	0.20	0.01	0.05	0.12
N=400	0.05	0.18	0.40	0.02	0.08	0.21	0.01	0.05	0.12
N=500	0.04	0.19	0.43	0.02	0.08	0.19	0.01	0.04	0.10
N=600	0.05	0.20	0.42	0.02	0.08	0.20	0.01	0.04	0.12
N=700	0.05	0.21	0.42	0.02	0.08	0.22	0.01	0.05	0.13
N=800	0.04	0.19	0.42	0.02	0.08	0.20	0.01	0.04	0.11
N=900	0.05	0.20	0.40	0.02	0.09	0.20	0.01	0.04	0.10
N=1000	0.05	0.21	0.42	0.02	0.08	0.20	0.01	0.04	0.11

Notes: For this simulation, we use $\sigma_{u_D} = 1$ and $\sigma_{u_Y} = 3$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_D/k_D^{\max} .

Table 8: Quantiles of the empirical distribution of k_Y/k_Y^{\max} for simulated models with sample size N and number of regressors k .

	$k = 10$			$k = 25$			$k = 50$		
	p50	p75	p90	p50	p75	p90	p50	p75	p90
N=100	0.01	0.04	0.10	0.01	0.03	0.06	0.01	0.03	0.07
N=200	0.01	0.04	0.08	0.01	0.02	0.04	0.01	0.01	0.03
N=300	0.01	0.03	0.07	0.01	0.02	0.04	0.00	0.01	0.03
N=400	0.01	0.03	0.07	0.00	0.01	0.03	0.00	0.01	0.02
N=500	0.01	0.03	0.07	0.00	0.01	0.03	0.00	0.01	0.02
N=600	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=700	0.01	0.03	0.07	0.00	0.01	0.03	0.00	0.01	0.02
N=800	0.01	0.03	0.08	0.00	0.01	0.03	0.00	0.01	0.02
N=900	0.01	0.03	0.07	0.00	0.01	0.03	0.00	0.01	0.01
N=1000	0.01	0.03	0.07	0.00	0.01	0.03	0.00	0.01	0.02

Notes: For this simulation, we use $\sigma_{u_D} = 1$ and $\sigma_{u_Y} = 3$ and $T = 1000$ (number of simulations). $p50$, $p75$ and $p90$ refer to the 50-th, 75-th and 90-th percentiles of the empirical distribution of k_Y/k_Y^{\max} .

for k_D and k_Y , respectively. With these bounds, we see from panel C in Table 2 that the probability that omitted variable bias can overturn the reported result lies between 0.03 and 0.1. Thus, there is no more than a 10% chance that omitted variable bias will overturn the reported result of the study. This is strong evidence of the robustness of the reported results.

7 Conclusion

Sensitivity analysis offers a way to reason about omitted variable bias when alternative approaches, e.g. instrumental variables or natural experiments, are not available. This can be extremely valuable in economics and the social sciences more generally where observational studies are key components of empirical analysis. In this paper, I have proposed some simple statistics to quantify the likelihood of omitted variable bias overturning reported results. The methodology is intuitive and easy to implement.

In any study which is open to the problem of omitted variable bias, the researcher only needs to choose a benchmark covariate (or a set of covariates) which is (are) included in the model and is (are) deemed to be an important variable(s) in determining treatment assignment and the outcome. Then the researcher needs to choose a significance level, e.g. $\alpha = 0.05$, for constructing confidence intervals for the true estimate. The methodology proposed in this paper will compute the probability that zero is contained in the $100(1 - \alpha)\%$ bias-adjusted confidence interval, and therefore that the reported results can be nullified once omitted variable bias is taken into account.

The analysis uses two sensitivity parameters: k_D , which is the relative strength of association of the residualized omitted variable, compared to the benchmark covariate, with the treatment variable; and k_Y , which is the relative strength of association of the residualized omitted variable, compared to the benchmark covariate, with the outcome variable. These strengths of association can be conceptualized in terms of total or partial R^2 measures.

For any particular study, it is possible compute the maximum permissible values of k_D and k_Y , which I denote in this paper as k_D^{\max} and k_Y^{\max} , respectively. We can construct a bias-adjusted confidence interval for each permissible combination of k_D and k_Y , and test whether zero is contained in it. The fraction of points where zero is contained in the bias-adjusted confidence interval gives us an estimate of the probability that taking account of omitted variable bias can overturn reported results.

If k_D and k_Y are allowed to take values over the full permissible, then we get the most conservative estimate of the probability that omitted variable bias can overturn the reported results. If the researcher is able to argue for some lower values to bound k_D and k_Y , then that will generate less conservative estimates of the same probability. Evidence from simulation studies reported in this paper suggest the following rules of thumb: (a) for empirical studies with large sample sizes ($N > 300$) and relatively large number of regressors ($k > 50$), then $0.1k_D^{\max}$ and $0.1k_Y^{\max}$ are reasonable bounds to use for computing the probability that omitted variable bias can overturn the reported results; (b) for empirical studies with lower number of regressors, $0.07k_D^{\max}$ and $0.4k_Y^{\max}$ can be used as reasonable bounds if it can be argued that the variance of the error

term in the treatment equation is higher than the variance of error term in the outcome equation; when the variance of the error term in the treatment equation is suspected to be lower than the variance of error term in the outcome equation, then $0.4k_D^{\max}$ and $0.07k_Y^{\max}$ can be used as reasonable bounds.

References

- Angrist, J. D. and Lavy, V. (1999). Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement. *Quarterly Journal of Economics*, 114(4):533–575.
- Card, D. (2001). Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems. *Econometrica*, 69(5):1121–1160.
- Cinelli, C. and Hazlett, C. (2020). Making Sense of Sensitivity: Extending Omitted Variable Bias. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1):39–67.
- Cornfield, J., Haenszel, W., Hammond, E. C., Lilienfeld, A. M., Shimkin, M. B., and Wynder, E. L. (1959). Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions. *JNCI: Journal of the National Cancer Institute*, 22(1):173–203.
- De Luca, G., Magnus, J. R., and Peracchi, F. (2019a). Comments on “Unobservable Selection and Coefficient Stability” and “Poorly Measured Confounders are More Useful on the Left Than on the Right”. *Journal of Business and Economic Statistics*, 37(2):217–222.
- De Luca, G., Magnus, J. R., and Peracchi, F. (2019b). Comments on “Unobservable Selection and Coefficient Stability: Theory and Evidence” and “Poorly Measured Confounders are More Useful on the Left Than on the Right”. *Journal of Business & Economic Statistics*, 37(2):217–222.
- Diegert, P., Masten, M. A., and Poirier, A. (2023). Assessing omitted variable bias when the controls are endogenous.
- Ding, P. and VanderWeele, T. J. (2016). Sensitivity Analysis Without Assumptions. *Epidemiology*, 27(3):368–377.
- Frank, K. A. (2000). Impact of a confounding variable on a regression coefficient. *Sociological Methods & Research*, 29(2):147–194.
- Greene, W. H. (2012). *Econometric Analysis*. Prentice Hall, seventh edition.

- Imai, K., Keele, L., and Yamamoto, T. (2010). Identification, Inference and Sensitivity Analysis for Causal Mediation Effects. *Statistical Science*, 25(1):51 – 71.
- Imbens, G. W. (2003). Sensitivity to exogeneity assumptions in program evaluation. *American Economic Review*, 93(2):126–132.
- Krauth, B. (2016). Bounding a Linear Causal Effect Using Relative Correlation Restrictions. *Journal of Econometric Methods*, 5(1):117–141.
- Oster, E. (2019). Unobservable Selection and Coefficient Stability. *Journal of Business and Economic Statistics*, 37(2):187–204.
- Rao, C. R., Toutenburg, H., Shalabh, and Heumann, C. (2008). *Linear Models and Generalizations: Least Squares and Alternatives*. Springer, Berlin. Third extended edition with contributions by Michael Schomaker.
- Rosenbaum, P. R. (2002). *Observational Studies*, pages 1–17. Springer New York, New York, NY.
- Rosenbaum, P. R. and Rubin, D. B. (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society. Series B (Methodological)*, 45(2):212–218.
- VanderWeele, T. J. and Ding, P. (2017). Sensitivity analysis in observational research: introducing the e-value. *Annals of Internal Medicine*, 167(4):268–274.
- Yule, G. U. (1911). *An Introduction to the Theory of Statistics*. C. Griffin and company, limited, London.
- Yule, G. U. and Kendall, M. G. (1948). *An Introduction to the Theory of Statistics*. C. Griffin and company, limited, London.

Appendix A Proofs

A.1 Proof of Lemma 1

Proof. Using results on the inverse of partitioned matrices, it can be shown (Rao et al., 2008, page 323) that

$$P_W = P_X + \frac{(I - P_X) Z Z' (I - P_X)}{Z' (I - P_X) Z}. \quad (41)$$

Using the definition of $Z^{\perp X}$, we see that

$$\frac{(I - P_X) Z Z' (I - P_X)}{Z' (I - P_X) z} = P_{Z^{\perp X}}, \quad (42)$$

where I use the fact that $(I - P_X)$ is also a projection matrix (onto the orthogonal complement of the column space of X) and hence symmetric and idempotent. Note that $P_X P_{Z^{\perp X}} = 0$ because

$$P_X P_{Z^{\perp X}} = X (X' X)^{-1} X' Z^{\perp X} [(Z^{\perp X})' Z^{\perp X}]^{-1} (Z^{\perp X})' = 0,$$

where I have used $X' Z^{\perp X} = 0$ (i.e. residuals are orthogonal to the regressors). \square

A.2 Proof of Theorem 1

Proof. Using lemma 1, we have

$$P_W M^0 P_W = P_W M^0 M^0 P_W = (M^0 P_X + M^0 P_{Z^{\perp X}})' (M^0 P_X + M^0 P_{Z^{\perp X}}).$$

This becomes

$$P_W M^0 P_W = P_X M^0 P_X + P_{Z^{\perp X}} M^0 P_{Z^{\perp X}} \quad (43)$$

because the cross product terms are zero,

$$P_X M^0 M^0 P_{Z^{\perp X}} = P_X M^0 P_{Z^{\perp X}} = P_X P_{Z^{\perp X}} = 0,$$

by the orthogonality of P_X and $P_{Z^{\perp X}}$ and I have used

$$\begin{aligned} M^0 P_{Z^{\perp X}} &= M^0 Z^{\perp X} [(Z^{\perp X})' Z^{\perp X}]^{-1} (Z^{\perp X})' \\ &= Z^{\perp X} [(Z^{\perp X})' Z^{\perp X}]^{-1} (Z^{\perp X})' \\ &= P_{Z^{\perp X}} \end{aligned}$$

where, because $Z^{\perp X}$ is a regression residual vector, we have $M^0 Z^{\perp X} = Z^{\perp X}$ (Greene, 2012, page 40).

I pre-multiply (43) by Y' , then post-multiply the result by Y , and finally divide through by $Y' M^0 Y$ to get

$$\frac{Y' P_W M^0 P_W Y}{Y' M^0 Y} = \frac{Y' P_X M^0 P_X Y}{Y' M^0 Y} + \frac{Y' P_{Z^{\perp X}} M^0 P_{Z^{\perp X}} Y}{Y' M^0 Y}.$$

Using (12) and (13), we get

$$R_{Y \sim X+Z}^2 - R_{Y \sim X}^2 = \frac{Y' P_{Z^\perp X} M^0 P_{Z^\perp X} Y}{Y' M^0 Y}.$$

Since the right hand side is $R_{Y \sim Z^\perp X}^2$, we get (17).

We proceed by subtracting $R_{Y \sim Z}^2$ from both sides of the above equality. Using (14), we get:

$$R_{Y \sim X+Z}^2 - R_{Y \sim X}^2 - R_{Y \sim Z}^2 = \frac{Y' P_{Z^\perp X} M^0 P_{Z^\perp X} Y - Y' P_Z M^0 P_Z Y}{Y' M^0 Y} \quad (44)$$

Now, using the definition of $\eta_{X,Y,Z}$ in (19), we see that the RHS of (44) is $\eta_{X,Y,Z}$. Hence, we get (18):

$$R_{Y \sim X+Z}^2 - R_{Y \sim X}^2 - R_{Y \sim Z}^2 = \eta_{X,Y,Z}.$$

□