University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2012 *Paper* 303

Sensitivity Analysis for Causal Inference Under Unmeasured Confounding and Measurement Error Problems

Iván Díaz* Mark J. van der Laan†

*Division of Biostatistics, University of California, Berkeley, idiaz@jhu.edu

†Division of Biostatistics, University of California, Berkeley, laan@berkeley.edu

http://biostats.bepress.com/ucbbiostat/paper303

Sensitivity Analysis for Causal Inference Under Unmeasured Confounding and Measurement Error Problems

Iván Díaz and Mark J. van der Laan

Abstract

In this paper we present a sensitivity analysis for drawing inferences about parameters that are not estimable from observed data without additional assumptions. We present the methodology using two different examples: a causal parameter that is not identifiable due to violations of the randomization assumption, and a parameter that is not estimable in the nonparametric model due to measurement error. Existing methods for tackling these problems assume a parametric model for the type of violation to the identifiability assumption, and require the development of new estimators and inference for every new model. The method we present can be used in conjunction with any existing asymptotically linear estimator of an observed data parameter that approximates the unidentifiable full data parameter, and does not require the study of additional models.

1 Introduction

A common problem in statistics and causal inference is the need to draw inferences and test hypotheses about phenomena based on partially observed data. Examples of this situation are given by causal inference, measurement error, and missing data problems.

Causal parameters are defined in terms of a counterfactual stochastic process that measures what would happen to a subject under a system that enforces the value of the exposure variable. Since the researcher can only observe the outcome under the actually observed exposure, causal inference problems are often presented as regular inference problems with a partially observed counterfactual process. A fundamental assumption that allows the identification of parameters in missing data and causal inference problems is that the missingness or exposure mechanism does not depend on unmeasured factors that are causally related to the outcome. This assumption has been referred to as missing at random [MAR Rubin, 1976], non-ignorability, selection bias, and sequential randomization assumption [SRA, see e.g., van der Laan and Robins, 2003] depending on the context and working paradigm, and is closely related to a more general assumption for estimation in missing data problems called coarsening at random [CAR, see e.g., Heitjan and Rubin, 1991, Gill et al., 1997]. Under violations to the CAR assumption, the full data parameter of interest cannot be identified only from the distribution of the observed data without additional untestable assumptions.

In measurement error problems the exposure or outcome variable is usually completely unobserved, and a surrogate variable is observed in its place. Causal parameters are often non parametrically unidentifiable if the exposure or outcome of interest is not measured or measured with error, unless the error measurement model is known.

A commonly used method for identifying the otherwise unidentifiable parameter of interest is to assume parametric models that encode stringent assumptions about the data generating mechanism. This approach has been widely used for causal inference, missing data, and measurement error problems. Despite their generalized use, it is widely accepted that parametric models are very often misspecified, and that incorrectly specified parametric models lead to unmeasurable amounts of bias in estimation of the parameter of interest.

An alternative approach, employed in this article, is to define a sensitivity parameter that quantifies the severity of the violation to the identifiability assumption (e.g., SRA, MAR, measurement error), test the hypothesis of interest for each value of this parameter, and translate subject-matter expert knowledge on the plausibility of each value of the sensitivity parameter into a decision about the hypothesis of interest. This approach has been extensively used in the statistical literature, and was originally developed by Rotnitzky et al. [1998]. In their original article, Rotnitzky et al. assume parametric models for the outcome regression and missingness (or treatment) mechanism, and find regular asymptotically linear (RAL) estimators of the parameters of the outcome regression model. In the same paper, Rotnitzky et al. introduce a sensitivity analysis in terms of a parametric model for the missingness mechanism. They propose to estimate the outcome regression model under several values of an unidentifiable parameter α that relates the unobserved outcome to the missingness mechanism, and decide on the plausibility of each α value based on subject-matter expert knowledge. In their approach, $\alpha = 0$ corresponds to a situation in which the CAR assumption is satisfied, and $\alpha > 0$ measures deviations from CAR in a particular direction. For each $\alpha > 0$, they present an estimator of the desired full data parameter, whose inference requires the development of new methods compared to the estimators developed for $\alpha = 0$. This approach is further developed in a series of subsequent articles [Scharfstein et al., 1999, Robins et al., 1999, Rotnitzky et al., 2001, Scharfstein and Robins, 2002]. In particular, Scharfstein

et al. [1999] extend the approach to allow the specification of the missingness mechanism in terms of a Cox proportional hazards model, and Robins et al. [1999] present a detailed study of the properties of this methodology in a variety of data structures and semiparametric models. Because the sensitivity parameter α introduced in their work depends on the functional form of the (semi) parametric model posed for the missingness mechanism, it may not have an intelligible interpretation that is easily communicable to the subject-matter expert, who is expected to judge on the plausibility of each of its values.

In this paper we propose a sensitivity analysis for drawing inferences about unidentifiable parameters in which the sensitivity parameter is given by a bound on the difference between an observed data parameter and the full data parameter. The observed data parameter may be defined as the parameter that would identify the full data parameter if the appropriate assumption had been satisfied. For instance, in missing or counterfactual data structures, the observed data parameter may be given by the parameter that would identify the causal parameter under the CAR assumption. In error measurement problems, the observed data parameter may defined as the parameter that identifies the causal effect assuming that the surrogate variable equals the unobserved variable. The estimators and confidence intervals we present come at no additional cost, only the specification based on subject-matter expert knowledge of a plausible range of values for the difference between the causal and observed data parameter.

The paper is organized as follows. In section 2 we describe the general estimation problem and present two motivating examples: the first one about effectiveness of a treatment for the Chagas disease under informative dropout, and the second one about the effect of physical activity on mortality in the elderly under measurement error. In section 3 we describe the proposed sensitivity analysis, and present the results of the analysis of the two examples. Finally, in section 4 we present a brief comparison with methods in the existing literature.

2 Data and examples

Let $X = \{X_1, \ldots, X_p\}$ be a set of observed temporally ordered random variables governed by a non parametric structural equation model (NPSEM) described by three elements: a set of functions $f_1, \ldots, f_p$, a set of parents of X_i denoted by $pa(X_i) \subseteq \{X_1, \ldots, X_{i-1}\}$, and a set of unobserved random variables $U_1, \ldots, U_p$. The NPSEM is given by $X_i = f_i(pa(X_i), U_i) : \ i = 1, \ldots, p$ [Pearl, 2000]. Let $F_0 \in \mathcal{M}^f$ denote the joint distribution of (U, X), and let $P_0 \in \mathcal{M}$ denote the distribution of X. Here $\mathcal{M}^f$ and $\mathcal{M}$ are models describing the set of allowed distributions for F_0 and P_0, respectively. The target parameter of the full data distribution is given by $\Psi^f : \mathcal{M}^f \to \mathbb{R}$, and its true value is denoted by $\psi_0^f = \Psi^f(F_0)$. Identification assumptions usually refer to the set of assumptions posed on the model $\mathcal{M}^f$ that must be made in order to write $\Psi^f(F_0) = \Psi(P_0)$ for some map $\Psi : \mathcal{M} \to \mathbb{R}$. For example, in a missing data problem where W is a set of covariates and Y is an outcome observed only when a missingness indicator C equals one, the assumption $Y \perp\!\!\!\perp C|W$ (often referred to as missing at random (MAR)) is necessary to prove that $E(Y) = E_W\{E(Y|C = 1, W)\}$, where the right hand side quantity is estimable from the observed data alone. Identification of parameters in this context is discussed in more detail by Pearl [2000], for example.

If the previous identifiability assumptions do not hold, the parameter Ψ_0^f cannot be estimated from a sample of the distribution P_0 of observed data. A possible solution, exploited in this paper, is to define an observed data parameter $\Psi : \mathcal{M} \to \mathbb{R}$ as an approximation to Ψ^f, and perform a sensitivity analysis in terms of a sensitivity parameter δ_0 satisfying $\psi_0 - \psi_0^f \leq \delta_0$. It is important to set up the problem in terms of a sensitivity parameter δ_0 that is equal to zero under CAR, so that it provides a true measure

of the amount of unmeasured confounding (see example 1 below). In a causal inference problem, if the randomization assumption is violated, Ψ may be defined by the parameter that would otherwise identify Ψ^f. If, on the other hand, the source of unidentifiability is measurement error in the exposure variable, Ψ may be given by the parameter that identifies Ψ^f under the (false) assumption that the surrogate exposure is equal to the unobserved exposure. The null hypothesis $H_0 : \psi_0^f \leq 0$ can then be tested for each value of δ_0, empowering the researcher with a tool to make decisions according to his subject-matter knowledge on the parameter δ_0. The amount and quality of knowledge available about δ_0 is specific to the problem, and depends on the definition of the parameter Ψ. Some examples will be discussed next.

Example 1. The Chagas disease is estimated to affect about 8 million people in Latin American endemic areas, with an additional 400.000 in the U.S. and Spain. Despite the public importance of the disease, few rigorous research studies exists about the efficacy of currently available treatments. Because of the long incubation periods of the disease (up to 30 years), the cost of the a well designed study to assess the efficacy of different treatments is prohibitive, and a meta analysis of the existing literature is necessary. The existing literature includes several observational studies with large number of drop-outs, lost-to-follow-up, and unmeasured confounding, as well as statistical methods whose conclusions are not reliable because they are not up to the state of the art in causal inference for observational studies. A thorough review of the literature about the Chagas disease rendered 19 studies comprising about 520 patients. A standard meta analysis of these studies was not possible because of several methodological issues: *a*) Many studies did not have a control group; *b*) There is very limited information about treatment allocation in all the manuscripts; *c*) None of the studies reported baseline/confounder information (e.g., baseline health status, exposure to the vector, etc.); and *d*) Many of the studies presented large numbers of lost to follow-up and drop-outs. In terms of treatment assignment, the previous points imply that these studies cannot be considered randomized, and that there may be an important amount of unmeasured confounding (e.g., treatment was allocated according to unmeasured baseline status of the infection). Additionally, since patients drop out of the study as a consequence of the worsening or improving of their health condition, lost to follow-up and drop out were often related to the unobserved outcome.

This problem can be set up in terms of the notation of this paper as follows. Assume the following NPSEM:

$$
\begin{aligned}
S &= f_S(U_S) \\
W &= f_W(S, U_W) \\
A &= f_A(S, W, U_A) \\
\Delta &= f_\Delta(S, W, A, U_\Delta) \\
Y &= f_Y(S, W, A, U_Y) \\
Y^* &= \begin{cases} \text{NA} & \text{if } \Delta = 0 \text{ and } A = 0 \\ \Delta Y & \text{otherwise} \end{cases}
\end{aligned}
$$

where S denotes the study, W denotes a set of unmeasured baseline characteristics, A denotes treatment allocation, Δ denotes a missingness indicator, Y denotes an indicator of cure, and Y^* is the partially observed outcome Y. Note that $Y^* = 0$ whenever $\Delta = 0$ and $A = 1$, so that we impute all treated missing patients as "not cured". The objective of this analysis is to estimate the causal effect among the

treated, given by

$$\begin{aligned}\Psi^f(F) =& E_F(Y_1 - Y_0|A=1)\\ \equiv& \Psi^{f,1}(F) - \Psi^{f,0}(F),\end{aligned}$$

where

$$\begin{aligned}\Psi^{f,a}(F) =& E_F(Y_a|A=1)\\ =& \sum_s E_F(Y_a|S=s,A=1)P(S=s|A=1),\end{aligned}$$

and $Y_a = f_Y(S, W, a, U_Y)$ is a counterfactual of Y^* obtained by setting $(\Delta, A) = (1, a)$ with probability one. Define the observed data parameter

$$\Psi(P) = E_P(Y^*|A=1) - E(Y|\Delta=1, A=0).$$

Under the randomization assumption that $(\Delta, A) \perp\!\!\!\perp Y_a$ for $a \in \{0,1\}$, $\Psi^{f,a}(F)$ would be identified in terms of the observed data distribution as $\Psi^a(P) = E_P(Y^*|\Delta=1, A=a)$. This randomization assumption is violated because both the missingness mechanism and the treatment mechanism depend on the unobserved variable W that is also causally related to the outcome. Note that under our imputation approach, $E_P(Y^*|A=1)$ is a conservative approximation of $\Psi^{f,1}$. That is

$$E(Y^*|A=1) = E(\Delta Y_1|A=1) \leq E(Y_1|A=1). \tag{1}$$

Since the randomization assumption does not hold and Ψ^f is not estimable from the observed data, we will approximate it with the observed data parameter $\Psi(P)$, and establish an upper bound δ_0 on $\psi_0 - \psi_0^f$ (here $\psi_0 = \Psi(P_0)$ and P_0 is the true distribution of the observed data) which will be used as sensitivity parameter. For that purpose, we note that given (1), we have

$$\begin{aligned}\psi_0 - \psi_0^f =& \{E_{F_0}(Y_0|A=1) - E_{P_0}(Y|\Delta=1,A=0)\} + \{E_{P_0}(Y^*|A=1) - E_{F_0}(Y_1|A=1)\}\\ \leq& E_{F_0}(Y_0|A=1) - E_{P_0}(Y|\Delta=1,A=0)\\ \equiv& \delta_0.\end{aligned}$$

Note that $\delta_0 = 0$ under CAR. We will now test the hypothesis of no effect of treatment for a range of plausible values of δ_0. In this example, knowledge about the plausibility of each value δ_0 can be obtained by finding independent bounds on each of its terms. One one hand, the parameter $E_{F_0}(Y_0|A=1)$ is the probability of cure for the treated population had they not been treated (also called probability of spontaneous cure from now on) and can be bounded based on subject-matter expert knowledge. On the other hand, $E_{P_0}(Y|\Delta=1, A=0)$ will be replaced by its most conservative value: zero.

Example 2. Tager et al. [1998] followed a group of people over 55 years of age living around Sonoma, CA, over a time period of about ten years as part of a longitudinal study of physical activity and fitness (Study of Physical Performance and Age Related Changes in Sonomans - SPPARCS). The goal in analyzing the data that were collected as part of this study is to examine the effect of baseline vigorous physical activity on subsequent five-year all-cause mortality. Vigorous physical activity is often measured in terms of Metabolic Equivalents of Tasks (MET), which is a physiological measure expressing the energy cost

of physical activities. One MET is approximately equal to the energy produced per unit surface area of an average person sitting quietly.

Since the actual amount of vigorous physical activity in METs is not measurable, Tager et al. used a questionnaire in which participants were asked how many hours during the past seven days they had participated in twelve common vigorous physical activities, such as jogging, swimming, bicycling on hills, or racquet-ball. A surrogate measure for LTPA (Leisure Time Physical Activity), is thus given by a continuous score based on the number of hours that the participants reported engagement in these activities during the last seven days, and the approximate intensity values in METs of each activity.

These data were previously analyzed by Bembom and van der Laan [2007] and Díaz and van der Laan [2012]. In particular, Díaz and van der Laan considered the question of assessing the causal effect of a population intervention that causes a shift of 12 METs in the distribution of vigorous physical activity—corresponding, for instance, to bicycling during three hours at less than 10MPH per week—assuming that the exposure was measured without error. Díaz and van der Laan showed that the effect of such intervention would be a reduction of the risk of all cause mortality in the elderly of 1.79% (0.71%).

This problem can be set up as follows. Assume that the following NPSEM holds:

$$W = f_W(U_W);\ A = f_A(W, U_A);\ A^* = f_{A^*}(A, W, U_{A^*});\ Y = f_Y(A, W, U_Y), \tag{2}$$

where A is the (unmeasured) true physical activity in METs, A^* is the surrogate physical activity discussed above, W represents a set of measured confounders (e.g., gender, age, health status, smoking, etc.), Y is an indicator of subsequent five-year mortality, and U_W, U_A and U_Y are exogenous random variables. For the sake of discussion assume that the randomization assumption holds (i.e., $U_A \perp\!\!\!\perp U_Y|W$), but the analysis presented here is also valid without it.

As explained by Díaz and van der Laan [2012], the intervention of interest in this case can be defined by the intervened NPSEM:

$$W = f_W(U_W);\ A_\epsilon = f_A(W, U_A) + \epsilon;\ A^*_\epsilon = f_{A^*}(A_\epsilon, W, U_{A^*});\ Y_\epsilon = f_Y(A_\epsilon, W, U_Y), \tag{3}$$

with $\epsilon = 12$, and the effect of such intervention on mortality would be identified by

$$\Psi^f(F) \equiv E_F(Y - Y_\epsilon) = E_P\{Y - E_P(Y|A + \epsilon, W)\},$$

if A had been measured. Clearly, since A is not observed the latter quantity is not estimable. Therefore, we will define the observed data parameter

$$\Psi(P) = E_P\{Y - E_P(Y|A^* + \epsilon, W)\}, \tag{4}$$

and perform a sensitivity analysis in terms of the difference

$$\delta_0 = \psi_0 - \psi_0^f = E_{P_0}\{E_{P_0}(Y|A + \epsilon, W) - E_{P_0}(Y|A^* + \epsilon, W)\},$$

which measures the difference between the true effect and the effect as approximated by $\Psi(P_0)$. As discussed by Hernán and Cole [2009], A^* cannot possibly have a causal effect on Y, and "the assumption implicit in many epidemiologic analyses is that the association between A^* and Y approximates the causal relation between A and Y". We therefore present a way of explicitly defining such association, and quantifying the extent to which that implicit assumption can be considered to hold.

3 Sensitivity analysis

In a general formulation of causal inference problems it is often of interest to establish whether an intervention of interest has a positive effect on the population. This hypothesis testing problem can be described in terms of the hypothesis system

$$H_0 : \psi_0^f \leq 0 \quad \text{vs} \quad H_1 : \psi_0^f > 0,$$

for an appropriately defined full data parameter $\psi_0^f = \Psi^f(F_0)$. Let P_n denote the empirical distribution of the observed data O, and let $\hat{\Psi} : \mathcal{M} \to \mathbb{R}$ be an asymptotically linear estimator of ψ_0. That is

$$\psi_n - \psi_0 = \frac{1}{n}\sum_{i=1}^{n} D(O_i) + o_P(1/\sqrt{n}),$$

where $\psi_n = \hat{\Psi}(P_n)$, and D is the influence function of $\hat{\Psi}$: a mean zero function of O that depends on P_0. Thus, the variance of an asymptotically linear estimator can be estimated as

$$\sigma_n = \frac{1}{n^2}\sum_{i=1}^{n} D_n(O_i)^2$$

where D_n is an estimator of D. We will use the statistic $T_n = \psi_n/\sigma_n$ to test the hypothesis H_0, rejecting H_0 if $T_n > c$.

3.1 Cut-off based on hypothesized value of sensitivity parameter.

Define a sensitivity parameter δ_0 such that $\delta_0 \geq \psi_0 - \psi_0^f$. In particular, δ_0 may be defined as $\psi_0 - \psi_0^f$, but the methodology that we present can also be used for an upper bound on $\psi_0 - \psi_0^f$. The interpretation of such upper bound may be more intelligible by a subject-matter expert, as in our example 1 above.

Consider the type I error probability

$$P_{H_0}(T_n > c) = P_{H_0}\{\psi_n - \psi_0 + (\psi_0 - \psi_0^f) + \psi_0^f > c\sigma_n\}.$$

Since $\psi_0^f \leq 0$ under H_0, we have

$$P_{H_0}(T_n > c) \leq P_{H_0}\{\psi_n - \psi_0 + (\psi_0 - \psi_0^f) > c\sigma_n\},$$

and since $\psi_0 - \psi_0^f \leq \delta_0$, we obtain

$$\begin{aligned} P_{H_0}(T_n > c) &\leq P_{H_0}\{\psi_n - \psi_0 + \delta_0 > c\sigma_n\} \\ &= P_{H_0}\{Z_n > c - d_0\}, \end{aligned}$$

where

$$Z_n = \frac{\psi_n - \psi_0}{\sigma_n} \xrightarrow{d} N(0,1) \text{ and } d_0 = \frac{\delta_0}{\sigma_n}.$$

We need a value c such that $c - d_0 \geq z_\alpha$, where $P(Z_n > z_\alpha) = \alpha$, so that the probability of type I error is less than or equal to α. Let t be the observed value of T_n. Consider $\Phi\{c - d_0\} = \alpha$, where Φ is the standard normal distribution. For a given value d_0, we can solve this equation in c giving $c(d_0) = z_\alpha + d_0$. Therefore, given a hypothesized value of d_0, we will reject the null hypothesis of no treatment effect if $t > c(d_0)$.

3.2 Determining all hypothesized values of sensitivity parameter under which the null hypothesis is rejected.

Let d^* be the value that solves $c(d_0) = t$, which is given by $d^* = t - z_\alpha$. This value d^* is critical since all the values of d_0 below d^* will lead to a rejection of the null hypothesis of no treatment effect. Equivalently, all values of δ_0 such that

$$\delta_0 \leq \sigma_n(t - z_\alpha)$$

result in rejection of the null hypothesis of no treatment effect at a type I error of α.

Thus, if it can be argued, based on subject matter knowledge and the observed data that $\delta_0 \leq \sigma_n(t - z_\alpha)$, then the test would reject the null hypothesis of no treatment effect with a probability of type I error less than or equal to α. Alternatively, a plot of $\delta(\alpha) = \sigma_n(t - z_\alpha)$ shows the bounds on δ_0 that lead to rejection of the null hypothesis of no effect at each probability of type one error α.

Example 1. In our working example, the test statistic T_n will be given by

$$T_n = \frac{\psi_n}{\sigma_n},$$

where ψ_n is the non parametric estimator of

$$\Psi(P) = \sum_s \Big\{ E(Y^*|S = s, A = 1)P(S = s|A = 1) - E(Y|S = s, \Delta = 1, A = 0)P(S = s|\Delta = 1, A = 0) \Big\}$$

obtained by replacing the population probabilities with sample probabilities. The standard error is computed as

$$\sigma_n = \sqrt{\sum_s \left\{ w_{s,1}^2 \frac{\psi_{n,s}^1(1 - \psi_{n,s}^1)}{n_{s,1}} + w_{s,0}^2 \frac{\psi_{n,s}^0(1 - \psi_{n,s}^0)}{n_{s,0}} \right\}},$$

where $n_{s,1}$ is the number of treated patients in study s, $n_{s,0}$ is the number of observed untreated patients in study s, $w_{s,a} = \hat{P}(S = s|A = a) = n_{s,a}/\sum_s n_{s,a}$, and $\psi_{n,s}^1$ is the sample probability of $Y^* = 1$ among treated patients, and $\psi_{n,s}^0$ is the sample probability of $Y = 1$ among observed, untreated patients. The computation of the quantities defined in the previous section results in the values $t = 23.35$, $\sigma_n = 0.0202$, and $\delta(0.05) = 0.4394$.

For the sake of interpretation we will now translate the previous bound on δ_0 into a bound on the more interpretable quantity $E_{F_0}(Y_0|A = 1)$. Recall that

$$\delta_0 = E_{F_0}(Y_0|A = 1) - E_{P_0}(Y|\Delta = 1, A = 0).$$

Replacing $E_{P_0}(Y|\Delta = 1, A = 0)$ by its most conservative value 0 implies that subject-matter knowledge that $E(Y_0|A = 1) \leq 0.4394$ would be enough to reject the hypothesis of no effect of treatment at a type I error probability smaller than $\alpha = 0.05$. In many of the studies it was documented that for ethical reasons the treatment was assigned to sicker patients, which would imply that $E(Y_0|A = 1) \leq E(Y_0|A = 0)$. The observation that $E_{P_0}(Y|\Delta = 1, A = 0)$ was estimated at 0.0184 is thus strong evidence that $E(Y_0|A = 1) \leq 0.4394$ holds.

Figure 1(a) shows the different bounds on δ_0 for rejecting the null hypothesis with different probabilities of type I error.

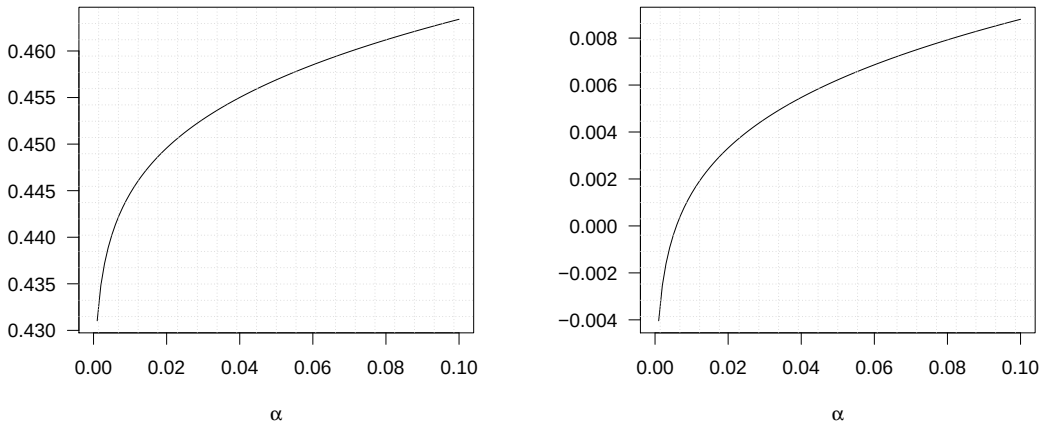

(a) Example 1.

(b) Example 2.

Figure 1: Upper bounds on δ_0 for which null hypothesis of no treatment effect is rejected at different probabilities of type I error α.

Example 2. Díaz and van der Laan [2012] presented a targeted minimum loss based estimator [TMLE, see e.g., van der Laan and Rubin, 2006, Rose and van der Laan, 2011] for the parameter $\Psi(P)$ defined in (4). TMLE is a loss-based semi-parametric estimation method that yields a substitution estimator of a target parameter of the probability distribution of the data that solves the efficient influence curve estimating equation, and thereby yields a double robust locally efficient estimator of the parameter of interest, under regularity conditions.

In the original paper, Díaz and van der Laan [2012] found $\psi_n = 0.0179$ and $\sigma_n = 0.0071$. Figure 1(b) shows the function $\delta_0(\alpha) = \sigma_n(t - z_\alpha)$, defined as the bound on δ_0 that leads to a rejection of the hypothesis of no effect for different probabilities of type I error α. This figure shows that if the difference between the probabilities of death $E\{E(Y|A + 12, W)\}$ and $E\{E(Y|A^* + 12, W)\}$ is as large as 0.8%, then the hypothesis of no effect of the intervention is rejected at a type I error probability of less than 0.1. In particular, if $\psi_0^f = \psi_0$, then the hypothesis of no effect of the intervention is rejected with a probability of type I error smaller than 0.01.

3.3 Estimator and confidence interval for each value of sensitivity parameter

As a by-product of the procedure described above, analogue to the point and interval estimators presented by Robins, Scharfstein, and Rotnitzky in their series of articles, if the sensitivity parameter is defined as $\delta_0 = \psi_0 - \psi_0^f$, it is possible to construct an asymptotically linear estimator of the causal parameter as a function of each hypothesized value δ_0. That is, if $\delta_0 = \delta$ is known, the estimate of ψ_0^f is given by $\psi_n^f(\delta) = \psi_n - \delta$, with standard error given by $\sigma_n^f = \sigma_n$. A $(1 - \alpha)100\%$ confidence interval is thus given by $(\psi_n - \delta - z_{\alpha/2}\sigma_n, \psi_n - \delta - z_{\alpha/2}\sigma_n)$, and can be computed for a range of values δ defined based on subject-matter expert knowledge at no additional analytical or computational cost.

3.4 A generalization of the sensitivity analysis

The sensitivity analysis presented in this section can be generalized as follows. Consider a function f such that it can be established that $\psi_0^f \geq f(\psi_0, \delta_0)$ for an observed data parameter ψ_0 (that needs not approximate ψ_0^f) and a sensitivity parameter δ_0. In such case, the untestable hypothesis of no causal effect $H_0 : \psi_0^f \leq 0$ implies the hypothesis $H_0^1 : f(\psi_0, \delta_0) \leq 0$, which can be tested for a range of user-given values of the sensitivity parameter δ_0. Rejecting H_0^1 with a probability of type I error smaller than α will thus lead to rejection of H_0 with at most the same probability of type I error.

The analysis presented above for example 1 is equivalent to this analysis with $f(\psi_0, \delta_0) = \psi_0 - \delta_0$, $\psi_0 = E(Y^*|A = 1) - E(Y|\Delta = 1, A = 0)$, and $\delta_0 = E(Y_0|A = 1) - E(Y|\Delta = 1, A = 0)$. Notice, however, that the choice of observed data parameter and sensitivity parameter is not unique, since the choice $\psi_0^* = E(Y^*|A = 1)$, and $\delta_0^* = E(Y_0|A = 1)$ would also satisfy the condition $\psi_0^f \geq f(\psi_0^*, \delta_0^*)$. Thus, the hypothesis $H_0 : \psi_0^f \leq 0$ implies the hypothesis $H_0^1 : E(Y^*|A = 1) \leq \delta_0^*$ which can be tested for a set of fixed values of δ_0^*. The latter choice of sensitivity and observed data parameters has two important consequences on the results of the analysis: first, for a given value of δ_0 the sample size is reduced because only the treated patients are used to test the hypothesis H_0; and second, the sensitivity parameter does not equal zero if the CAR assumption holds, which means that additional information about the location of the sensitivity parameter must be known. The convenience of a specific choice of sensitivity and observed data parameter must therefore be determined taking into account these consequences, as well as the simplicity of interpretation of the sensitivity parameter and the amount of information available about it.

4 Comparison with existing methods

Existing approaches for sensitivity analysis under violations to the CAR assumption require the specification of a parametric model for the treatment mechanism or measurement error that relates the observed exposure to the unmeasured counterfactuals through a sensitivity parameter α_0 [see e.g., Rotnitzky et al., 1998, Scharfstein et al., 1999, Robins et al., 1999, Rotnitzky et al., 2001, Scharfstein and Robins, 2002]. Estimation of the parameter of interest ψ_0^f is then carried out for a set of possible values of α_0 through a model-dependent and often complex estimator, and a decision about the original question of interest is made based on these estimates together with auxiliary subject-matter expert information about α_0. However, the parameter α_0 usually has a complex interpretation that is very difficult to communicate to the subject-matter expert, who is expected to provide auxiliary information about it. In addition, since parametric models are seldom correct, these sensitivity parameters are often ill-defined and do not even have the complex interpretation that is originally intended. Furthermore, even if these issues are overcome and a parametric model is known to contain the true measurement error or treatment mechanism, the development and computation of estimators for the causal quantity of interest are often overly and unnecessarily complex, and lead to sub-optimal estimators that do not enjoy desirable statistical properties like double robustness, which only hold in the case $\alpha = 0$ (i.e., when the CAR assumption holds). To illustrate some of these ideas, consider the case of example 2 above. It is common practice to assume a normal mean zero model for the measurement error:

$$A^* = A + \nu, \quad \nu \sim N(0, \alpha_0^2), \quad \nu \perp\!\!\!\perp A \quad \nu \perp\!\!\!\perp W. \tag{5}$$

The parameter of interest ψ_0^f can then be written as a function of the standard deviation α_0 of the measurement error:

$$\psi_0^f = E_0\{Y - E_0(Y|A + \epsilon, W)\},$$

where the outer expectation is taken with respect to the joint density of (A, W), and the density g_A of A conditional on W can be identified in terms of the density g_{A^*} of A^* given W. Though possible, this identification result is often a hard problem, since it requires using inverse characteristic functions. From model (5), note that the characteristic function of A conditional on W is given by

$$\varphi_A(t|W) = \frac{\varphi_{A^*}(t|W)}{\varphi_\nu(t)},$$

where $\varphi_{A^*}(t|W)$ and $\varphi_\nu(t)$ are the characteristic functions of A^* and ν conditional on W given by

$$\varphi_{A^*}(t|W) = \int_{\mathbb{R}} g_{A^*}(a|W) \exp(ita) da$$
$$\varphi_\nu(t) = \exp\left(-\frac{1}{2}\alpha_0^2 t\right).$$

Thus, using the inversion formula the density of A can be identified from the density of A^* (conditional on W) as

$$g_A(a|W) = \frac{1}{2\pi} \int_{\mathbb{R}} \exp\left(\frac{1}{2}\alpha_0^2 t - iat\right) \varphi_{A^*}(t|W) dt. \tag{6}$$

The expectation $E_0(Y|A, W)$ can be identified using similar arguments. It can be seen from these equations that the parameter (and thereby its efficient influence function) depends on α_0 in a very complex way, and thus a new evaluation of the estimator must be performed for each hypothesized value of α_0. In addition, because evaluation of the parameter requires the use of computer-intensive numerical methods for the computation of the integral in (6), this complex dependence of the parameter on α_0 represents an important disadvantage of this method compared to the method presented in section 3 above, in which an existing estimate of the observed data parameter was used to carry out the sensitivity analysis and no additional analytical or computational tools were necessary.

The approach presented in this paper overcomes the difficulties described above in various fronts. First, the sensitivity parameter is defined as a difference between the target quantity and an estimable quantity, thus providing a more straightforward account of the amount of unmeasured confounding as relevant for the specific outcome and parameter of interest. Second, since the parameter of interest is defined as $\psi_0^f = \psi_0 - \delta_0$, the estimate of ψ_0 has to be computed only once, providing an important advantage compared to methods in which an estimate must be computed for each value of the sensitivity parameter. Third, existing optimal estimators of ψ_0 may be used, avoiding the analytical complexity that α-specific estimators in the existing literature entail. Fourth, for each value of δ_0, the estimator of ψ_0^f enjoys all the statistical properties of the estimator of ψ_0 (e.g., double robustness and efficiency). Lastly, since the measurement error model presented in NPSEM (2) is completely non-parametric, the validity of the results of a sensitivity analysis for δ_0 is not compromised by possible misspecifications of the measurement error model.

An important similarity between both approaches is the need for auxiliary subject-matter information on a range of plausible values for the sensitivity parameter (either δ_0 or α_0). In our measurement error example, it may be the case that the subject-matter expert has auxiliary information about the standard

deviation α_0 of the measurement error, but not about the amount of unmeasured confounding δ_0. In the following subsection we will present a way of translating knowledge about a parameter α_0 in a parametric model for the treatment or measurement error mechanism into knowledge about the amount of unmeasured confounding $\delta_0 = \psi_0 - \psi_0^f$, so that the sensitivity analysis can be performed as presented in section 3.

4.1 Translating knowledge about α_0 into knowledge about δ_0

In certain situations it is possible that the amount of unmeasured confounding as measured by the difference $\psi_0 - \psi_0^f$ is completely unknown, and that subject-matter knowledge about it can only be obtained through a sensitivity parameter α_0 in a parametric model. For instance, Scharfstein et al. [1999] propose a sensitivity analysis in which knowledge about a sensitivity parameter α_0 that represents the log hazard ratio for drop-out between subjects with the same baseline covariates but who differ by 1 in their outcome is available. In our example 2, if the measurement error is known to follow model (5), a range of plausible values for α_0 may also be known.

In these cases, the range of plausible values for the parameter α_0 can be mapped into a range of plausible values for the sensitivity parameter $\delta_0 = \psi_0 - \psi_0^f$ by noting that the full data parameter can be written as a function of the observed data distribution and the parameter α as $\delta(\alpha, P_0) = \Psi(P_0) - \Psi^f(\alpha, P_0)$. In example 2, assuming model (5), application of this formula results in

$$\delta(\alpha, P_0) = E_{W,0} \int E_0(Y|A+\epsilon, W) g_A(a|W)\, da - E_0\{E_0(Y|A^*+\epsilon, W)\}, \tag{7}$$

where g_A is a function of g_{A^*} given by (6), and $E_0(Y|A+\epsilon, W)$ can be written as a function of $E_0(Y|A^*+\epsilon, W)$ and g_{A^*}. Note that (7) depends on α and $Q(P_0) = (E(Y|A^*, W), g_{A^*}(A^*|W), Q_W(W))$, where Q_W denotes the marginal distribution of W. For a given range of plausible values for α, the corresponding range of plausible values for δ can be found by plugging in α and a sensible estimate of $Q(P_0)$. In our example, estimates of $E(Y|A^*, W)$ and $g_{A^*}(A^*|W)$ may be obtained through machine learning methods (such as the super learner [van der Laan et al., 2007]), and the marginal distribution of W may be estimated with its empirical counterpart. Once the range of plausible values for δ_0 is obtained, the sensitivity analysis can be carried out as described in section 3 making full use of the advantages previously described.

Acknowledgements

We would like to thank Glenn Rice and Danong Chen from MetronomX for kindly making available the data set that was used in example 1 and for many valuable discussions about the Chagas disease that led to the development of the technique presented in this paper.

Mark van der Laan was supported by the NIH research grant (R01 A1074345), titled "Targeted Empirical Super Learning in HIV Research".

References

O. Bembom and M.J. van der Laan. A practical illustration of the importance of realistic individualized treatment rules in causal inference. *Electronic Journal of Statistics*, 2007.

Iván Díaz and Mark van der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, 68(2):541–549, 2012. ISSN 1541-0420. doi: 10.1111/j.1541-0420.2011.01685.x. URL http://dx.doi.org/10.1111/j.1541-0420.2011.01685.x.

R.D. Gill, M.J. van der Laan, and J.M. Robins. Coarsening at random: characterizations, conjectures and counter-examples. In D.Y. Lin and T.R. Fleming, editors, *Proceedings of the First Seattle Symposium in Biostatistics*, pages 255–94, New York, 1997. Springer Verlag.

D.F. Heitjan and D.B. Rubin. Ignorability and coarse data. *Annals of statistics*, 19(4):2244–2253, December 1991.

M. A. Hernán and S. R. Cole. Invited commentary: causal diagrams and measurement bias. *American Journal of Epidemiology*, 170(8):pp. 963–954, 2009.

J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.

J. M. Robins, A. Rotnitzky, and D. O. Scharfstein. 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*, IMA Volumes in Mathematics and Its Applications. Springer, 1999.

S. Rose and M.J. van der Laan. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

A. Rotnitzky, D. Scharfstein, S. Ting-Li Su, and J. Robins. Methods for conducting sensitivity analysis of trials with potentially nonignorable competing causes of censoring. *Biometrics*, 57(1):103–113, 2001. ISSN 1541-0420.

Andrea Rotnitzky, James M. Robins, and Daniel O. Scharfstein. Semiparametric regression for repeated outcomes with nonignorable nonresponse. *Journal of the American Statistical Association*, 93(444): 1321–1339, 1998.

D.B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.

D. O. Scharfstein and J. M. Robins. Estimation of the failure time distribution in the presence of informative censoring. *Biometrika*, 89(3):pp. 617–634, 2002.

D. O. Scharfstein, A. Rotnitzky, and J. M. Robins. Adjusting for nonignorable drop-out using semiparametric non-response models (with discussion). *Journal of the American Statistical Association*, 94: 1096–1146, 1999.

I.B. Tager, M. Hollenberg, and W. Satariano. Self-reported leisure-time physical activity and measures of cardiorespiratory fitness in an elderly population. 1998.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.

M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.