Review

Tyler J. VanderWeele
Departments of Epidemiology and Biostatistics
Harvard T.H. Chan School of Public Health

Propensity Scores

Propensity score: For treatment A and covariates C, the propensity score S(c) is defined by $S(c)=P(A=1|C=c)$.

In other words, the propensity score is the probability of a particular individual being assigned treatment conditional on some set of measured variables.

Example: Job training program is treatment; covariates are baseline income, age, gender, baseline health (ideally a sufficient set of confounders)

The propensity score has certain theoretical properties which make it particularly useful in trying to control for various confounders.

Propensity Scores

The propensity score has certain balancing properties.

Result (Rosenbaum and Rubin, 1983):

$$\text{If } S = P(A = 1|C) \text{ then } P(C|A = 1, S) = P(C|A = 0, S) \text{ i.e. } C \coprod A|S$$

Within strata of S the treated subjects and the control subjects will have the same distribution of the measured covariates C.

In other words, if we stratify on the propensity score then within each stratum, all the measured covariates tend to be balanced between the treatment and control groups.

In effect, as far as concerns *the measured covariates*, it is as though we have conducted a randomized trial within each stratum.

Unlike a randomized trial, however, stratification by use of a propensity score balances only those covariates that have been measured.

It does not balance those covariates that have not been measured.

Propensity Scores

In fact, if treatment is unconfounded conditional on C then it is also unconfounded conditional on S

Theorem (Rosenbaum and Rubin, 1983): $\text{If } Y_a \coprod A|C \text{ then } Y_a \coprod A|S$

If it suffices to adjust for C to estimate the effect of treatment then it suffices also instead to adjust only for the propensity score S.

If the counterfactual outcomes are independent of treatment conditional on C then the counterfactual outcomes are also independent of treatment conditional on S.

The advantage of adjusting for the propensity score S rather than C itself is that S is a single variable and stratification on a single variable is relatively straightforward

C may contain dozens of variables, making stratification and other methods of adjustment difficult

MSMs

Our two causal inference principles conflict!

Regression methods will not allow us to estimate the joint causal effects of A_1 and A_2 on Y in this case

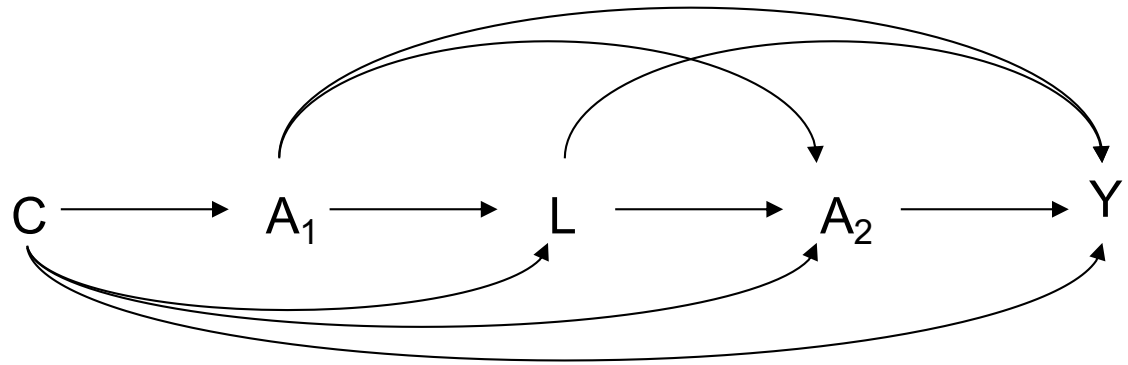

This problem will generally arise with time-varying treatment if there is any variable, such as L, that is both a confounder and an intermediate variable

Causal Inference with Longitudinal Data

Instead of regression (i.e. a model for the outcome conditional on the covariates) we will use a "marginal structural model" (again, a model for the counterfactual outcomes, Robins et al., 2000):

Let Y_{a1a2} be the counterfactual value of Y for an individual under an intervention to set A_1 to a_1 and A_2 to a_2

Regression: $E[Y|A_1=a_1, A_2=a_2,C=c] = \beta_0 + \beta_1 a_1 + \beta_2 a_2 + \beta_3' c$

MSM: $E[Y_{a1a2}] = \kappa + \gamma_1 a_1 + \gamma_2 a_2$

Because we do not observe $Y_{a^1a^2}$ for all possible values of a_1 and a_2 for all individuals we cannot fit the MSM directly

However we can fit the MSM using a weighting technique under certain assumptions. Specifically we need that:

Causal Inference with Longitudinal Data

(1) $Y_{a_1 a_2} \perp\!\!\!\perp A_1 \mid C$

(i.e. the effect of A_1 on the final outcome Y is unconfounded given C)

(2) $Y_{a_1 a_2} \perp\!\!\!\perp A_2 \mid \{C, A_1, L\}$

(i.e. the effect of A_2 on Y is unconfounded given baseline C, A_1 and the potential intermediate(s) denoted by L)

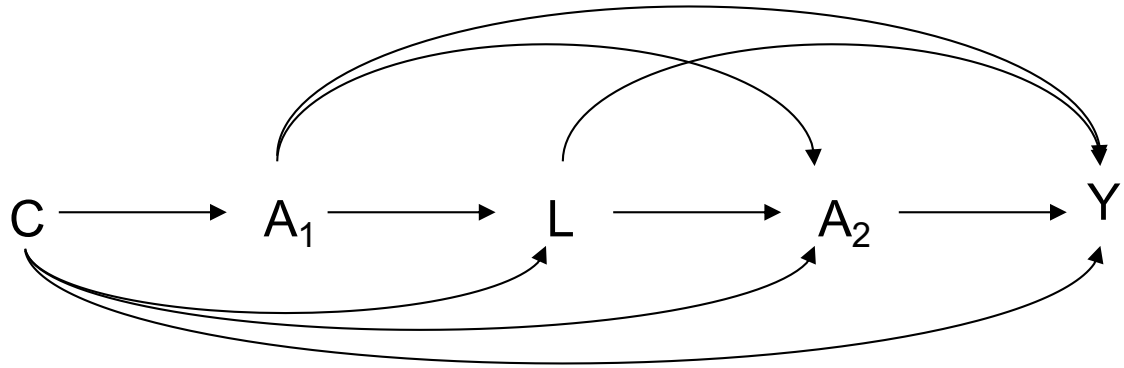

For the reasons given earlier, we generally wanted to include in C a measurement of the outcome (and possibly of the exposure A_0 as well); prior to the first measurement we are consider interventions on

Causal Inference with Longitudinal Data

MSM: $E[Y_{a1a2}] = \kappa + \gamma_1 a_1 + \gamma_2 a_2$

Robins (1999) showed that under these no-unmeasured-confounding assumptions we can obtain consistent estimators of κ , γ_1 and γ_2 (the parameters of the MSM) by fitting the regression model:

$$E[Y|A_1=a_1, A_2=a_2] = \kappa + \gamma_1 a_1 + \gamma_2 a_2$$

where each subject i is weighted by

$$\frac{P(A_1 = a_1^i)}{P(A_1 = a_1^i|C = c^i)} \times \frac{P(A_2 = a_2^i|A_1 = a_1^i)}{P(A_2 = a_2^i|A_1 = a_1^i, C = c^i, L = l^i)}$$

where a^i_1, a^i_2, c^i, l^i are the values for individual i of A_1, A_2, C and L respectively (C and L can be multivariate)

To fit a conditional MSM, conditioning on certain baseline covariates V: $E[Y_{a1a2}|V=v] = \kappa + \gamma_0 v + \gamma_1 a_1 + \gamma_2 a_2$ we follow the same procedure of fitting a weighted regression of $E[Y|A_1=a_1, A_2=a_2,V=v] = \kappa + \gamma_0 v + \gamma_1 a_1 + \gamma_2 a_2$ but where we condition on V in both the numerator and denominator of the weights

We can also include interaction terms between V and $A_1=a_1$ and/or $A_2=a_2$

Additional Time Points

The approach extends to more than two times of treatment
With three exposure periods to intervene on we can use the MSM:

$E[Y_{a1a2a3}] = \kappa + \gamma_1 a_1 + \gamma_2 a_2 + \gamma_3 a_3$

Confounding Assumption: At each period k, the baseline covariates C, and the history of the time-varying covariates, $L_1,\ldots,L_{k-1}$ and exposures up through time k-1, $A_1,\ldots A_{k-1}$ suffice to control for confounding of the effect of the exposure, at time k, A_k, on the outcome at each subsequent time i.e. $Y_{a1a2\ldots aT} \perp\!\!\!\perp A_k \mid (C, L_1,\ldots,L_{k-1}, A_1,\ldots A_{k-1})$

We add an additional inverse probability of treatment weight for each exposure/treatment period:

$$W = \prod_{k=1}^{t} \frac{p(A_k = a_k^i \mid A_1 = a_1^i, \ldots, A_{k-1} = a_{k-1}^i)}{p(A_k = a_k^i \mid A_1 = a_1^i, \ldots, A_{k-1} = a_{k-1}^i, C = c^i, L_1 = l_1^i, \ldots, L_{k-1} = l_{k-1}^i)}$$

Ignoring Time-Dependent Confounding

When can we look at the effects of trajectories of exposure and ignore time-varying confounders?

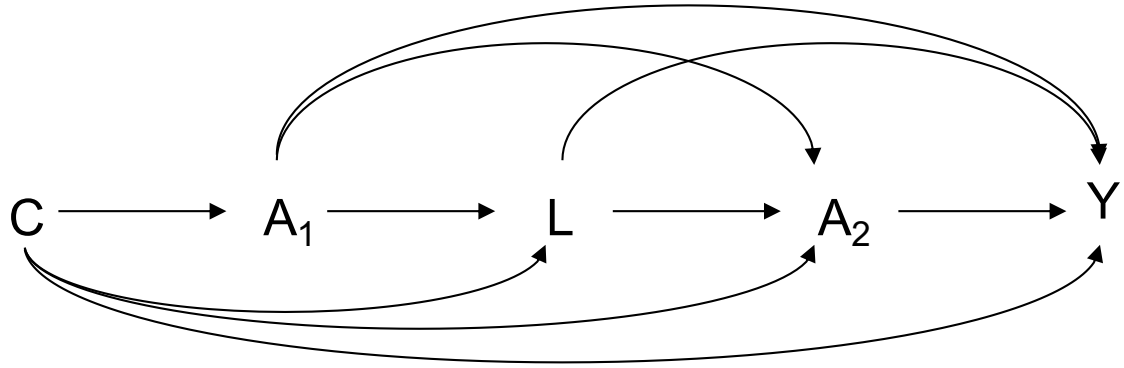

The problem arises when a variable L confounds the effect of subsequent exposure and is itself affected by prior exposure

The problems with traditional regression methods go away if:

- There is no effect of L on subsequent exposure, or
- There is no effect of L on the outcome, or
- There is no effect of prior exposure on L

We can examine evidence for these empirically; when one of the three effects is absent we do not need to use causal methods for time-dependent confounding

Sensitivity Analysis w/o Assumptions

Here we give an approach that in some sense makes “no assumptions” (Ding and VanderWeele, 2016)

Assume we have one or more unmeasured confounders U

We define the following parameters (always >= 1)

$$RR_{UY} = \max\left(\frac{\max_u P(Y=1|A=0,c,u)}{\min_u P(Y=1|A=0,c,u)}, \frac{\max_u P(Y=1|A=1,c,u)}{\min_u P(Y=1|A=1,c,u)}\right)$$

$$RR_{AU} = \max_u \frac{P(u|A=1,c)}{P(u|A=0,c)}$$

These are essentially the maximum effect of U on Y across exposure groups and the maximum risk ratio relating A and any category of U

If U is binary these are just ordinary risk ratios, but apply more generally

Note that the parameters condition on the measured covariates C

They capture the confounding after having already controlled for C

Sensitivity Analysis w/o Assumptions

Let RR_{UY} be the maximum risk ratio relating any two categories of U to Y conditional on covariates C and across exposure A

Let RR_{AU} be the maximum risk ratio relating A to any category of U

The largest factor by which such a U could reduce an observed risk ratio estimate is given by (Ding and VanderWeele, 2016, Epidemiology):

$$B = \frac{RR_{UY} * RR_{AU}}{(RR_{UY} + RR_{AU} - 1)}$$

We can divide the estimate and CI by B to obtain a "corrected" estimate and CI (i.e. the maximum a confounder could shift the estimate)

The result holds without making any assumptions about the structure of U

If the association is <u>protective</u> we multiply rather than divide by B

Note: Whenever the formula above is used B, RR_{UY}, RR_{AU} are ***always*** ≥ 1

Sensitivity Analysis w/o Assumptions

We can also ask how much confounding would explain away an estimate and its confidence interval

With an observed risk ratio of RR, we have that if RR_{UY} and RR_{AU} are greater than (VanderWeele and Ding, 2017):

E-value = RR + sqrt[RR*(RR-1)]

Then this could suffice, but weaker joint confounder associations could not

We can apply this in a routine manner to both the estimate and the confidence interval limit closest to the null

If RR<1, take inverses first and apply the formula

We might call this the E-value (*evidence* for causality)

RR = 10.73 (95% CI: 8.02, 14.36)

E-Value for Estimate: 20.9; E-value for CI: 15.5

Prevalence Specification

If we are willing to specify the prevalence (and make further assumptions) we can sometimes obtain smaller bias, and also get the exact bias

Result (Schlesselman,1978) : If there is no unmeasured confounding given (C,U) i.e. $Y_a \perp\!\!\!\perp A \mid (C,U)$ and U is a binary unmeasured confounder with the same risk ratio, γ, on Y for both exposed and unexposed so that $\gamma = \frac{E(Y|a,c,U=1)}{E(Y|a,c,U=0)}$

Then we have that: $B_{mult}(c) = \frac{1+(\gamma-1)P(U=1|a_1,c)}{1+(\gamma-1)P(U=1|a_0,c)}$

We can use the bias formula by specifying:
(i) γ = the effect of U and
(ii) the prevalence of U amongst the exposed and unexposed

Difference Scale and Continuous Outcomes

The bias for the causal effect conditional on C on the difference scale:

$B_{add}(c) = \{E(Y|a_1,c) - E(Y|a_0,c)\} - \{E(Y_{a_1}|c) - E(Y_{a_0}|c)\}$

Result. Suppose that we have a binary unmeasured confounder U with the effect of U on Y on the difference scale the same for the exposed and unexposed with $\gamma = E(Y|a,c,U=1) - E(Y|a,c,U=0)$ then the conditional bias is:

$$B_{add}(c) = \gamma\delta$$

where δ = E(U|A=1,c)-E(U|A=0,c) is the difference in expectations of U between the exposed and unexposed conditional on C=c (prevalence difference if U is binary)

Note these parameters are conditional on / control for measured C

Measurement Error

Suppose we have measured the prevalence of X* as p_{X^*}

Then the prevalence p_X of the true exposure X is given by:

$$p_X = (p_{X^*} - 1 + Sp) \ / \ (Se + Sp - 1)$$

$$Var(\hat{p}) \approx \frac{1}{\left(Se + Sp - 1\right)^2} \frac{p_{X^*}(1 - p_{X^*})}{n}$$

The variance assumes known Se and Sp

In general, sensitivity and specificity need to be obtained by a validation study in which we can measure the gold standard of the exposure

This could be considerably smaller than the primary sample measuring X*

In this case, the variance is more complex

Measurement Error

(1) Non-differential measurement error of a binary exposure will be biased towards the null (provided Se+Sp>1), but may be biased away for a categorical exposure; for a continuous exposure estimates will be biased towards the null under independent additive error

(2) Non-differential measurement error of a binary outcome will be biased towards the null; additive non-differential measurement of a continuous outcome will be unbiased

(3) When biases towards the null are problematic (e.g. side effects) we need sensitivity analysis or validation studies

(4) Measurement error for a confounder will not always give results between the crude and the true association adjusted for the true confounder; if monotonicity is violated it may not be between these two

(5) With differential misclassification we can get biases away from the null and need to be careful; we can sometimes address this in sensitivity analysis or validation data

Additive vs. Multiplicative Interactions

For some time in the epidemiologic literature, there was debate as to which scale one should assess interactions on (Blot and Day, 1979; Saracci, 1980; Rothman et al., 1980)

The general historical consensus was:

(1) Often the additive scale is of greatest public health importance
    It allows one to discern whether the effect would be different in different subgroups (Rothman et al., 1980)

(2) The additive scale also seemed to correspond to the more biological notion of synergism as conceived of by Rothman (1976) *[later in lecture]*

(3) However, sometimes the multiplicative scale (or neither scale) may be the one that more naturally corresponds to the biological mechanisms (Siemiatycki and Thomas, 1981); though in these cases the additive scale is still important for assessing public health impact

Both additive and multiplicative interaction measures should be reported

There is <u>no good reason</u> not to report additive as well

Attributing Effects to Interactions

We can also do this on a ratio scale using excess relative risks:

$(RR_{11} - 1) = (RR_{10} - 1) + (RR_{01} - 1) + RERI$

And we could again calculate what portion of the of the joint effect is attributable to each of these effects:

$$\frac{RR_{10} - 1}{RR_{11} - 1}$$

$$\frac{RR_{01} - 1}{RR_{11} - 1}$$

$$\frac{RERI}{RR_{11} - 1}$$

These three proportions will sum to 100%

Note the third component $RERI/(RR_{11} - 1)$ differs from Rothman's attributable proportion measure $RERI/(RR_{11})$

Rothman's captures the proportion of *disease* in the doubly exposed due to interaction; the measure here the proportion of the *effect* of both exposures due to interaction

Interaction and Effect Heterogeneity

There are two quantities which may be of interest when we consider subgroups defined by some variable Q:

First, we may be interested in how the treatment effects vary by strata of Q:

$E[Y_1 - Y_0 \mid Q = q]$ ("effect heterogeneity" if this varies with q)

This would be of interest in **choosing groups to target** for treatment

Second, we may be interested in whether the treatment effect would vary if we were to intervene on Q

$E[Y_{1q} - Y_{0q}]$ ("causal interaction" if this varies with q)

where Y_{aq} denotes the counterfactual outcome for an individual if, possibly contrary to fact treatment, we set to treatment to a and Q to q

We might be interested in this second quantity if we were thought we might **intervene** on Q to make treatment A more effective or if we were interested in the mechanisms by which A and Q interacted

Mechanistic Interaction

Result: If the effects of G and E on D are unconfounded one can test for a sufficient cause interaction by testing:

$$p_{11} - p_{10} - p_{01} > 0$$

This condition can be expressed as RERI > 1

Result: If the effects of G and E on D are unconfounded and if both G and E have positive monotonic effects on the outcome then one can test for a sufficient cause interaction by testing:

$$p_{11} - p_{10} - p_{01} + p_{00} > 0$$

i.e. positive additive interaction

This condition can be expressed as RERI > 0

Definitions

Robins and Greenland (1992) and Pearl (2001) proposed the following counterfactual definitions for direct and indirect effects:

Controlled direct effect: The controlled direct effect comparing treatment level A=1 to A=0 intervening to fix M=m

$$CDE(m) = Y_{1m} - Y_{0m}$$

Natural direct effect: The natural direct effect comparing treatment level A=1 to A=0 intervening to fix $M=M_0$

$$NDE = Y_{1M_0} - Y_{0M_0}$$

Natural indirect effect: The natural indirect effect comparing the effects of $M=M_1$ versus $M=M_0$ intervening to fix A=1

$$NIE = Y_{1M_1} - Y_{1M_0}$$

Properties of Direct and Indirect Effects

A total effect decomposes into a direct and indirect effect:

$$Y_1 - Y_0 = Y_{1M_1} - Y_{0M_0}$$

$$= (Y_{1M_1} - Y_{1M_0}) + (Y_{1M_0} - Y_{0M_0})$$

$$= \text{NIE} + \text{NDE}$$

The <u>definitions</u> of natural direct and indirect effect do not presuppose no interactions between the effects of the exposure and the mediator on the outcome

The <u>effect decomposition</u> of a total effect into a natural direct and indirect effect also does not presuppose no interaction between the effects of the exposure and the mediator on the outcome

Natural direct and indirect effects are useful for effect decomposition; in general, controlled direct effects are not

Identification of Direct and Indirect Effects

To estimate average natural direct and indirect effects we need:

(1) There are no unmeasured exposure-outcome confounders given C

(2) There are no unmeasured mediator-outcome confounders given (C,A)

(3) There are no unmeasured exposure-mediator confounders given C

(4) There is no mediator-outcome confounder affected by exposure
(i.e. no arrow from A to C_2)

For controlled direct effects, only assumptions (1) and (2) are needed

Note (1) and (3) are guaranteed when treatment is randomized

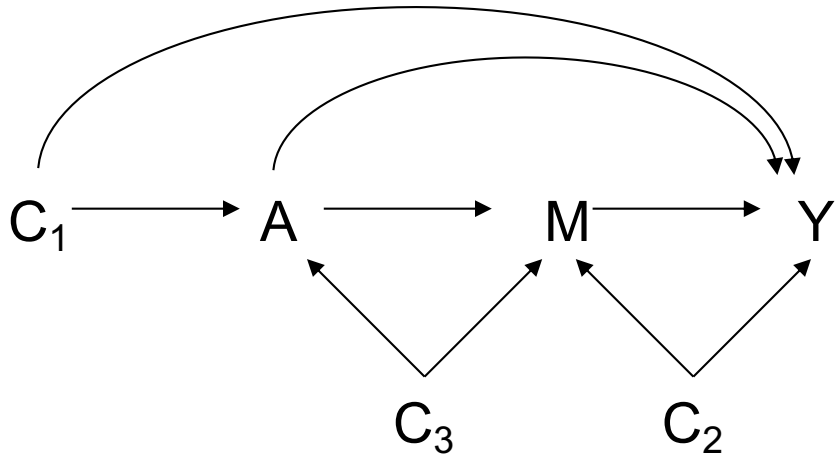

Identification of Direct and Indirect Effects

Under assumptions (1) and (2) the controlled direct effect conditional on the covariates is given by:

$$E[\ CDE(m) \mid c\ ] = E[Y|A=1,m,c] - E[Y|A=0,m,c]$$

Under (1)-(4) the conditional natural direct and indirect effects are:

$$E[\ NDE \mid c\ ] = \Sigma_m \{E[Y|A=1,m,c] - E[Y|A=0,m,c]\}\ P(M=m|A=0,c)$$

$$E[\ NIE \mid c\ ] = \Sigma_m\ E[Y|A=1,m,c]\ \{P(M=m|A=1,c) - P(M=m|A=0,c)\}$$

These are the effects within strata of the covariates

We could take averages over each stratum weighted by the probability P(C=c) to get population averages of the effects

Regression for Causal Mediation Analysis

We use regressions that accommodate exposure-mediator interaction:

$E[Y|A=a,M=m,C=c] = \theta_0 + \theta_1 a + \theta_2 m + \theta_3 am + \theta_4' c$

$E[M|A=a,C=c] = \beta_0 + \beta_1 a + \beta_2' c$

Under assumptions (1)-(4), and provided our models are correctly specified, we can combine the estimates from the two models to get the following formulas for direct and indirect effects, comparing exposure levels a and a* (VanderWeele and Vansteelandt, 2009):

$CDE(m) = (\theta_1+\theta_3 m)(a-a^*)$

$NDE = (\theta_1+\theta_3(\beta_0+\beta_1 a^*+\beta_2' E[C]))(a-a^*)$

$NIE = (\theta_2\beta_1+\theta_3\beta_1 a)(a-a^*)$

If the conditional NDE were of interest then we would have:

$E[Y_{aM_{a^*}} - Y_{a^*M_{a^*}} | C=c] = (\theta_1+\theta_3(\beta_0+\beta_1 a^*+\beta_2' c))(a-a^*)$