

The sensitivity of linear regression coefficients' confidence limits to the omission of a confounder

Carrie A. Hosman* and Ben B. Hansen†

*Statistics Department
University of Michigan
Ann Arbor, MI 48109-1107
e-mail: ben.b.hansen@umich.edu*
and

Paul W. Holland

*Paul Holland Consulting Corporation
703 Sayre Drive
Princeton, NJ 08540*

Abstract: Omitted variable bias can affect treatment effect estimates obtained from observational data due to the lack of random assignment to treatment groups. Sensitivity analyses adjust these estimates to quantify the impact of potential omitted variables. This paper presents methods of sensitivity analysis to adjust interval estimates of treatment effect — both the point estimate and standard error — obtained using multiple linear regression. Central to our approach is what we term *benchmarking*, the use of data to establish reference points for speculation about omitted confounders. The method adapts to treatment effects that may differ by subgroup, to scenarios involving omission of multiple variables, and to combinations of covariance adjustment with propensity score stratification. We illustrate it using data from an influential study of health outcomes of patients admitted to critical care.

Keywords and phrases: causal inference, hidden bias, observational study.

1. Introduction

1.1. Methodological context

In a common use of multiple linear regression, one regresses an outcome variable on an treatment variable and adjustment variables, then interprets the fitted treatment-variable coefficient as an estimate of the treatment's effect on the outcome. The interpretation relies on the assumptions of the linear model and some assumption to the effect that there either are no unmeasured confounders or at least none that demand adjustment. (Ignorability of treatment assignment (Rubin, 1978; Holland, 1988, Appendix) is one such assumption; there are many

*Supported by NSF grant # SES-0753164.

†Supported by a NSA grant (# 07Y-177) and a Ford Minority Dissertation Fellowship.

variants.) The linearity assumptions are often testable given the data, but the remaining assumption is not. When regression results are questioned, it's often the non-confounding assumption that is the focus of doubt.

Because the issue arises even with the most thorough observational studies, adjusting for any number of covariates, it fuels cynicism about observational research. If the possibility of unmeasured variable bias can't be removed, then why bother with potential confounders, particularly those that are difficult to measure, or not obvious threats? It might be clear that the damage from omitting a confounder W would be reduced by adjustment for available correlates of W yet, because introducing these correlates would draw attention to the absence of W , not at all clear that effecting the additional adjustments would enhance the credibility of the research. Plainly, the problem here is not the methodological strategy of broadly adjusting for relevant baseline characteristics but an absence of, or lack of awareness of, suitable methods with which to quantify benefits of more comprehensive confounder controls.

Sensitivity analyses, procedures quantifying the degree of omitted variable bias needed to nullify or reverse key conclusions of a study, can help. Sensitivity analysis methods for various models and data structures are proposed in [Cornfield et al. \(1959\)](#), [Rosenbaum and Rubin \(1983\)](#), [Rosenbaum \(1988\)](#), [Copas and Li \(1997\)](#), [Robins et al. \(2000\)](#), [Scharfstein and Irizarry \(2003\)](#), [Marcus \(1997\)](#), [Lin et al. \(1998\)](#), [Frank \(2000\)](#) and [Imbens \(2003\)](#), among others, the last four bearing closest resemblance to the approach to be presented here. Invariably the methods start by in some way quantifying relationships between hypothetical omitted variables and included ones, go on to give an algorithm for converting these parameters into impacts on estimates, p -values or confidence limits, and then leave to researchers themselves the task of deciding what parameter values are plausible or relevant. Here we develop a method following the first parts of this general recipe but then doing a bit more to help researchers calibrate intuitions about speculation parameters.

The resulting technique possesses a combination of advantages making it both uniquely practicable and conducive to certain insights. First, it applies to inferences made with ordinary multiple regression, as we show in [Section 2](#), as well as to inferences made with regression in combination with propensity-score stratification, a topic discussed in [Section 4.3](#). Second, it quantifies relationships between omitted and included variables in terms intrinsic to multiple regression, permitting intuitions for the relationships to be calibrated with a few additional regression fits ([§ 2](#)). ([Angrist and Imbens \(2002, §4\)](#) suggest such calibration in a related context.) Third, it represents effects of an omission with just two quantities, one tracking confoundedness with the treatment variable and the other measuring conditional association with the outcome ([§ 2](#)). Fourth, there are important practical benefits to sensitivity analysis based on both of these quantities — “dual” sensitivity analyses, in [Gastwirth et al.'s terminology \(1998\)](#) — in that analysis based only on confoundedness with the treatment variable is likely to overstate sensitivity. Our method makes this plain ([§ 3](#)), as we will demonstrate with a case study to be introduced presently, although some other methods may obscure it. Fifth, it gives closed-form characterizations of how

confidence intervals, as opposed only to estimates or hypothesis tests, could be changed by inclusion of the omitted confounder (§ 3.2). Sixth, the method readily adapts to analyses in which several omissions are suspected, or where interactions with the treatment variable are used to handle possible effect heterogeneity (§ 4). Seventh, the same application brings to light certain practical advantages of the use of propensity scores which to our knowledge have not previously been noted (§ 4.3).

1.2. A case study

Our application is to Connors et al.'s (1996) highly influential, and controversial, study of the critical-care procedure known alternately as Swan-Ganz, pulmonary artery or right heart catheterization (RHC). RHC is a procedure to perform continuous measurements of blood pressure in the heart and large blood vessels of the lungs. Introduced in 1970, it became standard procedure without first being tested in clinical studies, as might be expected today, and empirical assessments that were subsequently conducted failed to uncover evidence that it improved medical outcomes (e.g., Gore et al., 1987; Zion et al., 1990). However, these studies were criticized for insufficient confounder controls. Using a large sample, good measures and extensive adjustments for background variables, Connors et al. (1996) echoed the disappointments of the earlier assessments and went further, finding RHC to worsen, rather than improve, both mortality and the duration of treatment in critical care. Each of these studies used nonrandomized data and is in some degree vulnerable to omitted variable bias. Although the results of subsequent randomized trials have been largely consistent with this picture (Rhodes et al., 2002; Sandhan et al., 2003; Shah and Stevenson, 2004; Richard and et al, 2003; Harvey et al., 2005), the procedure remains a staple of critical care, and the surrounding debate continues. This paper examines how the omission of covariates from Connors et al.'s data might skew point and interval estimates of RHC's effect on length of stay, in the process shedding light on the degree of protection from omitted variables afforded by included ones.

1.3. The SUPPORT data

Connors et al.'s (1996) data come from the Study to Understand Prognoses and Preferences for Outcomes and Risks of Treatments (SUPPORT). The study collected data on the decision-making and outcomes of seriously ill, hospitalized adult patients. Patients included in the study had to meet certain entry criteria and a predefined level of severity of illness. All 5735 SUPPORT patients who were admitted to an ICU in the first 24 hours of study were analyzed together, for reasons detailed in Connors et al. (1996). Data include initial disease categories upon admission, physiological measurements, and demographic information. Patients were coded as having had RHC if it was performed within the first 24 hours of the study.

Length-of-stay is the difference between a patient’s recorded dates of entry and exit from the hospital. Here we study the omitted variable sensitivity of Connors et al.’s finding that RHC increases costs, by lengthening stays in the hospital. For cost analysis it is logical to compare lengths of stay irrespective of whether those stays ended in death, as Connors et al. do and as we do also. Since a medical procedure could, in principle, shorten stays only by increasing mortality, comparisons like this one speak directly to economic effects but not to health effects of the procedure. (A study focused on health outcomes, as opposed to resource utilization, would most naturally begin by analyzing survival and continue with analyses of patient experience, including duration of stay in the hospital, that address the issue of censoring by death [e.g., Rubin, 2006]. Such analyses require methods other than ordinary multiple regression, however, and sensitivity analysis for them is somewhat beyond the scope of this paper.)

1.4. First-pass regression results

Length-of-stay is right-skewed. We log-transform it before regressing it on RHC and covariates. There are approximately 50 covariates for which regression adjustments might be considered; a backward stepwise selection procedure reduces this number to 19, and estimates the RHC effect as 0.11: taking the log transformation into account, RHC seems to increase lengths of stay by about $100 * (\exp(.11) - 1) = 12\%$. To reflect variability added by variable selection (Faraway, 1992), we ran a nonparametric bootstrap to determine null quantiles for regression parameters’ t -statistics; in this case, however, bootstrap t -quantiles were similar to quantiles of the appropriate t -distribution. Either way, a 95% confidence interval for the RHC effect ranges from .06 to .16, encompassing increases of 6% up to 18%.

2. Effect estimates accounting for omitted variables

To understand how an omitted variable, W , could affect the coefficient on the treatment variable, Z , in a hypothetical regression of a given outcome on these and other variables,

$$Y = \alpha \mathbf{X}^T + \beta Z + \delta W + e, \quad (1)$$

it is well to begin by examining how included variables affect the treatment coefficient in the regression that was actually performed,

$$Y = a \mathbf{X}^T + bZ + e. \quad (2)$$

This process lends context to an accompanying sensitivity analysis.

2.1. Omitted variable bias

Perhaps the most familiar way of comparing regressors is in terms of their effects on R^2 ; we begin in this mode. Of 33 regressors (columns of the design matrix

corresponding to the set of 19 covariates, some of which are categorical with multiple levels), the one that moves R^2 the most is “DNR status,” an indicator of whether there was a do-not-resuscitate order on file for the subject on the day that he first qualified for inclusion in the study. Without it, R^2 would be reduced from .141 to .112; in contrast, removal of the next-most predictive regressor reduces R^2 only to .131. On this basis one might expect DNR status to have a relatively large effect on inferences about the RHC effect, and in a sense it does: omitting it from the regression equation increases the RHC coefficient estimate from .112 to .143, 1.2 standard errors. For comparison, consider a regressor which contributes more modestly to the outcome regression: omission of “bleeding in the upper GI tract,” for instance, reduces R^2 only by .001, the 28th smallest of 33 such reductions, and removing it moves the treatment coefficient only a few percent of an SE. Blurring the distinction between parameters and estimates just a little, we refer to this difference as bias, the bias that would have been incurred by omitting the DNR variable. It appears that R^2 does track omitted-variable bias.

There is a simple relationship between a covariate W 's contribution to explaining the response, as measured by its impact on R^2 , and the bias incurred by its omission. According to Proposition 2.1, this bias is the product of the standard error on Z , the treatment variable, as calculated without regression adjustment for W ; another term to be explained presently; and $\rho_{y \cdot w | z \mathbf{x}}$, the partial correlation of W and the response given Z and remaining covariates, X . In turn, $\rho_{y \cdot w | z \mathbf{x}}^2$ equals $[(1 - R_{\text{no } W}^2) - (1 - R_{\text{with } W}^2)] / (1 - R_{\text{no } W}^2)$, the proportionate reduction in unexplained variance when W is added as a regressor (Christensen, 1996, Ch. 6). That is to say, the square of the bias due to omitting a covariate is linear in the fraction by which that covariate would reduce unexplained variance. Returning to our example, DNR status consumes 3.3% of outcome variation that would otherwise be unexplained, which is 16 inter quartile ranges greater than the third quartile of $\rho_{y \cdot w | z \mathbf{x}}^2$'s associated with the 33 available regressors; the $\rho_{y \cdot w | z \mathbf{x}}^2$ for upper GI bleeding, in contrast, is .1%, just above the first quartile.

The degree to which an omitted variable can bias effect estimates through the value of $\rho_{y \cdot w | z \mathbf{x}}^2$ depends on the value of R^2 with which we begin. In our example, the value of R^2 is relatively small; when the R^2 value is bigger, changes in R^2 have a more pronounced effect. Imagining another regression with a higher baseline R^2 and thus less unexplained outcome variation, consider an unmeasured covariate that reduces R^2 by the same percentage as the DNR variable in our regression. As a result of having less unexplained variation to start, the proportionate reduction in unexplained variation, $\rho_{y \cdot w | z \mathbf{x}}^2$, must be larger than that of DNR in our example. Such a variable would consequently have a larger effect on the estimate of omitted variable bias.

The remaining factor in Proposition 2.1's expression for omitted variable bias expresses the degree to which the variable is confounded with treatment. This confounding turns out to figure more centrally in omitted variable bias, despite the fact that measurements of it arise less often in regression analysis

than do $\rho_{y \cdot w | z \mathbf{x}}$ or R^2 . A straightforward descriptive measurement falls out of the regression of Z on the remaining regressors from the outcome regression: their confoundedness with Z is reflected in their t -statistics, the ratios of their coefficient estimates to their conventional standard errors. Note that these are here considered as descriptive, not inferential, statistics; were this auxiliary regression instead being used for statistical inference, a nonlinear model may be required for binary Z , in which case these nominal standard errors may not be appropriate. Denoting by t_W the t -ratio attaching in this way to a regressor W , by b the RHC-coefficient in the absence of W and by β the same coefficient when W is included, the bias due to omission of a regressor decomposes as follows.

Proposition 2.1. *If $R_{y \cdot z \mathbf{x}}^2 < 1$ and t_W is finite, then*

$$\hat{b} - \hat{\beta} = \text{SE}(\hat{b})t_W\rho_{y \cdot w | z \mathbf{x}}. \quad (3)$$

A proof is given in Section 3.2.

Besides being a strong predictor of the response, DNR status ranks highly in terms of its confounding with the treatment: if RHC is regressed on the covariates then its t -statistic is -6.5 , the fifth-largest in magnitude. (Table 1 displays t_W values, along with a few square roots of F statistics, to be explained in Section 4.1, for multicategory variables. The magnitudes of the F statistics are not directly comparable to those of the t -statistics.) Consistent with there being little bias associated with removing it, upper GI bleeding is as weakly related to the treatment variable as it was to the response; its t is only -0.7 , the fifth-smallest in magnitude. The strongest confounder among the adjustment variables is the PaO₂/FIO₂ ratio, a reflection of oxygenation in the lungs, with $t = -15.3$. On the other hand, its relatively small contribution toward explaining variation in the outcome, $\rho_{y \cdot w | z \mathbf{x}}^2 = .0007$, limits the effects of its removal on the treatment coefficient; here $\hat{b} - \hat{\beta}$ is .011, or 40% of an SE. Although the effect on the treatment coefficient of excluding the PaO₂/FIO₂ variable is tempered by its small corresponding $\rho_{y \cdot w | z \mathbf{x}}^2$, (3) says that an actual omitted variable this confounded with treatment could contribute up to 15.3 standard errors' worth of bias (depending on how much more outcome variation it explains). If this possibility is not merely notional, it follows that uncertainty assessments based in the usual way on $\text{SE}(\hat{b})$ are quite incomplete.

Proposition 2.1 suggests that the degree of an omitted variable's confounding with the treatment plays a potentially much larger role in determining the bias from its omission than does its conditional association with the response. As it can explain no more than 100% of otherwise unexplained variation, in general $\rho_{y \cdot w | z \mathbf{x}}^2 \leq 1$, ensuring in turn that $|\hat{b} - \hat{\beta}|/\text{SE}(\hat{b}) \leq |t_W|$. In contrast, nothing bounds the t -ratios t_W — covariates highly collinear with the treatment can introduce biases that are arbitrarily large, at least in principle. We refer to an unsigned statistic $|t_W|$ as the *treatment confounding* of W a covariate or potential covariate W .

Table 1 reports t_W and $\rho_{y \cdot w | z \mathbf{x}}$ values that actually occur in the SUPPORT data, placing observed covariates in the role of W one at a time. With these

data, t_W is often of a magnitude to make large biases possible (although this is somewhat tempered by relatively small $\rho_{y \cdot w | z \mathbf{x}}$'s). The calculations permit the statistician and his audience to calibrate intuitions for t_W against actual variables and their relations to the treatment. We refer to this as *benchmarking* of treatment-confounding.

TABLE 1
Selected included covariates' relation to treatment and response variables, given remaining included variables

	confounding with RHC ($ t_W $) (rounded)	% decrease in unexplained variation by adding W ($100\rho_{y \cdot w z \mathbf{x}}^2$)
income*	6.8	0.3
primary initial disease cat.*	48.1	3.4
secondary initial disease cat.*	20.2	0.8
<i>Comorbidities illness:</i>		
renal disease	2.1	0.2
Upper GI bleed	0.7	0.1
<i>Day 1 Measurements:</i>		
APACHE score	5.1	0.1
white blood cell ct	0.5	0.0
heart rate	2.5	0.0
temperature	2.3	0.1
PaO2/FIO2	15.4	0.1
albumin	2.3	0.7
hematocrit	3.3	0.9
bilirubin	2.2	0.1
sodium	3.1	0.1
PaCo2	6.8	0.2
DNR	6.5	3.3
PH	3.7	0.3
<i>Admit Diagnosis Categories:</i>		
neurology	5.9	0.2
hematology	3.6	0.1

* Categorical variables with multiple levels. (Strictly speaking, these rows of the table give $F_W^{1/2}$ in the middle column, not t_W ; See Section 4.1.)

2.2. Omitted variables and the treatment's standard error

The representation (3) of omitted-variable bias, as the product of $SE(\hat{b})$ with other factors, lends hope that a meaningful combinations of errors due to variable omissions and sampling variability might be bounded by a multiple of $SE(\hat{b})$. A tempting solution is simply to add to the normal 97.5 percentile 1.96, or to whatever other multiplier of the SE has been found appropriate to describe likely sampling error, a quantity large enough to bound plausible values of $|t_W \rho_{y \cdot w | z \mathbf{x}}|$. This solution will often miss the mark, however, since omitting a covariate affects standard errors as well as coefficient estimates. This is illustrated in Table 2, which shows effects on estimation of the RHC-coefficient of omitting one-by-one the regressors discussed in § 2.1.

Excluded Variable	Estimate	Standard Error	
		df=5700	df = 50
no exclusion	0.112	0.0260	0.278
DNR order	0.143	0.0264	0.206
GI bleed	0.112	0.0260	0.274
PaO2/FIO2	0.122	0.0255	0.115

TABLE 2

RHC-coefficient and its standard error after removing specified variables.

According to Table 2, adjusting the RHC coefficient for covariates including DNR order gives a smaller standard error, .0260, than does excluding DNR order and adjusting for remaining covariates only (SE=.0264). In contrast, inclusion of PaO2/FIO2 among the adjustment variables has the effect of increasing the standard error, from .0255 to .0260. Including or excluding upper GI tract bleeding leaves the standard error unchanged, suggesting that omitted variables associated with little bias should have little effect on the standard error. This turns out to be true, in a limited sense, as can be read from Proposition 2.2: the same statistics governing the bias due to omission of a covariate W also govern its effects on the standard error, in such a way that when both t_W and $\rho_{y \cdot w|z\mathbf{x}}$ are small then including or excluding W has little effect on the Z -coefficient's standard error.

Proposition 2.2. *If $R_{y \cdot z\mathbf{x}}^2 < 1$ and t_W is finite, then*

$$\text{SE}(\hat{\beta}) = \text{SE}(\hat{b}) \sqrt{\frac{df + t_W^2}{df - 1}} \sqrt{1 - \rho_{y \cdot w|z\mathbf{x}}^2}. \quad (4)$$

Here $df = n - \text{rank}(X) - 1$, the residual degrees of freedom after Y is regressed on X and Z .

Proposition 2.2 will be strengthened in Proposition 4.3, which is proved in the Appendix.

Whereas $|t_W|$ and $|\rho_{y \cdot w|z\mathbf{x}}|$ both associate directly with the size of the bias from omitting W , they act in opposite directions on the standard error — reflecting the fact that increasing R^2 tends to reduce coefficients' standard errors, except when it's increased at the expense of introducing colinearity among regressors. The difference explains why omitting some variables increases standard errors whereas omitting others decreases them.

It also reaffirms that omitted variables' effects on statistical inferences are only incompletely reflected in omitted variable bias. In the context of the present example, variable omissions have only modest effects on the standard error; but this is a consequence of the sample being quite large ($n = 5735$) relative to the rank of the covariate matrix (33). With $df=5700$, even a t_W of 16 inflates the standard error by at most 2%, according to (4). In moderate or small samples, standard errors will be more sensitive. The third column of Table 2 presents standard errors as they would be if the sample had been much smaller, such that $df = 50$, but with the same sample means and covariances. While the

omission of a variable like GI bleed, which is only weakly related with already-included variables, still leaves the treatment’s standard error much the same, variables that are either moderately or strongly confounded with the treatment now cause wide shifts in the magnitude of the standard error. Adjustment for the PaO₂/FIO₂ measurement, for example, now increases the treatment’s standard error by a whopping 59%.

2.3. One sensitivity parameter or two?

Should we be inclined to consider worst-case scenarios, Proposition 2.2 reaffirms Proposition 2.1’s message that the omitted variable’s treatment-confounding, not its potential to increase R^2 , most demands our attention. The $\rho_{y.w|z\mathbf{x}}$ -contribution to $\text{SE}(\hat{\beta})$ is a factor bounded above by 1, whereas the t_W -factor can again be arbitrarily large — as could the t_W -contribution to omitted-variable bias. The greater potential disturbance from confounding with the treatment than from large new contributions to the outcome regression seems to be a feature multiple regression analysis shares with other statistical techniques; indeed, sensitivity analysis methods which set out to limit inference errors in terms of a single sensitivity parameter parametrize confounding with treatment, not predictivity of the response, countenancing arbitrarily strong response-predictors in the role of the hypothetical omitted variable (*e.g.*, Rosenbaum, 1988).

What would such an analysis suggest about the sensitivity of our RHC results to the omission of a confounder not unlike included variables? PaO₂/FIO₂ was sufficiently confounded with treatment to indicate as many as 16 standard errors’ worth of bias. Yet its inclusion in the regression adjustment moves the treatment effect by less than half a standard error. To have moved $\hat{\beta}$ so much it would have had to consume all of the variation in Y that was not explained without it, something no one would have expected it to do. To the contrary, inspection of other variables’ contributions to R^2 , as enabled by Table 1, suggests that .1 (for instance) would be a quite generous limit on $|\rho_{y.w|z\mathbf{x}}|$. This in turn would restrict omitted variable bias due to a treatment-confounder as strong as PaO₂/FIO₂ to 1.6 SEs — still a meaningful addition to the uncertainty estimate, if a less alarmist one. Rather than reducing the number of sensitivity parameters by permitting $\rho_{y.w|z\mathbf{x}}$ to fall anywhere within its *a priori* limits, it is more appealing to retain $\rho_{y.w|z\mathbf{x}}$, restricting it within generous bounds of plausibility.

3. Sensitivity Intervals

Taken together, Propositions 2.1 and 2.2 enable a precise, closed-form description of the union of interval estimates $\{\hat{\beta} \pm q\text{SE}(\hat{\beta}) : |t_W| \leq T, \rho_{y.w|z\mathbf{x}}^2 \leq D\}$, for any nonnegative limits T, D on omitted variables’ treatment-confounding and contributions to reducing unexplained variation — the collection of Z -slopes falling within the confidence interval after addition of a covariate W such that $-T \leq t_W \leq T$ and $\rho_{y.w|z\mathbf{x}}^2 \leq D$. Such a union of intervals is itself an interval: following Rosenbaum (2002), we call it a *sensitivity interval*; and following Small

(2007), we refer to the determining set of permissible values for $(t_W, \rho_{y \cdot w | z \mathbf{X}}^2)$ as a *sensitivity zone*. The mapping of sensitivity zones to sensitivity intervals is given in Proposition 3.1.

For $\rho^2 \in [0, 1]$ and $df_w > 0$ define f by

$$f(t, \rho^2, q, df_w) = \begin{cases} t|\rho| + q\sqrt{(1 + \frac{1}{df_w}(1 + t^2))(1 - \rho^2)}, & \rho^2 \leq g(t, q, df_w) \\ \sqrt{t^2 + q^2 \cdot (1 + \frac{1}{df_w}(1 + t^2))}, & \rho^2 > g(t, q, df_w); \end{cases} \quad (5)$$

where $g(t, q, df_w) = \sin^2\{\arctan[tq^{-1}(1 + df_w^{-1}(1 + t^2))^{-1/2}]\}$.

Proposition 3.1. *Let Y, \mathbf{X}, Z and W be as in (1) and (2), with both regressions fit either by ordinary least squares or by weighted least squares with common weights. Let $\rho_{y \cdot w | z \mathbf{X}}$ and t_W be as defined in § 2.1; suppose $R_{y \cdot z \mathbf{X}}^2 < 1, |t_W| < \infty$. Fix $q > 0$. Then*

$$\hat{\beta} \pm qSE(\hat{\beta}) \subset \hat{b} \pm f(|t_W|, \rho_{y \cdot w | z \mathbf{X}}^2, df_w, q)SE(\hat{b}) \quad (6)$$

where f is given by (5) and $df_w = df - 1$ is the residual degrees of freedom in the regression of Y on \mathbf{X}, Z , and W .

Proposition 3.1 expresses the solution to a constrained optimization problem combining the bias and variance expressions in Propositions 2.1 and 2.2. The function g defines the two cases of the optimization problem: in the case where $\rho^2 \leq g(t, q, df_w)$, the solution is on a boundary of the constraint set defined by known quantities and the speculated value of treatment confounding. The explicit solution in each case, given by the function f , allows the calculation of sensitivity limits around the given effect estimate. The proof of Proposition 3.1 appears in the Appendix.

3.1. Pegging the boundaries of the sensitivity zone

Because our analysis began with covariate selection, there are a number of deliberately omitted variables that can be used to peg at least one boundary of the sensitivity zone. As one might expect, the covariates eliminated by the stepwise procedure add little to those variables that were included in terms of prediction of the response, and it is too optimistic to suppose of a genuinely unmeasured confounder that its contribution to the outcome regression would be no greater than that of measured covariates that variable selection would put aside. On the other hand, confounding with Z plays little role in common stepwise procedures like the one we used, and the deliberately omitted variables can be used to guide intuitions about plausible values of t_W . We selected six of these whose partial associations with RHC spanned the full range of such associations among stepwise-eliminated covariates and used their t_W -values to delimit the first dimension of several sensitivity zones. Table 3 delineates the t_W -part of the

sensitivity zone accordingly, choosing the bound T on treatment-confounding to coincide with the magnitude of confounding, conditional on stepwise-selected covariates, between the treatment and each of our six covariates.

The benchmarking method, using known variables to determine plausible values of t_W , informs targeted speculation about the potential effects of omitted variable bias. Using existing information in this way, we can speculate about the effects of omitted covariates that are of a similar nature to measured covariates — t_W benchmarks extracted from partial demographic information might reasonably predict the t_W values that would attach to additional demographic variables, were they available. To calibrate intuitions about omitted variables that are different in kind from included ones, reference values for t_W might also be obtained from external data sets.

TABLE 3

95% sensitivity intervals for the treatment coefficient, with the putative unobserved variable's treatment-confounding ($|t_W|$) hypothesized to be no greater than the treatment-confounding of 6 deliberately omitted variables. The decrease it would bring to the variance of response residuals is hypothesized to be no greater than either of 2 index values, 1% and 10%, or is not restricted.

Variable	Treatment confounding benchmark		% decrease in unexplained variation ($100\rho_{y \cdot w z\mathbf{x}}^2$)		
			1%	10%	Unrestricted
Insurance class	12.2	most	(0.03, 0.20)	(-0.04, 0.26)	(-0.21, 0.43)
Respiratory eval.	8.9	some	(0.04, 0.19)	(-0.01, 0.23)	(-0.12, 0.35)
Mean blood press.	8.6	some	(0.04, 0.19)	(-0.01, 0.23)	(-0.12, 0.34)
Cardiovascular eval.	8.5	some	(0.04, 0.19)	(-0.01, 0.23)	(-0.11, 0.34)
Weight (kg)	6.1	some	(0.04, 0.18)	(0.01, 0.21)	(-0.05, 0.28)
Immunosuppression	0.4	least	(0.06, 0.16)	(0.06, 0.16)	(0.06, 0.16)

Many analysts will have sharper intuition for potential covariates' effect on R^2 , making it relatively easy to set plausible limits on $\rho_{y \cdot w|z\mathbf{x}}^2$. Table 3 considers $\rho_{y \cdot w|z\mathbf{x}}^2 \leq .01$ or $.10$, which may be useful as general starting points. In the present case study, for instance, when included covariates are removed and put in the role of W , $\rho_{y \cdot w|z\mathbf{x}}^2 = .01$ corresponds approximately to the second most predictive of them, and the strongest single predictor (DNR order) gives $\rho_{y \cdot w|z\mathbf{x}}^2 = .03$. It appears that $\rho_{y \cdot w|z\mathbf{x}}^2 \leq .1$ is a rather conservative bound: it is difficult even to find sets of included covariates that jointly contribute so much to the outcome regression as this. Only by simultaneously placing all of the covariates into the role of W , leaving the intercept alone in the role of X , does $\rho_{y \cdot w|z\mathbf{x}}^2$ reach $.10$. Restoring these variables into the regression, and barring scenarios of still-omitted variables for which $\rho_{y \cdot w|z\mathbf{x}}^2$ exceeds $.1$, Table 3 shows the result of a positive treatment effect to be sensitive to omitted confounding on par with some of the strongest of the included confounders, but insensitive to confounding weaker than that. If benchmarking leads to accurate guesses about the values of treatment confounding and reduction in unexplained variance, and if the linear model would hold were the omitted confounder added to the regressors, then 95% sensitivity intervals will have 95% coverage, despite the variable omission.

3.2. Basis for sensitivity formulas

Propositions 2.1, 2.2, and 3.1 extend better-known descriptions of bias in regression coefficients' point estimates due to variable omission (e.g., Seber, 1977, p.66) to interval estimates. They also have antecedents in earlier literature on numerical adjustment of multiple regression results for the addition or removal of a covariate (Cochran, 1938). Of the three, Proposition 2.1's proof is the most illuminating. It also conveys the flavor of the others, which appear in the appendix.

Consider \mathbf{X} to be a matrix containing a column of 1's (or columns from which a column of 1's can be recovered as a linear combination) and let Y , Z , and W be column vectors of common extent, equal to the number of rows of \mathbf{X} . An inner product is defined as $(A, B) := \sum w_i a_i b_i / \sum w_i$, where w_i is a quadratic weight for the i th observation (in the case of unweighted least squares regression, $w_i \equiv 1$). Write $\mathbf{1}$ for the column vector of 1s. For vectors A , B , and C , let $\text{Pj}(A|B, C)$ represent the projection of A into the subspace spanned by B and C . Variances and covariances are defined as follows: $\sigma_{ab \cdot c} := (A - \text{Pj}(A|C), B - \text{Pj}(B|C))$, $\sigma_{a \cdot c}^2 = \sigma_{aa \cdot c}$; $\sigma_{ab} = \sigma_{ab \cdot \mathbf{1}}$, $\sigma_a^2 = \sigma_{a \cdot \mathbf{1}}^2$. Partial correlations are then given as: $\rho_{ab} := \sigma_{ab} / (\sigma_a \sigma_b)$; $\rho_{ab \cdot c} := \sigma_{ab \cdot c} / (\sigma_{a \cdot c} \sigma_{b \cdot c})$. Denote the degrees of freedom available for estimating b as $df = n - m - 1$, where $m = \text{column.rank}(\mathbf{X})$. The nominal standard error estimates for \hat{b} and $\hat{\beta}$ (cf. (2) and (1)) are then

$$\text{SE}(\hat{b}) = df^{-1/2} \frac{\sigma_{y \cdot z\mathbf{x}}}{\sigma_{z \cdot \mathbf{x}}} \quad \text{and} \quad \text{SE}(\hat{\beta}) = (df - 1)^{-1/2} \frac{\sigma_{y \cdot z\mathbf{x}\mathbf{w}}}{\sigma_{z \cdot \mathbf{x}\mathbf{w}}}. \quad (7)$$

Proof of Proposition 2.1. To show $\hat{b} - \hat{\beta} = \text{SE}(\hat{b}) t_W \rho_{y \cdot w|z\mathbf{x}}$, write

$$\text{Pj}(W|Z, \mathbf{X}) =: B^* Z + C^* \mathbf{X}^t. \quad (8)$$

Using (2), (1), and (8), the bias in the treatment coefficient can be written as $\hat{b} - \hat{\beta} = B^* \hat{\delta}$, a well-known result (Seber, 1977, p. 66).

Write $W^{\perp\mathbf{x}}$ for $W - \text{Pj}(W|\mathbf{1}, \mathbf{X})$, $Z^{\perp\mathbf{x}}$ for $Z - \text{Pj}(Z|\mathbf{1}, \mathbf{X})$, $Y^{\perp z\mathbf{x}}$ for $Y - \text{Pj}(Y|\mathbf{1}, Z, \mathbf{X})$, and $W^{\perp z\mathbf{x}}$ for $W - \text{Pj}(W|\mathbf{1}, Z, \mathbf{X})$. Then $\text{Pj}(W^{\perp\mathbf{x}}|Z^{\perp\mathbf{x}}) = B^* Z^{\perp\mathbf{x}}$, and $\text{Pj}(Y^{\perp z\mathbf{x}}|W^{\perp z\mathbf{x}}) = \hat{\delta} W^{\perp z\mathbf{x}}$. These formulas imply $B^* = \sigma_{wz \cdot \mathbf{x}} / \sigma_{z \cdot \mathbf{x}}^2$ and $\hat{\delta} = \sigma_{yw \cdot z\mathbf{x}} / \sigma_{w \cdot z\mathbf{x}}^2 = \rho_{yw \cdot z\mathbf{x}} \sigma_{y \cdot z\mathbf{x}} / \sigma_{w \cdot z\mathbf{x}}$, so that $\hat{b} - \hat{\beta} = B^* \hat{\delta}$ can be written as the product of $\sigma_{y \cdot z\mathbf{x}} / \sigma_{z \cdot \mathbf{x}}$, $\sigma_{wz \cdot \mathbf{x}} / (\sigma_{z \cdot \mathbf{x}} \sigma_{w \cdot z\mathbf{x}})$ and $\rho_{yw \cdot z\mathbf{x}}$. Introducing mutually cancelling factors of $(df)^{\pm 1/2}$ to the first and second of these and applying (7) turns this into the product of $\text{SE}(\hat{b})$, $(df)^{1/2} \sigma_{wz \cdot \mathbf{x}} / (\sigma_{z \cdot \mathbf{x}} \sigma_{w \cdot z\mathbf{x}})$ and $\rho_{yw \cdot z\mathbf{x}}$. But t_W is just the ratio of $\sigma_{wz \cdot \mathbf{x}} / \sigma_{w \cdot \mathbf{x}}^2$ to $\sigma_{z \cdot \mathbf{x}\mathbf{w}} / [(df)^{1/2} \sigma_{w \cdot \mathbf{x}}]$, which simplifies to the second of these terms. The result follows. \square

4. Extensions

As it is presented in Section 2, our method explores sensitivities of covariance-adjusted estimates of a main effect to the omission of a single covariate. It may appear to be limited, then, to effect estimates made by linear covariate

adjustment, without interaction terms or other allowances for heterogeneity of treatment effects, and to hidden bias scenarios involving omission of a single variable, rather than several. Such an appearance would be misleading.

4.1. Several variables omitted at once

Suppose now that W denotes not one but several omitted variables, or that it represents a single nominal variable with 3 or more levels, so that its encoding in terms of a design matrix would require 2 or more columns, and 2 or more degrees of freedom. Results previously presented still describe potential effects of W 's omission, if t_W is reinterpreted in a natural way. (The sensitivity parameter $\rho_{y \cdot w|z\mathbf{x}}^2$ retains its original interpretation, as the proportionate decline in unexplained variance from including W as a regressor.)

When Z is regressed on X and a multivariate W , there is no one W -coefficient and corresponding t -statistic. The natural analogue of such a statistic is the ANOVA F -statistic comparing regression fits with and without W , F_W ; for univariate W , $F_W = t_W^2$, as is well known. When $\text{rank}(W) > 1$, define the omitted variables' treatment confounding, again denoted t_W , as the positive square root of $[k(df)/(df + 1 - k)]F_W$. Proposition 2.1 then gives the following.

Corollary 4.1. *Suppose $R_{y \cdot z\mathbf{x}}^2 < 1$, t_w^2 is finite, and $\text{rank}(W) = k > 1$. Then $(\hat{b} - \hat{\beta})^2 \leq \hat{V}(\hat{b})[k(df)(df + 1 - k)^{-1}]F_W\rho_{y \cdot w|z\mathbf{x}}^2$, or equivalently*

$$|\hat{b} - \hat{\beta}| \leq \text{SE}(\hat{b})t_W|\rho_{y \cdot w|z\mathbf{x}}|.$$

Proof of Corollary 4.1. Without loss of generality, W is uncorrelated with Z and X : if not, replacing W with $W - \text{Pj}(W|X, Z)$ leaves Z -coefficients and their standard errors unchanged. Define $\tilde{W} = \text{Pj}(Y^{\perp\mathbf{x},z}|W)$, where $Y^{\perp\mathbf{x},z} = Y - \text{Pj}(Y|X, Z)$. Again without loss of generality, $W = (W, W_2, \dots, W_k)$, where $\tilde{W} \perp (W_2, \dots, W_k)$. Writing

$$\text{Pj}(Y|Z, \mathbf{X}, W) =: \hat{\alpha} + \hat{\beta}Z + \hat{\gamma}\mathbf{X}^T + \hat{\delta}_1\tilde{W} + \hat{\delta}_2W_2 + \dots + \hat{\delta}_kW_k, \quad (9)$$

it is immediate that $\hat{\delta}_2, \dots, \hat{\delta}_k = 0$, since W_2, \dots, W_k are orthogonal to $\text{Pj}(Y^{\perp\mathbf{x},z}|W)$, and hence orthogonal to $Y^{\perp\mathbf{x},z}$. Projecting (9) onto the span of Z, \mathbf{X} , and then equating the Z -coefficient in what results with the Z -coefficient in (2) yields

$$\hat{\beta} + \hat{\delta}_1B_1^* = \hat{b}, \quad (10)$$

where B_1^* is defined by $\text{Pj}(\tilde{W}|Z, X) = B_1^*Z + C^*X$. In other words, \hat{b} and $\hat{\beta}$ are related just as they would have been had W been of rank 1, rather than k , consisting only of \tilde{W} .

We record some entailments of the definitions of $\rho_{y \cdot w|z\mathbf{x}}$, t_W and F_W in a lemma, proved in the appendix:

Lemma 4.2. *Suppose $R_{y.zx}^2 < 1$, t_W^2 is finite, and $\text{rank}(W) = k$. Then*

$$1. \rho_{y.w|zx}^2 = \rho_{y.\bar{w}|zx}^2, \quad 2. t_W^2 \leq k \frac{df}{df+1-k} F_W$$

The desired result now follows from (10), Proposition 2.1, and Lemma 4.2. \square

When $\text{rank}(W) > 1$ we have the following variant of Proposition 2.2, proved in the appendix.

Proposition 4.3. *Suppose $R_{y.zx}^2 < 1$, t_W^2 is finite, and $\text{rank}(W) = k$, $k > 1$. Then*

$$\hat{V}(\hat{\beta}) = \hat{V}(\hat{b}) \left[\frac{df + t_W^2}{df - k} \right] (1 - \rho_{y.w|zx}^2). \quad (11)$$

Because the sensitivity intervals in Proposition 3.1 follow algebraically from the bias and standard error representations (3) and (4), they work for W of arbitrary rank. The proofs of Proposition 3.1 and the following are essentially the same.

Proposition 4.4. *In the setting of Proposition 3.1 except with $\text{rank}(W) > 1$, (6) holds with $df_w = df - \text{rank}(W)$ and $t_W = \{[k(df)/(df + 1 - k)]F_W\}^{1/2}$.*

Some of the hypothetical omissions discussed in Section 2 are of the type for which Proposition 4.4 is needed. The variable ‘‘Insurance class’’ appearing in Table 3, for example, is a nominal variable with 6 categories, consuming 5 degrees of freedom when added as a regressor. Its treatment-confounding, t_W , was calculated as the appropriately rescaled square-root of the F -statistic comparing the linear regression of Z on X and it to the regression of Z on X alone, about 2.24.

4.2. Treatment effects differing by subgroup

Recall from Section 2.1 that of the 33 X -variables selected as covariates for the regression of length-of-stay on RHC, DNR status most reduced R^2 . Patients with do-not-resuscitate orders suffered 25% greater mortality during the study than other patients, probably making it inevitable that their outcomes on this variable should systematically differ from patients without such orders. It is natural to suspect that effects of RHC might differ for them as well. In this case our linear models require interactions between RHC and DNR status — and perhaps other interactions as well, for that matter, but it suffices for us to restrict attention to the treatment interaction with a single binary moderating variable, as all issues pertaining to sensitivity analysis arise in this simplest case. Marcus (1997) explores related problems.

Supplementing the regression of length-of-stay on covariates and RHC with an interaction between RHC and DNR status gave quite revealing results. The additional right-hand-side variable, an indicator of RHC and DNR simultaneously, bears a coefficient of $-.43$ and a t -statistic of -5 : it appears that the model

devoting a single term to the treatment obscured significant effect heterogeneity. Correspondingly the main RHC coefficient, interpretable as the effect for patients without DNR orders, is larger (+.15) than it was in earlier analyses without interactions (+.11); its standard error increases slightly, from .026 to .027.

To subject these results to sensitivity analysis, we again use index values to limit the impact of the hypothesized omitted covariate on R^2 : $\rho_{y \cdot w|z\mathbf{x}}^2 \leq .10$ remains a generous limit, as simultaneously adding all 33 covariates to the regression of length of stay on RHC (now interacted with DNR status) decreases unexplained variation by only slightly more, about 11%. To set suitable limits on the omitted variable’s treatment confounding, imagine for the moment that the inclusion of an interaction term had been handled somewhat differently: rather than adding ZX_1 , the product of RHC and DNR indicators, add in $\tilde{X} = ZX_1 - P_j(ZX_1|X, Z)$. The coefficient of this variable lacks any straightforward interpretation, but its addition to the right hand side of the equation has precisely the same effect on remaining coefficients as would the addition of ZX_1 itself. To benchmark treatment-confounding, we would then regress Z on X and \tilde{X} . By construction, however, \tilde{X} is orthogonal to Z and X , so that it itself earns a t -statistic of 0 in this fitting, and its inclusion has no effect other than to remove a single residual degree of freedom. In other words, the t_W -benchmarks extracted by regressing Z on X alone, used for sensitivity analysis of the RHC effect in the absence of interactions, serve just as well here after multiplication by the factor $[(df + 1)/(df)]^{1/2}$ (which in this case is effectively 1). The first 2 columns of Table 4 use benchmarks gathered in this way, finding the conclusion that RHC increases lengths of time in the hospital for patients without DNR orders to be a bit less sensitive to hidden bias than was the analogous conclusion for the analysis assuming homogeneous treatment effects, in Table 3.

Variable	Treatment confounding benchmarks found with estimated					
	effect on patients w/out DNR order		effect on patients with DNR order		ETT-weighted average effect	
	t_W	$\rho_{y \cdot w z\mathbf{x}}^2 \leq 0.1$	t_W	$\rho_{y \cdot w z\mathbf{x}}^2 \leq 0.1$	t_W	$\rho_{y \cdot w z\mathbf{x}}^2 \leq 0.1$
Insurance Class	12.2	(-0.01, 0.30)	11.9	(-0.74, 0.17)	12.2	(-0.03, 0.27)
Respiratory eval.	8.9	(0.02, 0.27)	8.1	(-0.64, 0.08)	-4.1	(0.03, 0.20)
Mean blood press.	-8.6	(0.02, 0.27)	-8.2	(-0.64, 0.08)	-5.1	(0.03, 0.21)
Cardiovascular eval.	-8.5	(0.02, 0.27)	-7.7	(-0.63, 0.07)	-7.1	(0.01, 0.23)
Weight (kg)	6.1	(0.05, 0.25)	6.4	(-0.60, 0.03)	-5.3	(0.03, 0.21)
Immunosuppression	0.4	(0.09, 0.20)	0.4	(-0.44, -0.12)	-5.9	(0.02, 0.22)

TABLE 4
Sensitivity intervals for subgroup effects and for weighted average effects

When we instructed it to include the RHC-DNR interaction among the explanatory variables, our software might equally well have added an indicator of RHC and the absence of DNR, $Z(1 - X_1)$. In this case, the main effect would be interpretable as the effect of RHC for patients *with*, rather than without, DNR orders. It follows that had that effect been the object of our interest, we could construct a sensitivity analysis for it in the same manner as just above,

by persuading our regression program to expand the interaction differently. In actuality things are still simpler than that; we really only need to take ordinary care in interpreting the regression results, and somewhat modify the benchmarking equation. The effect for patients with DNR orders is the sum of the main RHC effect and the RHC-and-DNR effects, $.148 + -.430 = -.28$, with estimated variance equal to the sum of the two estimated variances and twice their covariance, $.0065 = (.080)^2$. In parallel, to benchmark treatment confounding for this analysis, regress the sum of the main RHC indicator and the interaction term, the product of RHC and DNR indicators, on covariates. This gives somewhat different results than did the benchmarking for the RHC effect on patients without DNR orders, which omitted the interaction term from the left-hand side of its regression equation; compare the middle and left columns of Table 4.

The same approach yields a sensitivity analysis for any target parameter representable as a linear combination of main effect and interaction terms. Take the effect of treatment on the treated, or ETT, parameter, considered from within a model permitting treatment effects to vary within specified subgroups. If the groups are, for simplicity, patients with and without DNR orders, then since 7% of patients receiving RHC had DNR orders the ETT can be represented as the main effect plus .07 times the DNR-RHC interaction effect, estimated as $.148 + (.07) - .430 = .118$, with corresponding standard error .026. For benchmarking, regress on covariates the RHC indicator plus .07 times the product of RHC and DNR indicators, with results as given in the rightmost two columns of Table 4.

4.3. Propensity-adjusted estimates of treatment effects

Regression adjusts between-group comparisons by attempting to remove adjustment variables' contributions from the outcomes before comparing them. In contrast, adjustments based on propensity scores attempt to divide the sample into strata within which treatment and control subjects have similar distributions of covariates. We estimated propensity scores using all 50 of the SUPPORT data's covariates in a logistic regression (Rosenbaum and Rubin, 1984), finding six equally-sized subclasses made treatment-control differences on the covariates jointly insignificant at level $\alpha = .10$ (Hansen and Bowers, 2008). One can couple such a stratification with linear modeling to estimate treatment effects. In the simplest variant, responses are regressed on the treatment and fixed stratum effects. Fitting such a model to the SUPPORT data gives an RHC effect similar to what was estimated after ordinary covariance adjustment, as in Section 1.4, but with somewhat larger standard errors.

The main assumption for this model is that so far as the outcome, length of stay, is concerned, the only systematic difference between the RHC and non-RHC patient in a propensity stratum is RHC itself. Relax this assumption in favor of another to the effect that so far as differences between outcomes and their projections onto an omitted variable are concerned, within strata RHC and non-RHC patients do not systematically differ. Were the omitted variable

TABLE 5

Sensitivity intervals for the treatment effect after ordinary covariance and propensity-score adjustment, illustrating that propensity adjustment better limits sensitivity to the omission of adjustment variables. For covariance adjustment, t_W is limited by the confoundedness with the treatment of 6 variables that had been eliminated by a preliminary variable-selection procedure, as in Table 3; for propensity adjustment, limits on treatment confounding are set by separately removing each of these and calculating their t_W 's after propensity adjustment for remaining variables.

	OLS regression			Propensity adjusted regression		
	$ t_w $	$\rho_{y \cdot w z \mathbf{x}}^2 \leq .01$	$\rho_{y \cdot w z \mathbf{x}}^2 \leq .1$	$ t_w $	$\rho_{y \cdot w z \mathbf{x}}^2 \leq .01$	$\rho_{y \cdot w z \mathbf{x}}^2 \leq .1$
Insurance class	12.2	(0.03, 0.20)	(-0.04, 0.26)	8.6	(0.02, 0.18)	(-0.03, 0.23)
Respiratory eval	8.9	(0.04, 0.19)	(-0.01, 0.23)	3.1	(0.04, 0.17)	(0.02, 0.19)
Mean blood press.	8.6	(0.04, 0.19)	(-0.01, 0.23)	6.8	(0.03, 0.18)	(-0.01, 0.22)
Cardiovascular eval	8.5	(0.04, 0.19)	(-0.01, 0.23)	5.4	(0.03, 0.17)	(0, 0.20)
Weight (kg)	6.1	(0.04, 0.18)	(0.01, 0.21)	5.1	(0.03, 0.18)	(0, 0.21)
Immunosuppression	0.4	(0.06, 0.16)	(0.06, 0.16)	0.5	(0.04, 0.16)	(0.04, 0.16)

to become available, we could adjust by adding it to the explanatory side of the linear model. Without it, we can do sensitivity analysis.

Benchmarking treatment-confounding levels takes a bit more effort than before: rather than simply regressing Z on covariates and recording their t - or F -statistics, we have to account in some way for the propensity stratification. We do this by removing the covariates, one at a time, from the propensity model, after each removal subclassifying the sample into sextiles, as before, but now using the modified propensity score; then regressing Z on the withheld covariate and on the propensity strata in order to associate a t - or F -statistic with that covariate. Results of this process appear in Table 5.

The results exhibit a striking pattern: adjustment based on propensity scores gives causal inferences that are far less sensitive to omitted variables than does regression-based covariate adjustment. With it, one can expect less residual confounding with the treatment than with covariate adjustment, as seen in smaller t_W -values on the propensity-score side of the table. In a sense, propensity scores focus on confounding with the treatment, whereas covariate adjustment focuses on covariates and the response. Recall from Section 2 that while both matter to omitted variable bias, confounding with the treatment is both more difficult to pin down and potentially more pernicious. It stands to reason that while propensity adjustment may pay a slight penalty up front, in terms of somewhat larger standard errors than covariance adjustment, it offers a greater return downstream, in reduced sensitivity to hidden bias.

5. Summary

For effect estimates adjusted for covariates using ordinary least squares, impacts of covariate omission on point estimates and on standard errors have been represented in terms of two statistics relating the omitted variable to included ones, a measure of how adding the variable would affect R^2 and a measure of its association with the treatment variable given included variables. We refer to

the latter as the omitted variable's treatment-confounding measurement. When generous limits on how the omitted variable would affect R^2 can be defended, they yield far less pessimistic assessments of sensitivity than would be possible without such a limit. Unlike the sensitivity "parameter" pertaining to R^2 , plausible limits on the treatment confounding parameter are unlikely to emerge from intuition alone; on the other hand, it is straightforward and informative to determine study-specific benchmarks for it using available data.

The changes to the treatment coefficient's point and error estimates that the addition of an omitted covariate would cause have been represented as multiples of its standard error. So these representations yield error appraisals accounting for certain hidden biases in familiar terms, as a multiple of the SE. The method adapts readily to scenarios of multivariate omission, heterogeneous treatment effect, and combinations of regression with propensity scores.

References

- ANGRIST, J. and IMBENS, G. (2002). Comment. *Statist. Sci.*, **17** 304–07.
- CHRISTENSEN, R. (1996). *Plane Answers to Complex Questions: The Theory of Linear Models*. Springer.
- COCHRAN, W. G. (1938). The omission or addition of an independent variate in multiple linear regression. *Supplement to the Journal of the Royal Statistical Society*, **5** 171–176.
- CONNORS, A. J., SPEROFF, T., DAWSON, N., THOMAS, C., HARRELL, F. E. J., WAGNER, D., DESBIENS, N., GOLDMAN, L., WU, A., CALIFF, R., FULKERSON, W. J., VIDAILLET, H., BROSTE, S., BELLAMY, P., LYNN, J. and KNAUS, W. (1996). The effectiveness of right heart catheterization in the initial care of critically ill patients. support investigators. *J. Amer. Med. Assoc.*, **276** 889–97.
- COPAS, J. B. and LI, H. G. (1997). Inference for non-random samples. *J. Roy. Statist. Soc. Ser. B*, **59** 55–95. With discussion and a reply by the authors.
- CORNFIELD, J., HAENSZEL, W., HAMMOND, E., LILIENTHAL, A., SHIMKIN, M. and WYDNER, E. (1959). Smoking and lung cancer: Recent evidence and a discussion of some questions. *Journal of the National Cancer Institute*, **22** 173–203.
- FARAWAY, J. J. (1992). On the cost of data analysis. *Journal of Computational and Graphical Statistics*, **1** 213–229.
- FRANK, K. A. (2000). Impact of a confounding variable on a regression coefficient. *Sociological methods and research*, **29** 147–194.
- GASTWIRTH, J., KRIEGER, A. and ROSENBAUM, P. (1998). Dual and simultaneous sensitivity analysis for matched pairs. *Biometrika*, **85** 907–920.
- GORE, J., GOLDBERG, R., SPODICK, D. and ET AL (1987). A community-wide assessment of the use of pulmonary artery catheters in patients with acute myocardial infarction. *Chest*, **92** 721–727.
- HANSEN, B. B. and BOWERS, J. (2008). Covariate balance in simple, stratified and clustered comparative studies. *Statistical Science*, **23** 219–236.

- HARVEY, S., HARRISON, D. and ET AL (2005). Assessment of the clinical effectiveness of pulmonary artery catheters in management of patients in intensive care (pac-man): a randomised controlled trial. *Lancet*, **366** 472–7.
- HOLLAND, P. (1988). Causal inference, path analysis, and recursive structural equations models. *Soc. Methodol.*, **18** 449–484.
- IMBENS, G. W. (2003). Sensitivity to exogeneity assumptions in program evaluation. *American Economic Review* 126–132.
- LIN, D., PSATY, B. and KRONMAL, R. (1998). Assessing the sensitivity of regression results to unmeasured confounders in observational studies. *Biometrics*, **54** 948–963.
- MARCUS, S. M. (1997). Using omitted variable bias to assess uncertainty in the estimation of an AIDS education treatment effect. *J. Educ. and Behav. Stat.*, **22** 193–201.
- RHODES, A., CUSACK, R., NEWMAN, M. and ET AL (2002). A randomized, controlled trial of the pulmonary artery catheter in critically ill patients. *Intensive Care Medicine*, **348** 5–14.
- RICHARD, C. and ET AL (2003). Early use of the pulmonary artery catheter and outcomes in patients with shock and acute respiratory distress syndrome: a randomized controlled trial. *J. Amer. Med. Assoc.*, **290** 2732–4.
- ROBINS, J. M., SCHARFSTEIN, D. and ROTNITZKY, A. (2000). Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical Models in Epidemiology, the Environment and Clinical Trials* (M. Halloran and D. Berry, eds.). Springer, New York, 1–94.
- ROSENBAUM, P. R. (1988). Sensitivity analysis for matching with multiple controls. *Biometrika*, **75** 577–581.
- ROSENBAUM, P. R. (2002). *Observational Studies*. Springer.
- ROSENBAUM, P. R. and RUBIN, D. B. (1983). Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *J. R. Statist. Soc. B*, **45** 212–218.
- ROSENBAUM, P. R. and RUBIN, D. B. (1984). Reducing bias in observational studies using subclassification on the propensity score. *J. Am. Statist. Assoc.*, **79** 516–524.
- RUBIN, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *Ann. Statist.*, **6** 34–58.
- RUBIN, D. B. (2006). Causal Inference Through Potential Outcomes and Principal Stratification: Application to Studies with “Censoring” Due to Death. *Statist. Sci.*, **21** 299–309.
- SANDHAN, J. D., HULL, R. D., BRANT, R. F., KNOX, L., PINEO, G. F., DOIG, C. J., LAPORTA, D. P., VINER, S., PASSERINI, L., DEVITT, H., KIRBY, A., JACKA, M. and THE CANADIAN CRITICAL CARE CLINICAL TRIALS GROUP (2003). A Randomized, Controlled Trial of the Use of Pulmonary-Artery Catheters in High-Risk Surgical Patients. *New England Journal of Medicine*, **348** 5–14.
- SCHARFSTEIN, D. O. and IRIZARRY, R. A. (2003). Generalized additive selection models for the analysis of studies with potentially nonignorable missing

- outcome data. *Biometrics* 601–613.
- SEBER, G. (1977). *Linear Regression Analysis*. John R. Wiley and Sons.
- SHAH, M. and STEVENSON, L. (2004). Evaluation study of congestive heart failure and pulmonary artery catheterization effectiveness: the escape trial. *American Heart Association Scientific Sessions, New Orleans*.
- SMALL, D. S. (2007). Sensitivity Analysis for Instrumental Variables Regression With Overidentifying Restrictions. *J. Am. Statist. Assoc.*, **102** 1049.
- ZION, M., BALKIN, J., ROSENMAN, D. and ET AL. (1990). Use of pulmonary artery catheters in patients with acute myocardial infarction. *Chest*, **98** 1331–1335.

Appendix A: Proofs of theoretical results

Proof of Lemma 4.2. Under the conditions of the lemma, (1) and (2) can be established: (1) In a regression of Z and \mathbf{X} on Y , adding \tilde{W} has the same effect on the Z -coefficient and model R^2 as adding W . Thus, (1) holds.

(2) Furthermore, $\tilde{W} \in \text{span}(W)$, so \tilde{W} explains no more variation in Z than does W . Thus, $R^2_{z \cdot w|x} \geq R^2_{z \cdot \tilde{w}|x}$, which implies $\frac{\rho^2_{z \cdot w|x}}{1 - \rho^2_{z \cdot w|x}} \geq \frac{\rho^2_{z \cdot \tilde{w}|x}}{1 - \rho^2_{z \cdot \tilde{w}|x}}$. From the definition of the ANOVA F-statistic: $F_W = [(\sigma^2_{z \cdot \mathbf{x}} - \sigma^2_{z \cdot \mathbf{x}\mathbf{w}})/(k)]/[\sigma^2_{z \cdot \mathbf{x}\mathbf{w}}/(df + 1 - k)]$, which implies $F_W = [(df + 1 - k)/k][\rho^2_{z \cdot w|x}/(1 - \rho^2_{z \cdot w|x})]$

In Section 4.1, t_W is defined for multivariate W . Since $k = \text{rank}(W)$, in case of univariate W , $t_W^2 = F_W$. As $\text{rank}(\tilde{W}) = 1$, $F_{\tilde{W}} = t_{\tilde{W}}^2 = df \frac{\rho^2_{z \cdot \tilde{w}|x}}{1 - \rho^2_{z \cdot \tilde{w}|x}}$. The result follows. \square

Proof of Proposition 4.3. To relate $\text{SE}(\hat{b})$ and $\text{SE}(\hat{\beta})$, begin by rewriting some of the variance terms: $\sigma^2_{y \cdot z\mathbf{x}\mathbf{w}} = \text{Var}(Y^{\perp z\mathbf{x}\mathbf{w}}) = \text{Var}(Y^{\perp z\mathbf{x}} - \text{Pj}(Y^{\perp z\mathbf{x}}|W^{\perp z\mathbf{x}})) = \sigma^2_{y \cdot z\mathbf{x}} - \sigma^2_{y \cdot z\mathbf{x}} \rho^2_{y \cdot z\mathbf{x}}$, which implies $\sigma_{y \cdot z\mathbf{x}\mathbf{w}} = \sigma_{y \cdot z\mathbf{x}}(1 - \rho^2_{y \cdot z\mathbf{x}})^{-1/2}$.

Since $\rho^2_{z \cdot w|x}$ is the proportionate reduction in residual variance when W is added to the regression of Z on X , $\sigma^2_{z \cdot \mathbf{x}\mathbf{w}} = \sigma^2_{z \cdot \mathbf{x}}(1 - \rho^2_{z \cdot w|x})$.

Thus, $\text{SE}(\hat{\beta}) = (df - k)^{-1/2}(\sigma_{y \cdot z\mathbf{x}\mathbf{w}}/\sigma_{z \cdot \mathbf{x}\mathbf{w}})[(1 - \rho^2_{y \cdot z\mathbf{x}})/(1 - \rho^2_{z \cdot w|x})]^{1/2}$, or equivalently, $\text{SE}(\hat{b})(df/df - k)^{1/2}[(1 - \rho^2_{y \cdot z\mathbf{x}})/(1 - \rho^2_{z \cdot w|x})]^{1/2}$.

Now, $(1 - \rho^2_{z \cdot w|x})^{-1} = 1 + k(df + 1 - k)^{-1}F_W$. By the relationship of F_W to t_W^2 , it follows that $(1 - \rho^2_{z \cdot w|x})^{-1} = (df + t_W^2)(df^{-1})$, which combines with the SE relationship to give (11). \square

Proof of Proposition 3.1. There are two cases. First suppose $\rho^2_{y \cdot w|z\mathbf{x}} > g(t_w, q, c)$. For a fixed q and t_w , by Proposition 2.1 and Proposition 2.2, the desired confidence interval $\hat{\beta} \pm q\text{SE}(\hat{\beta})$ can be represented as $\hat{\beta} - q\text{SE}(\hat{\beta}) = \hat{b} + l(\arcsin \rho_{y \cdot w|z\mathbf{x}})$ and $\hat{\beta} + q\text{SE}(\hat{\beta}) = \hat{b} + u(\arcsin \rho_{y \cdot w|z\mathbf{x}})$, where $\arcsin \rho_{y \cdot w|z\mathbf{x}} \in (-\pi/2, \pi/2)$ and $l(\theta) := a \sin \theta - d \cos \theta$ and $u(\theta) := a \sin \theta + d \cos \theta$ for $a := -t_W \text{SE}(\hat{b})$, and

$d := ((t_W^2 + df)/(df - 1)SE(\hat{b}))^{1/2}$. By calculus, $u(\cdot)$ can be seen to have its maximum over $(-\pi/2, \pi/2)$ at $\arctan(a/d) \in (-\pi/2, \pi/2)$. By algebra and trigonometric identities, that maximum is $(a^2 + d^2)^{1/2}$. Similarly, over $(-\pi/2, \pi/2)$, $l(\cdot)$ takes its minimum value of $-(a^2 + d^2)^{1/2}$ at $\arctan -(a/d)$.

Now suppose $\rho_{y \cdot w | z \mathbf{x}}^2 \leq g(t_w, q, c)$. For some ρ such that

$$0 \leq \rho \leq \sin \arctan(|t_W|[(t_W^2 + df)/(df - 1)]^{-1/2}),$$

let $\rho_{y \cdot w | z \mathbf{x}} \in [-\rho, \rho]$. Then, $\arg \max_{\rho_{y \cdot w | z \mathbf{x}}} u(\arcsin \rho_{y \cdot w | z \mathbf{x}}) = \rho$ if $t_W < 0$, otherwise it is $-\rho$ if $t_W > 0$. Similarly, $\arg \min_{\rho_{y \cdot w | z \mathbf{x}}} l(\arcsin \rho_{y \cdot w | z \mathbf{x}}) = -\rho$ if $t_W < 0$, otherwise it is ρ if $t_W > 0$.

To verify the relationship for u , note $u(\arcsin(\cdot))$ takes only one maximum as its argument ranges over $[-1, 1]$ at $\sin \arctan |t_W|[(t_W^2 + df)/(df - 1)]^{-1/2} \geq \rho$. Thus, $u(\arcsin \rho_{y \cdot w | z \mathbf{x}})$ must take its maximum as $\rho_{y \cdot w | z \mathbf{x}}$ ranges over $[-\rho, \rho]$ at one of the endpoints of that interval. The analogue for l is established by similar reasoning. \square

Acknowledgements

The authors thank: Derek Briggs, Stephen Fienberg, John Gargani, Tom Love, Tanya Moore, Paul Rosenbaum, Nanny Wermuth and an anonymous reviewer, for helpful comments; and Connors et al., particularly C. Thomas and F. Harrell, for making their data publicly available.