Practical Sensitivity Analysis

Kevin A. Clarke [†]
University of Rochester

Please respect the first draft status of this paper, and do not cite without permission.

Abstract

This paper is intended to serve as a practical guide to sensitivity analysis in econometric research. Sensitivity analysis is a technique that allows a researcher to assess the impact of an unmeasured variable on a coefficient of interest. I discuss a variant of sensitivity analysis that is useful for econometrics, provide computer code and an example, and discuss the sensitivity of sensitivity analysis to the choices made by researchers.

August 2006

[†] Assistant Professor, Department of Political Science, University of Rochester, Rochester, NY 14627-0146. Tel.: (585) 275-5217; Fax: (585) 271-1616; Email: `kevin.clarke@rochester.edu`. I thank Fabiana Machado and Jeremy Kedziora for excellent research assistance.

1 Introduction

Most of the data analysis performed in political science has, often implicitly, causal statement as its goal. For example, we wish to aver, understandably, that dyadic democracy causes an absence of war. Once we have our result in hand, we turn to assessing how robust that result is to various kinds of biases. First and foremost among these is omitted variable bias. The sum total of our efforts at robustness checking, however, amounts to controlling for whatever additional variables we have at hand. If the estimated coefficient changes little in a regression with 15 additional covariates, we celebrate.

Unfortunately, omitted variable bias is a subject about which we know less than we think. Standard textbook treatments of the subject analyze a situation—choosing whether or not to include a single relevant covariate or a single set of covariates—that is unlikely to occur in practice. Under this contrived scenario, however, we understand omitted variable bias perfectly. When correlated with the other variables in a specification, including the omitted variable or variables unambiguously decreases the bias on the estimated coefficient of interest. We do not live in a contrived world, however.

In a more realistic scenario—choosing whether to include a subset of the set of omitted variables—there is precious little information to guide us. Whether or not the inclusion of this subset of omitted variables increases or decreases the bias on the estimated coefficient of interest is generally unknown and depends upon a host of factors (Clarke 2005; Clarke 2006). Simply adding the variables at hand to our specifications, therefore, provides no meaningful prophylactic, and new strategies are required. One of the promising avenues of research comes from statistics and is known as sensitivity analysis. The technique has a long history, dating back to a discussion of the evidence that smoking causes lung cancer in 1959 (Cornfield *et al.* 1959). More recently, the technique has been developed in different directions by Rosenbaum and Rubin (1983), Manski (1990), Lin, Psaty, and Kronmal (1998), and Rosenbaum (2002), among others.

The goal of this article is to introduce political scientists to one variant of sensitivity analysis that is compatible with the practices of empirical political science. This approach was developed by Imbens (2003) and is explicitly based on Rosenbaum and Rubin (1983). After describing the method, I provide an example and a discussion of the choices an applied researcher must make when using the technique. I also discuss some examples that may complicate the interpretation of the results of a sensitivity analysis. Finally, I discuss directions for future work and provide an example of the computer code used to do the analyzes included herein.

2 The Need for Sensitivity Analysis

Deciding which variables should be included in a statistical model is one of the unsolved, and probably unsolvable, problems in econometrics. It is, of course, well-known that relevant omitted variables bias regression results. The standard response to this knowledge has been to include additional control variables under the belief that the inclusion of every additional variable serves to reduce the potential threat from omitted variable bias.

Unfortunately, reality is more complicated, and the control variable strategy does not protect us from omitted variable bias. Despite the demonstrations in textbooks, we are never in the position of choosing between including the final relevant variable (or set of variables) or not. Rather, we are faced with choosing between including an additional relevant variable (or set of variables) out of a larger set of relevant omitted variables that we either do not know or cannot measure. The effect of adding to a statistical specification some, but not all, of these relevant omitted variables is generally unknown.

To put the argument in mathematical terms, consider a data generating process in scalar notation,

$$Y_i = \beta_0 + \beta_1 X_{i1} + \beta_2 X_{i2} + \beta_3 X_{i3} + \epsilon_i, \quad \epsilon_i \sim N(0, \sigma^2),$$

and two misspecified models,

$$\begin{aligned} &\text{Model 1:} \qquad Y_i = \beta_{01} + \beta_{11} X_{i1} + \epsilon_{i1}, \\ &\text{Model 2:} \qquad Y_i = \beta_{02} + \beta_{12} X_{i1} + \beta_{22} X_{i2} + \epsilon_{i2}. \end{aligned}$$

The claim, based on the logic of control variables, is that the bias on $\hat{\beta}_{11}$, the estimated coefficient on X_1 in model 1, is greater than the bias on $\hat{\beta}_{12}$, the estimated coefficient on X_1 in model 2. Letting the bias on $\hat{\beta}_{11}$, $E[\hat{\beta}_{11}] - \beta_1$, be denoted as $b(\hat{\beta}_{11}, \beta_1)$, and the bias on $\hat{\beta}_{12}$, $E[\hat{\beta}_{12}] - \beta_1$ be denoted as $b(\hat{\beta}_{12}, \beta_1)$, the mathematical argument is that

$$b(\hat{\beta}_{11}, \beta_1) \geq b(\hat{\beta}_{12}, \beta_1).$$

The mathematics of regression analysis, however, do not support the conclusion that the bias on $\hat{\beta}_{11}$ is necessarily greater than or equal to the bias on $\hat{\beta}_{12}$. The inclusion of additional relevant variables *may* reduce the bias on the X_1 coefficient, but it may also have the opposite effect and increase the bias on the X_1 coefficient. Short of knowing all omitted relevant variables, the researcher cannot know which is the case. Thus, in order to have confidence that an effect that we observe is real, and not the product of an omitted variable, we need to use a technique such as sensitivity analysis.[1] (See Clarke 2005 and Clarke 2006 for more information.)

3 Sensitivity Analysis

Statistical discussions of sensitivity analysis often come cloaked, sometimes hermetically, in the language and arcana of causal analysis, which dates back to Rubin (1974). The actual ideas, however, are quite straightforward. Observational studies attempt to determine the effect of a particular cause in the absence of the benefits of randomization. Because a researcher cannot randomly assign people to smoke, and those who smoke are likely to be different from those who do not smoke, differences in the dependent variable cannot necessarily be attributed to the putative cause—smoking. In observational studies, therefore, means other than randomization must be found for ensuring the two groups area similar enough for comparison. The common techniques are, of course, regression and variants (covariance adjustment) and matching (nonparametric covariance adjustment).

Whether or not this literature applies broadly to political science is an open question as many of our studies do not fall into the statistical definition of observational. Rosenbaum (2002, 1), for instance, argues that observational studies concern "treatments, interventions, or policies and the effects they cause..." and that a study without a treatment "is neither an experiment nor an observational study." The distinction goes back to Holland (1986)'s insistence that every unit be "potentially exposable" to the cause. Attributes of units therefore cannot be causes. The question is whether typical political science type variables such as democracy can be considered causes and whether studies that investigate the effects of democracy can be considered

[1] It should be noted that matching does not solve this particular problem, and the same demonstration can be performed using matching.

observational. Rephrasing Holland, is democracy an attribute of units or a cause that can act on units?

For the purposes of this article, I assume that variables such as democracy are causes that can act on units. States, after all, can have different regime types, and at worst, such cases fall into Holland's grey area concerning voluntary human activity (states, for instance, choose to be democracies). This assumption means that we can treat many political science studies as observational. More importantly, however, while the following discussion is sometimes couched in the language of causal analysis, the distinction between what is and is not an observational study is not particulary meaningful for the particular brand of sensitivity analysis that is the subject of this paper.

The goal of sensitivity analysis is to provide a sense of how large an effect an omitted variable or variables would have to have in order to invalidate a finding. That is, sensitivity analysis provides a quantitative statement that in order to explain away a particular association, one would need a hidden or unobserved bias of a certain size (Rosenbaum 2002). The canonical example comes from biostatistics and concerns smoking and lung cancer. Cornfield *et al.* (1959, 194) demonstrate that if cigarette smokers have 9 times the risk of nonsmokers for lung cancer but only because of some as yet unknown factor X (and thus smoking is not a causal factor), then the proportion of smokers with factor X must be at least 9 times greater than the proportion of nonsmokers with factor X. To explain away an association as strong as that between smoking and lung cancer, then, it is necessary to hypothesize a hidden bias with a very large magnitude. Given that the existence of such a bias is unlikely, we gain confidence in the reliability of our finding, and controlling for every possible omitted variable is unnecessary.

The practical approach to sensitivity analysis taken here comes from Imbens (2003) who, in turn, bases his work on a method pioneered by Rosenbaum and Rubin (1983). Much of the work done on sensitivity analysis has focused on statistical tests that are rarely used in political science such as McNemar's Test, the Sign Test, the Signed Rank Test, and the Rank Test (Rosenbaum 2002). Imbens's approach is of interest because he takes an explicitly econometric approach to sensitivity analysis and formulates the problem in terms of regression.[2] The analysis is conducted by making assumptions about the

[2] Lin, Psaty, and Kronmal (1998) also take a regression approach to sensitivity analysis, but their method is more involved.

effect of an omitted variable on the dependent variable and on an independent variable of interest. Let the possibly omitted variable be U_i, the variable of interest be W_i, and the other covariates be $\mathbf{X}_i$. Both U and W are assumed to be 0,1 for simplicity.

The distribution of the variable of interest, W_i, given the possibly omitted variable and the other covariates is assumed to be logistic,

$$\Pr(W = 1|\mathbf{X}, U) = \frac{\exp(\gamma'\mathbf{X} + \alpha U)}{1 + \exp(\gamma'\mathbf{X} + \alpha U)}. \tag{1}$$

Furthermore, we assume that the distribution of the dependent variable, Y, is normal, given U and $\mathbf{X}$,

$$Y|\mathbf{X}, U \sim N(\tau w + \boldsymbol{\beta}'\mathbf{X} + \delta U, \sigma^2). \tag{2}$$

The trick of sensitivity analysis is to choose values for α, the effect of the possibly omitted variable on the variable of interest, and δ, the effect of the possible omitted variable on the dependent variable, and calculate the maximum-likelihood estimate, $\hat{\tau}$, of the effect of the variable of interest. So, by varying α and δ, we can get a range of estimates for the effect of the variable of interest on the dependent variable.

The maximum likelihood estimator of τ comes from maximizing the following log-likelihood function,

$$\begin{aligned} L(\tau, \boldsymbol{\beta}, \sigma^2, \gamma, \alpha, \delta) &= \sum_{i=1}^{n} \ln \left[\frac{1}{2} \left(\frac{1}{\sqrt{2\pi\sigma^2}} \right) \times \exp \left(-\frac{1}{2\sigma^2} (Y_i - \tau W_i - \boldsymbol{\beta}'\mathbf{X}_i)^2 \right) \right. \\ &\times \frac{(\exp(\gamma'\mathbf{X}_i))^{W_i}}{1 + \exp(\gamma'\mathbf{X}_i)} + \frac{1}{2} \left(\frac{1}{\sqrt{2\pi\sigma^2}} \right) \\ &\times \exp \left(-\frac{1}{2\sigma^2} (Y_i - \tau W_i - \boldsymbol{\beta}'\mathbf{X}_i - \delta)^2 \right) \\ &\times \left. \frac{(\exp(\gamma'\mathbf{X}_i + \alpha))^{W_i}}{1 + \exp(\gamma'\mathbf{X}_i + \alpha)} \right]. \end{aligned}$$

Imbens argues that the sensitivity parameters α and δ are not easily interpretable. He solves this problem by translating the sensitivity parameters into partial R^2s and comparing "the amount of variation that is explained by the unobserved covariate relative to the amount not explained by the observed covariates" [128]. The proportion of the previously unexplained variation in Y that is explained by the unobserved covariate is

$$\begin{aligned} R^2_{Y,par} &= \frac{R^2_Y(\alpha,\delta) - R^2_Y(0,0)}{1 - R^2_Y(0,0)} \\ &= \frac{\hat{\sigma}^2(0,0) - \hat{\sigma}^2(\alpha,\delta)}{\hat{\sigma}^2(0,0)}, \end{aligned}$$

where $R^2_Y(\alpha,\delta) = 1 - \hat{\sigma}^2(\alpha,\delta)/\Sigma_Y$ and $\Sigma_Y = \Sigma_i(Y_i - \bar{Y})^2/N$.

For the logistic regression, there is no natural R^2 so Imbens uses the implicit R^2,

$$R^2_W(\alpha,\delta) = \frac{\hat{\gamma}(\alpha,\delta)'\mathbf{\Sigma_X}\hat{\gamma}(\alpha,\delta) + \alpha^2/4}{\hat{\gamma}(\alpha,\delta)'\mathbf{\Sigma_X}\hat{\gamma}(\alpha,\delta) + \alpha^2/4 + \pi^2/3},$$

where $\mathbf{\Sigma_X}$ is the sample covariance matrix of the observed covariates $\mathbf{X}$ with the constant term omitted. The partial R^2 is then

$$R^2_{W,par} = \frac{R^2_W(\alpha,\delta) - R^2_W(0,0)}{1 - R^2_W(0,0)}.$$

To perform the sensitivity analysis, we construct pairs of partial R^2s, $R^2_{W,par}$ and $R^2_{Y,par}$, for pairs of α and δ so that the coefficient on the variable of interest, $\hat{\tau}$, changes by a preset amount. "If the set of all such values does not include reasonable values of the partial R^2 values," then the sign of the estimated coefficient on the variable of interest is judged robust (Imbens 2003, 128). Reasonableness is judged by comparing these partial R^2 to pairs of partial R^2 values from observed covariates.

4 Performing a Sensitivity Analysis

The previous paragraph suggests the two major issues that confront anyone wishing to perform a sensitivity analysis. The first issue is by what preset amount the coefficient on the variable of interest, $\hat{\tau}$, should change. The second is how to chose values of α and δ such that they move $\hat{\tau}$ by that preset amount.

As an example of his method, Imbens uses the well-known job-training program data analyzed by Lalonde (1986). As he knows the experimental estimate of the average effect of the job-training program ($1,672), he can make the judgement that if this average effect were to change by more than $1,000, then the program would be judged to have no effect. In many political science applications, however, we are more concerned with whether or not we get the sign on the coefficient of interest correct rather than the actual value of the coefficient. Thus, for political science, it seems that zero is an appropriate choice for many applications.[3] If the estimated coefficient shows a positive effect, we should ask how large the effect of an unmeasured covariate would have to be in order to turn the coefficient negative. Similarly, if the estimated coefficient shows a negative effect, we should ask how large the effect of an unmeasured covariate would have to be in order to turn the estimated coefficient positive. This choice is important because, as I demonstrate in Section 6, it can change the results of the analysis.

The second issue is how to choose values of α and δ that move $\hat{\tau}$ by the preset amount. Although Imbens is silent on this issue, it seems that a grid search is a reasonable, if not efficient, way to proceed. We can specify a range of values for α and the same set of values for δ. We then obtain values for $\hat{\tau}$ by maximizing the log-likelihood function in the previous section for every pair of α and δ. For those pairs that change $\hat{\tau}$ to the preset amount, we calculate the partial R^2s and report them.

Running this kind of analysis is extremely computationally intensive. Say we let α range between -10 and 10 by distances of 0.2. At the same time, we let δ range between -10 and 10 by distances of 0.2. Thus, there are 10201 pairs of values to assess, meaning that we have to maximize the log-likelihood function 10201 times. If we were to let α and δ range between -40 and 40(by 0.2), we would have to maximize the log-likelihood function 160801 times. Unfortunately, we can not always go with a smaller number because, depending on the preset amount chosen for the change in $\hat{\tau}$, we

[3]Note that sensitivity analysis is not a test to check whether or not the coefficient *is* zero. Rather, sensitivity analysis measures how large the effect of an unmeasured covariate would have to be in order to produce an estimated coefficient of zero.

may need larger values of α and δ to produce the desired change. When the desired amount of change cannot be reached by any combination of α and δ, the resulting plot of these pair of partial R^2 is, of course, empty.

The Appendix includes computer code for running a sensitivity analysis in R on a 24-machine Linux cluster. For a data set with 500 observations and α and δ ranging between -10 and 10, the analysis takes 15 minutes to run. With a larger data set or particularly with a larger grid search, the analysis can take many hours to run even on a cluster. Running such an analysis on a desktop machine is likely to take days.

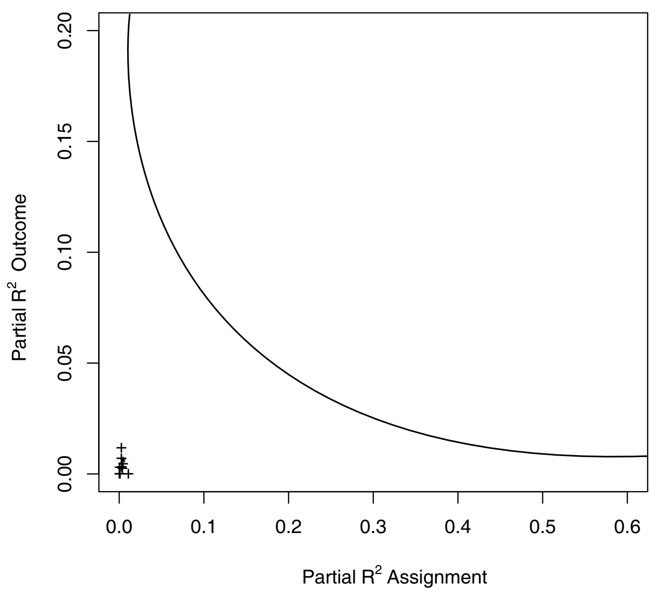

Figure 1: Replication of Imbens's Figure 1.

As a test of the code in the Appendix, I replicated Imbens's analysis of the venerable Lalonde data. Figure 1 above replicates Imbens's Figure 1 [130] almost perfectly. The code therefore appears to work. The curve is produced by choosing the lowest pairs of R^2 values that create *at least* the \$1,000 change for which Imbens is looking.

5 An Example

To make these ideas concrete and transparent, I revisit an analysis by Buhaug and Gates (2002), the objective of which is to "examine factors that determine location and scope of civil wars" [420-421]. Location is defined as the distance between the capital city and the conflict center point, and scope is defined as the geographic domain of the conflict zone, measured as the circular area centered around the conflict center and covering all significant battle zones (rounded to the nearest 50-km interval).[4] The data set includes 265 civil conflicts in the period 1946-2000.

The authors hypothesize that international borders are related to the size of a conflict zone because borders are valuable to the leaders of a rebellion. Rebels attempt to push conflicts to borders because neighboring states may provide safe refuge from government troops, and borders are natural places for rebels to gain access to the weapons and resources trade. 51% of the conflicts in the sample extend to or cross the border of the conflict-ridden state. The results from Buhaug and Gates's model 5 [427] are in Table 1.[5] Their results show that a conflict that abuts an international border is, on average, roughly 10 percentage points larger than if the conflict does not abut an international border.

Table 1: Relative Scope ($N = 246$)

Variable	Coefficient	Standard Error
Location	5.64	1.235**
Land area	-14.88	1.226**
Duration	0.77	0.302**
Border	9.49	4.545**
Resource	17.51	5.533**
Constant	91.62	6.897**
R^2	0.37	

** $p \leq 0.05$.

[4]The *relative* scope of conflict is the conflict zone as a proportion of total land area.

[5]The additional covariates include location (the distance from the conflict center to the capital center), land area (the size of the country), duration (the duration of the conflict), and resource (whether or not the conflict zone contains natural resources).

The question we want to answer is how large an unobserved covariate or covariates would have to be to destroy the substantive significance of the estimated coefficient on border. We can imagine, for instance, any number of non-geographic variables that might be correlated with scope of conflict and border. Whether or not the rebel group constitutes an ethnic faction is just one such example.[6] Such non-geographic variables might not be included in the specification for a host of reasons. The data on these variables might not be available, or perhaps they are measured with significant error, or the concepts might not be measurable, or including such variables may simply be a distraction for a researcher interested solely in the effects of geographic variables. Formal sensitivity analysis can be used in any of these instances to help assess the robustness of the coefficient on border.

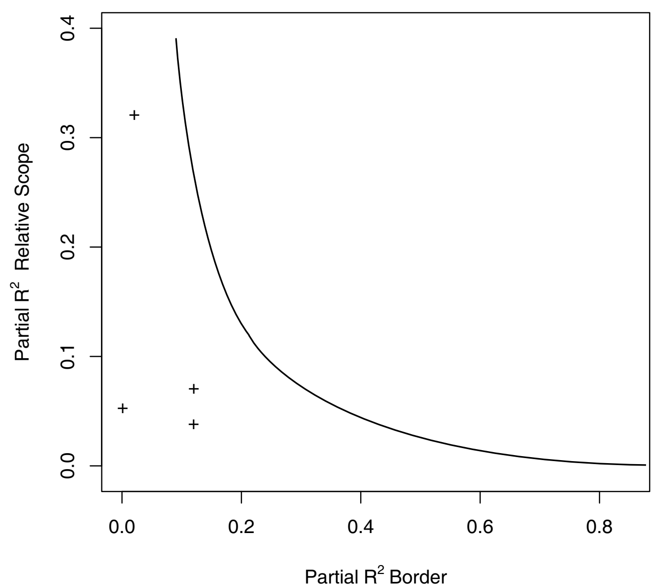

Figure 2: How large a hidden bias would have to be in order to make the coefficient on border zero.

Note that the standard error on border is roughly 4.5. If the observed coefficient were to decrease by two standard errors, it would be essentially zero, and border would not have any substantive significance. Thus, we want to

[6]Ethnic groups often cross state lines, and a rebel ethnic minority funded by an ethnic majority elsewhere would serve to widen the scope of the conflict.

know how large an effect a non-geographic variable or variables would have to have in order to render border spurious. To answer this question, we treat border as W and construct our pairs of partial R^2s for pairs of our sensitivity parameters, α and δ, that decrease the estimated coefficient on border by two standard errors. The results are in Figure 2.

The curve in Figure 2 describes how strongly the unobserved covariate would have to be correlated with the scope of conflict and whether or not the conflict abuts a border in order to make the coefficient on border be zero. An unobserved covariate would have to explain, for example, 20% of the variation in border and 15% of the variation in scope not explained by the observed covariates to make the coefficient on border zero. Now, compare the effects of the observed covariates. All lie below the curve. Thus, in order for an unobserved covariate to wipe out the effect of border on the scope of conflict, it would have to explain more of the variation in border and scope than location of the conflict, the size of the state, the duration of the conflict, and whether or not the conflict zone includes natural resources. Given that the existence of a variable of that importance is unlikely, we judge the coefficient on border to be robust. It is possible, therefore, to assess the possible impact of omitted variables without resorting to including large numbers of control variables. The Buhaug and Gates (2002) regression includes a manageable number of covariates, and we can be reasonably certain that incorporating non-geographic variables into the specification would not change the results on border.

6 The Sensitivity of Sensitivity Analysis

A perfectly accurate method of determining whether or not a variable is missing from a specification will never be found. The best we can hope for is a reasonably reliable guide that will serve to either increase or decrease our confidence in a particular result. Sensitivity analysis may be able to serve that function. Like any statistical technique, sensitivity analysis can, however, provide misleading answers under certain conditions. Fully understanding sensitivity analysis means beginning to understand the conditions under which the procedure may be misleading. Two cases are of particular interest, and these are depicted in Figure 3.

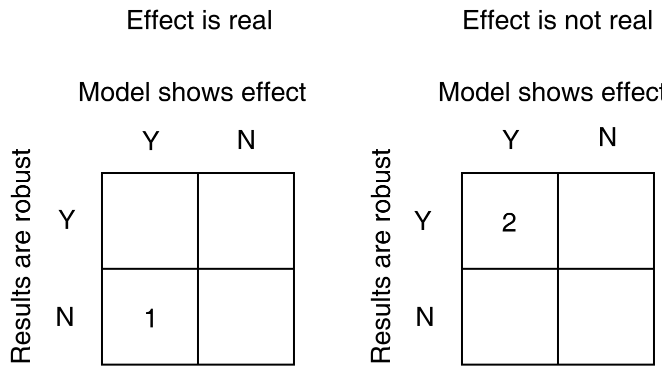

Figure 3: Two cases where the analysis might be misleading.

In the first case, the variable of interest has a true effect on the dependent variable. The results of a regression also show that the variable of interest has an effect. The problem comes when the sensitivity analysis indicates that an omitted variable might exist that could change the results. Thus, the sensitivity analysis incorrectly sends the researcher looking for nonexistent omitted variable. In the second case, the variable of interest does not have a true effect, and the estimated regression indicates that it does. The problem occurs when the sensitivity analysis indicates that the results are robust. In this case, the sensitivity analysis incorrectly lends confidence to false results. To understand the conditions under which these misleading results might occur, we have to understand how the two pieces of the graph—the curve and the "pluses"—move in relation to choices the researcher makes.

The curve comprises pairs of partial R^2s for values of α and δ such that the estimated effect of $\hat{\tau}$ changes by the preset amount. The major determinant of the placement of this curve is, of course, that preset amount. The larger the preset amount, the larger an effect an omitted variable would have to have in order to create that kind of change. Thus, the larger the preset amount, the more the curve moves toward the top right-hand corner of the graph. The smaller the preset amount, the smaller an effect an omitted variable would have to have in order to create that kind of change. Thus, the smaller the preset amount, the more the curve moves toward the lower left-hand corner of the graph. The effect of the preset cutoff can be seen in Figure 4.

The pairs of partial R^2s for the observed covariates, the +s in Figure 2, are

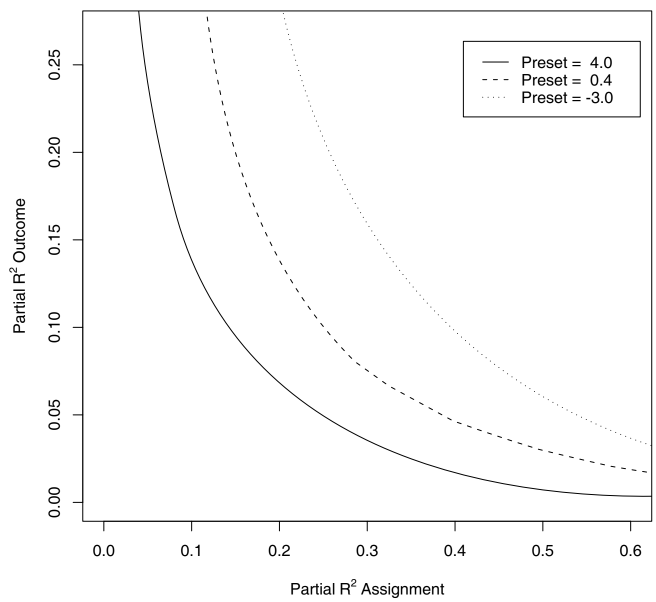

Figure 4: Curve placement in relation to the preset change.

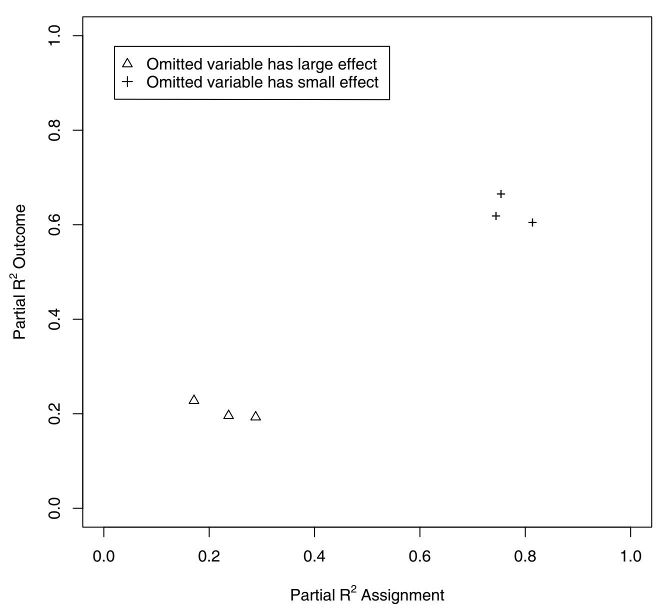

Figure 5: Curve placement in relation to the cutoff.

moved both by the size of $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, the coefficients on the observed covariates in the assignment and outcome equations, respectively, and the impact of any possibly omitted covariate on W and the dependent variable, Y. Figure 5 shows how the partial R^2s of the included variables react to the effect of the omitted variable. The larger that effect (on both the assignment equation and the outcome question), the more these values are found in the lower left-hand corner of the graph. The smaller that effect, the more these values are found the upper right-hand corner of the graph. These effects are to be expected given the definition partial R^2—the proportion of the previously unexplained variance in Y that is explained by the variable in question. The larger the effect of the omitted variable, the smaller the amount of the previously unexplained variance in Y that can be explained by the variable in question.

Table 2: Case 1 ($N = 500$)

Variable	Misspecified	Correct
W	6.97***	5.53***
X1	5.69***	6.01***
X2	5.91***	6.10***
X3	5.91***	5.91***
U		3.54***
Constant	1.24***	0.54***
R^2	0.88	0.90

*** $p \leq 0.01$.

Taking these two effects into account, we can construct examples where sensitivity analysis can mislead a researcher. Table 2 shows regression results from both a misspecified model and a correctly specified model. The data were generated according to Equations 1 and 2. In both equations, the variable of interest, W, has a large and significant effect on the dependent variable. The practitioner, of course, only sees the misspecified results.

The results of the sensitivity analysis on these data are in Figure 6. The plot indicates that there exists an omitted variable that would change the results. While an omitted variable does indeed exist, we know that it does not change the results on the variable of interest. Why this occurs is no mystery. First, the amount of the preset change was set quite low. The coefficient on W was

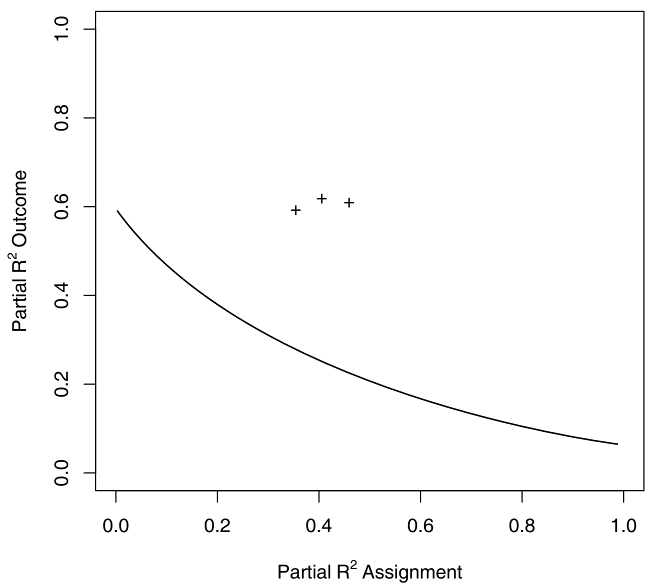

Figure 6: Example 1.

set at 5, and the preset change in $\hat{\tau}$ was set at 1. Thus, the curve is moved toward the lower left-hand corner of the plot. Second, while the omitted variable has a real effect, it is not large compared to the included variables (it is roughly the same size). Thus, the 3 pluses are in the upper right-hand corner of the plot. Note that the sensitivity analysis is not wrong; the code is doing what it is supposed to do. The analysis is misleading because we set the amount of preset change too low.

Table 3 shows regression results from another set of models: one misspecified model and one correctly specified. Again, the data were generated according to Equations 1 and 2. The misspecified equation shows that the variable of interest, W, has a large and significant effect. The correctly specified equation shows that W has neither a large nor a significant effect. Once again, the practitioner only sees the misspecified results.

The results of the sensitivity analysis are in Figure 7. The plot indicates that the observed results are fairly robust to omitted variables, although we know that actually is not the case. Again, the explanation is straightforward. The amount of preset change was set high—a change of 6 points. The curve is therefore pushed to the upper right-hand corner of the figure. Second,

Table 3: Case 2 ($N = 500$)

Variable	Misspecified	Correct
W	-8.69***	0.87
X1	3.82***	3.03***
X2	3.98***	3.02***
X3	3.60***	2.62***
U		12.64***
Constant	9.23***	0.01
R^2	0.50	0.78

*** $p \leq 0.01$.

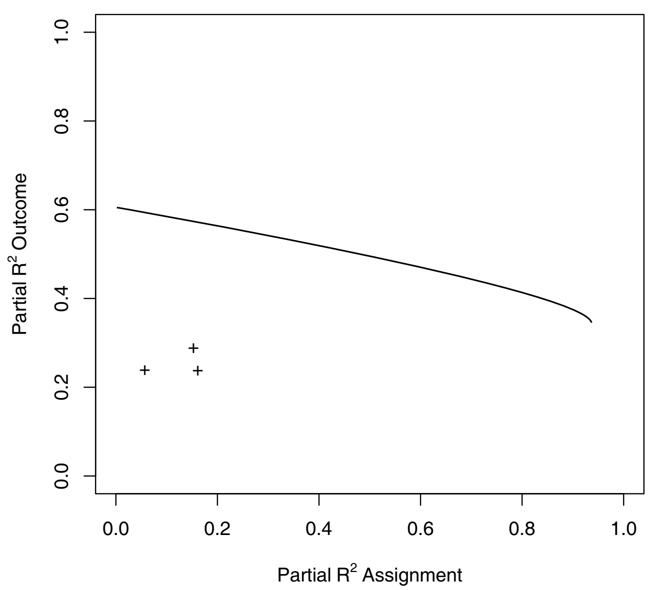

Figure 7: Example 2.

the effect of the omitted variable is large compared to the included variables (roughly four times the size). The pluses are therefore in the lower left-hand corner of the plot. Again, the analysis is not wrong *per se*; it is simply reflecting the state of the world (the size of the omitted variable), and the choices we as the researcher made.

7 Conclusion

Sensitivity analysis is a potentially useful tool in the statistical arsenal of political scientists. It can provide valuable information and obviates the need for failed strategies such as the inclusion of numerous control variables. The downsides are that the technique is computationally intensive (access to parallel computing is still rare in the field), and, as with any statistical technique, the procedure must be applied with care and sensitivity to achieve meaningful results. On balance, however, sensitivity analysis appears to complement quantitative political science, and it should see more use in the future.

To ensure that happens, Imbens-type sensitivity analysis should be, and can easily be, extended to other kinds of models. Two binary choice models, for example, would seem to be a straightforward extension. It would also be useful to fill in the remaining interesting cells in Figure 3 to get a more complete picture of how careful a researcher must be when using this procedure.

A Sensitivity Code for R

```
setClass('imbens',
representation(Sigma2='numeric', Tau='numeric',
                Gamma='numeric', VC='matrix'))



imbens <- function(alpha,delta,xmat,stvaln) {

    logl <- function(b,alpha,delta,X,W,Y) {
        tau <- b[length(b)]
        lns2 <- b[length(b)-1]
        s2 <- exp(lns2)
        gamma <- b[1:ncol(X)]
        beta <- b[(ncol(X)+1):(length(b)-2)]
        llik <- log(
            0.5*(1/sqrt(2*pi*s2)
               *
               exp(-(1/(2*s2))*(Y-tau*W-X%*%beta)^2)
               *
               ((exp(X%*%gamma))^W)/(1+exp(X%*%gamma)))
            +
            0.5*((1/sqrt(2*pi*s2))
               *
               exp(-(1/(2*s2))*(Y-tau*W-X%*%beta-delta)^2)
               *
               ((exp(X%*%gamma+alpha))^W)/(1+exp(X%*%gamma+alpha)))
            )
    sum(llik)
    }

    imbens.mle <- optim(stvaln,logl,hessian=F,method="BFGS",
    control=list(fnscale=-1,trace=1,maxit=2500,reltol=1e-17),
    alpha=alpha, delta=delta,X=xmat,W=W,Y=Y)

    sigma2 <- exp(imbens.mle$par[length(stvaln)-1])
```

```
    tau <- imbens.mle$par[length(stvaln)]
    gamma <- imbens.mle$par[2:ncol(xmat)]
    vc <- as.matrix(var(xmat[,-1]))

    stvaln <- imbens.mle$par

    result <- new('imbens',Sigma2=sigma2,Tau=tau,Gamma=gamma,VC=vc)
    class(result) <- 'imbens'
    result
}

setMethod('summary', signature(object='imbens'),
    definition=function(object, ...){
    Sigma2 <- object@Sigma2
    Tau <- object@Tau
    table <- cbind(Sigma2,Tau)
    colnames(table) <- c('Sigma^2', 'Tau')
    print(table)
    }
)
```

B Cluster Code for R

```
source("sensitivityMLE.R")

load("Data.RData")

constant <- 1
vars <- with(DG2.1,na.omit(cbind(constant,X1,X2,X3,W,Y)))
xmatt <- vars[,1:4]
W <- vars[,5]
Y <- vars[,6]

starting <- function(xmat){
  coef.lm <- lm(Y~xmat+W-1)$coefficients
  coef.logit <- glm(W~xmat-1)$coefficients
  coef.lm <- coef.lm[-length(coef.lm)]
  startv <- c(coef.logit,coef.lm,20,1)
  startv
}

stvalnt <- starting(xmatt)

aldelval <- cbind(rep(seq(-10,10,.2),101),rep(seq(-10,10,.2),each=101))

ncomb <- nrow(aldelval)

Izero <- imbens(0,0,xmatt,stvalnt)
S2zero <- Izero@Sigma2
R2wzero <- (t(Izero@Gamma)%*%Izero@VC%*%Izero@Gamma)/
           (t(Izero@Gamma)%*%Izero@VC%*%Izero@Gamma+((pi^2)/3))

rsquare <- function(i){

  Imb <- imbens(aldelval[i,1],aldelval[i,2],xmatt,stvalnt)
  Tau <- Imb@Tau
  S2 <- Imb@Sigma2
  R2w <- (t(Imb@Gamma)%*%Imb@VC%*%Imb@Gamma+(((aldelval[i,1])^2)/4))/
```

```
         (t(Imb@Gamma)%*%Imb@VC%*%Imb@Gamma+(((aldelval[i,1])^2)/4)
         +((pi^2)/3))
  R2p <- (S2zero-S2)/(S2zero)
  R2wp <- (R2w-R2wzero)/(1-R2wzero)
  R2p <- round(R2p,digits=3)
  R2wp <- round(R2wp,digits=3)
  vals <- cbind(Tau, R2p, R2wp, S2)
  vals

}

usecluster <- TRUE

if (usecluster) {
  library(snow)
  nclus <- 24
  clus1 <- makePVMcluster(nclus)
  clusterExport(clus1,c("aldelval","imbens","S2zero","R2wzero",
                "stvalnt","xmatt","W","Y"))
  clusterCall(clus1,source,file="sensitivityMLE.R")

  parsapply <- function(cluster,...) {
    parSapply(cluster,...)
  }
} else {
  parsapply <- function(cluster,...) {
    sapply(...)
  }
}

system.time(
            rsc <- parsapply(clus1,1:ncomb,rsquare)
            )


rsc.all <- cbind(aldelval,t(rsc))
colnames(rsc.all) <- c("alpha","delta","Tau","R2p","R2wp","S2")
save (rsc.all,file="Datarsc.RData")
```

```
Tau <- round(rsc.all[,"Tau"],1)

Tau.in <- (Tau==0)
rsc.in <- rsc.all[Tau.in,]

totalR2w <- R2wzero

partials <- function(xmatp){
  stvalnp <- starting(xmatp)
  imb <- imbens(0,0,xmatp,stvalnp)
  gammas <- as.matrix(imb@Gamma)
  vxb <- var(xmatp[,-1]%*%gammas)
  exclR2wp <- vxb/(vxb + pi^2/3)
  parR2w <- (totalR2w - exclR2wp)/(1 - totalR2w)

  s2p <- imb@Sigma2
  parR2 <- (s2p-S2zero)/s2p
  par <- cbind(parR2w,parR2)
  par
}

X1R2p <- partials(xmatt[,-2])
X2R2p <- partials(xmatt[,-3])
X3R2p <- partials(xmatt[,-4])


covp <- rbind(X1R2p,X2R2p,X3R2p)

ob1<-rsc.in[,4]
ob2<-rsc.in[,5]

save(ob1, file="DataR2p.RData")
save(ob2, file="DataR2wp.RData")

pdf(file="Datagraph.pdf")
plot(rsc.in[,5],rsc.in[,4],xlim=c(0,1),
     xlab=expression("Partial"~ R^2 ~ "Assignment"),
     ylim=c(0,1),ylab=expression("Partial" ~  R^2 ~ "Outcome"))
```

```
points(covp[,1],covp[,2],pch="+")
graphics.off()
```

References

Buhaug, Halvard, and Scott Gates. 2002. "The Geography of Civil War." *Journal of Peace Research* 39:417–433.

Clarke, Kevin A. 2005. "The Phantom Menace: Omitted Variable Bias in Econometric Research." *Conflict Management and Peace Science* 22:341–352.

Clarke, Kevin A. 2006. "Return of the Phantom Menace: Omitted Variable Bias in Econometric Research."

Cornfield, Jerome, William Haenszel, E. Cuyler Hammond, Abraham M. Lilienfeld, Michael B. Shimkin, and Ernst L. Wynder. 1959. "Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions." *Journal of the National Cancer Institute* 22:173–203.

Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81:945–960.

Imbens, Guido W. 2003. "Sensitivity to Exogeneity Assumptions in Program Evaluation." *The American Economic Review* 93:126–132.

Lalonde, Robert. 1986. "Evaluating the Econometric Evaluations of Training Programs with Experimental Data." *American Economic Review* 76:604–620.

Lin, D.Y., B.M. Psaty, and R.A. Kronmal. 1998. "Assessing the Sensitivity of Regression REsults to Unmeasured Confounders in Observational Studies." *Biometrics* 54:948–963.

Manski, Charles. 1990. "Nonparametric Bounds on Treatment Effects." *American Economic Review* 80:319–323.

Rosenbaum, Paul R. 2002. *Observational Studies*. New York: Springer.

Rosenbaum, P.R., and D.B. Rubin. 1983. "Assessing Sensitivity to an Unobserved Binary Covariate in an Observational Study with Binary Outcome." *Journal of the Royal Statistical Society, Series B* 45:212–218.

Rubin, Donald. 1974. "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies." *Journal of Educational Psychology* 66:688–701.