The Front Door Criterion

And why it is pretty cool

Demetrios Papakostas

Grad Stats Club

March 30, 2023


Back Door
Front Door
IG:@IJFXL

Directed acyclic graphs for causal inference

The use of directed acylcic graphs (DAGs) for causal inference owes a lot to Judea Pearl.

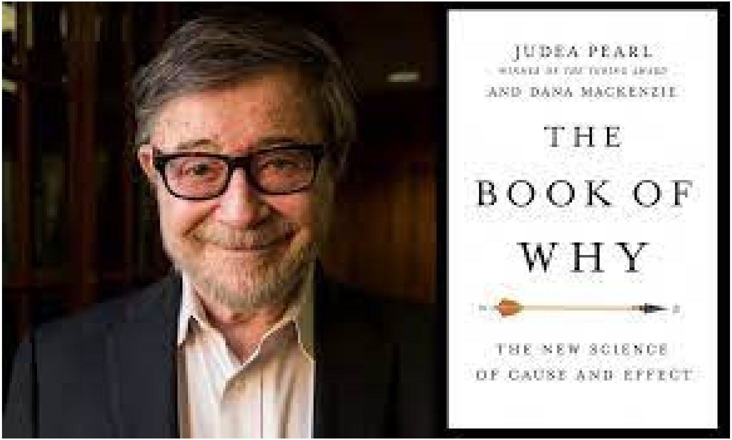

Pretty nifty approach...as we will see

Terminology

- A **collider** is a variable who has arrows pointing into it from multiple variables.
- A **node** is a variable in a causal graph
- A **line** is an edge of the graph

Terminology: Part II

- A **descendant** of a node is a child, or grandchild, etc. If we have arrows that point to directly, or downstream, from a variable X_i to another X_j, then we say X_j is a descendant of X_i
- A **parent** of a node is the variable whose arrow comes into the node of interest.

The Back-door criterion

First, assume we have a graph with a treatment, D, outcome, Y, and a collection of other variables (say $X_1, X_2, \ldots, X_n$), the set of which we call S. The back-door criterion algorithm can be expressed as:

1. Identify all (undirected) paths between D and Y.
2. Consider all variables along each of the paths and make sure at least one is "blocked". A variable Q is blocked if
    - Q is not a collider and is in S
    - or Q is a collider and neither Q nor any of its descendants are in S

The Back-door criterion

First, assume we have a graph with a treatment, D, outcome, Y, and a collection of other variables (say $X_1, X_2, \ldots, X_n$), the set of which we call S. The back-door criterion algorithm can be expressed as:

1. Identify all (undirected) paths between D and Y.
2. Consider all variables along each of the paths and make sure at least one is "blocked". A variable Q is blocked if
   - Q is not a collider and is in S
   - or Q is a collider and neither Q nor any of its descendants are in S

- An undirected path is a sequence of edges (ignoring directionalities), but with constraint that the path has an edge pointing into D.
- We need to find variables in S that are not colliders and ensure colliders (and their descendants) are not in the conditioning set S of variables we are interested in.

Further explanation

- If every path contains at least one blocked variable, we say D and Y are "d-separated" by S, and thus we have found a set of variables that satisfies the "backdoor criterion".
- If we do condition on a collider, we must also condition on a parent of the collider that is also on the path between D and Y.
- Conditioning on S means any relationship we find between D and Y we can interpret as causal. Failure to condition on S will contaminate estimates causal estimates.

A Do-calculus primer

- If the backdoor criterion is satisfied, then the causal effect of D on Y is identified, given by the relation:

$$\Pr(Y \mid \text{do}(D)) = \sum_{s} \Pr(Y \mid D, S) \Pr(D) \tag{1}$$

- Where the "do"-operator" is introduced as an operator allows us to artificially set a treatment at a certain value [Pea00]. Under the do($\cdot$) framework, we can return a valid estimate of the ATE as

$$\text{E}(Y = 1 \mid \text{do}(D = 1)) - E(Y = 1 \mid \text{do}(D = 0)) \tag{2}$$

- do($D = d$) refers to an exogeneous intervention, i.e. setting D at 1 or 0 (in the binary treatment) case and seeing what *would* have happened if every individual were treated (or not).

An analysis of some DAGs

Does this satisfy the back-door criterion?

Would conditioning on X satisfy the back-door criterion?

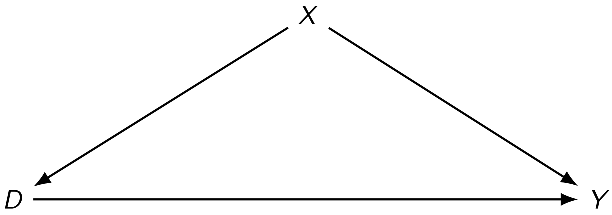

Does this satisfy the back-door criterion?

Would conditioning on X satisfy the back-door criterion?

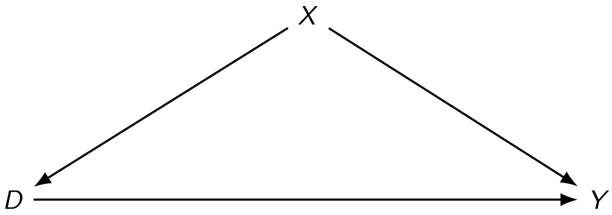

YES

A second example

Is conditioning on W okay for causal identification of D on Y?

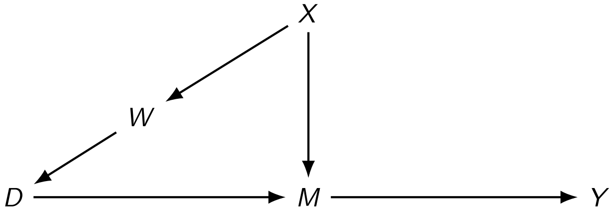

A second example

Is conditioning on W okay for causal identification of D on Y?

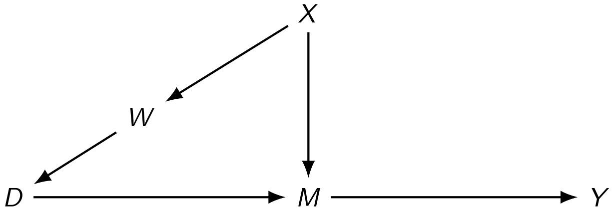

Yes, this will gives us an unbiased estimate of the **total effect** of D on Y. Holding M constant would give us the **direct effect**.

Does this satisfy the back-door criterion?

Would conditioning on X satisfy the back-door criterion?

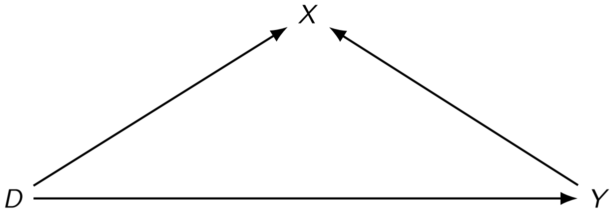

Does this satisfy the back-door criterion?

Would conditioning on X satisfy the back-door criterion?

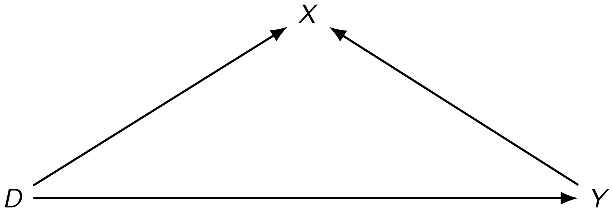

No, X is a **collider**. An example collider is conditioning on lung inflammation may open up an association between COVID and lung cancer that is spurious.

Another scarf example

- Why is it that all attractive people who you date are **jerks**? (see right)
- Mean people aren't necessarily attractive nor the other direction.
- BUT, if you condition on who you date (being nice and being attractive both have an arrow pointing to whether or not you are a desirable person to date) we notice a spurious association between being a jerk and being attractive! You wouldn't date someone you don't find attractive AND is a jerk!

An instrument!

What if we condition on Z?

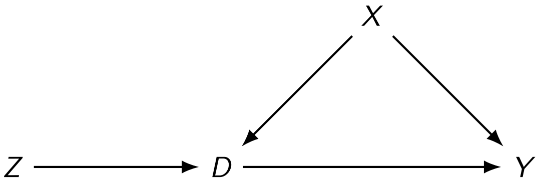

An instrument!

What if we condition on Z?

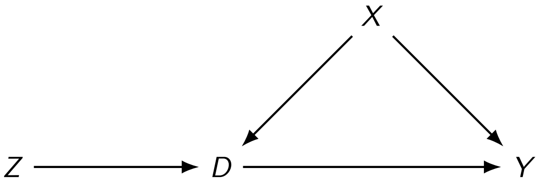

Not good! We won't *bias* our estimate, but the variance will increase.

What about here?

Would conditioning on X_2 be okay?

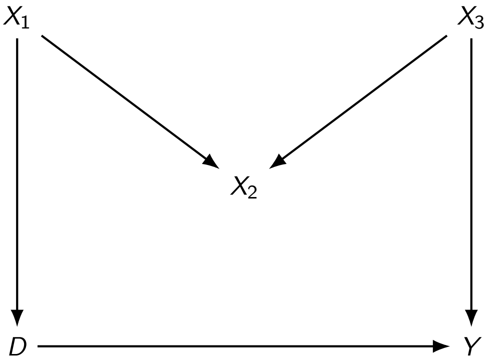

What about here?

Would conditioning on X_2 be okay?

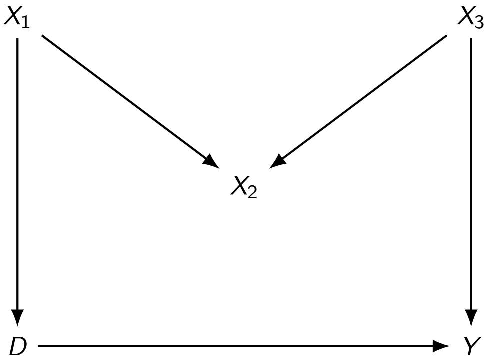

NO! Conditioning on X_2 (a collider) opens a backdoor path from D to Y since both X_1 and X_3 point into X_2, and thus an association between them is spurred.

(Hint: You can check in R if you are not sure!)

But...

- Drawing the graphs (assuming we actually can) and applying the back-door algorithm is all fine and dandy, but what if we do **not observe** variables that will satisfy the back-door criterion?

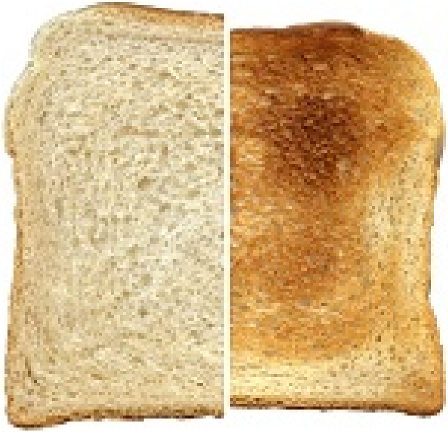

Let’s try the front door instead

The front door criterion

- The front door criterion, originally formulated in [Pea95], is a useful workaround if we have unobserved confounding.

- Denote D as the treatment, Y is the outcome, U is unobserved confounding, and M is a mediator (this is not to be confused with what is typically done in "mediation analysis").

Front door criterion: conditions

The following conditions were established by Pearl to proceed via a "Front-door" analysis.

- M intercepts all directed paths from D to Y, i.e. $D \perp\!\!\!\perp Y \mid M$ (conditional independence)
- There is no backdoor path between D and M.
- Every back-door path between M and Y is blocked by D.

Intuitively, the idea of the front door analysis is to find a mediating variable, M, that is not affected by the unobserved confounding. Typically conditioning on a mediator is undesirable, as it blocks the effect of D on Y

Visually

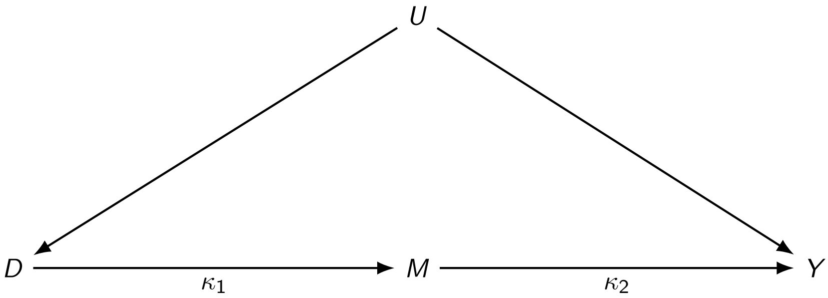

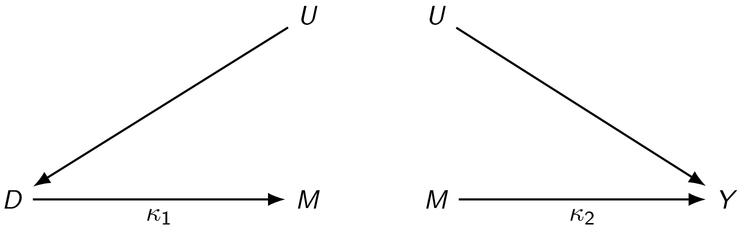

Estimand of Interest

- The $\text{do}(\cdot)$ notation refers to the causal intervention of setting a given variable at a certain level.

- What we want to estimate is:

$$\Pr(Y \mid \text{do}(D)) = \sum_{M} \Pr(M \mid \text{do}(D)) \times \Pr(Y \mid M, \text{do}(D)).$$

- However, observing $\text{do}(D)$ (or $\text{do}(M)$) is unlikely outside of an experiment, but as [Pea95] showed it is actually possible, given the conditions above are met, to re-formulate this expression using only observed data.

Derivation

- If we use the conditions, we can reformulate the desired estimand.
- Because there is no backdoor path from the treatment D to the mediator M, $\Pr(M \mid do(D)) = \Pr(M \mid D)$
- Because D blocks all back-door paths between M and Y, then $\Pr(Y \mid do(M)) = \sum_D \Pr(Y \mid D, M) \times \Pr(D)$.
- Finally, because by condition that M intercepts all directed paths from D to Y, then $\Pr(Y \mid M, do(D)) = \Pr(Y \mid do(M))$

The front door criterion equation

- Substituting the expressions gives us the following equation:

$$Pr(Y \mid D) = \sum_{M} \Pr(M \mid D) \sum_{D'} \Pr\big(Y \mid D', M\big) \times \Pr\big(D'\big)$$

- This is the **front door criterion** equation.

- Additionally, $\Pr(D_i \mid M_i) > 0$ for all units i, i.e. a positivity assumption. The mediator can't solely be determined by the treatment.

What did we just do?

- Essentially, we get the effect of D on M which automatically satisfied the back door criterion. (Left side plot on 20)
- To get a valid causal estimate of M on Y, we must condition on D, as that blocks the only backdoor path between M and Y. (Right side plot on **??**)

How to estimate

- The following regression framework is used (see [BBW19] for a nice overview):

$$M_i = \kappa_1 D_i + \varepsilon_i$$
$$Y_i = \kappa_2 M_i + \lambda D_i + \tilde{\varepsilon}_i$$

- and the ATE is given by

$$\text{ATE} = E(Y \mid \text{do}(D)) = \hat{\kappa}_1 \times \hat{\kappa}_2.$$

Is this ever used?

- Not really... the conditions necessary are quite daunting!
- Let's think as a graduate statistics club if we can brainstorm some together

A demonstration in R

- We will demonstrate this cool method and play around in R
- This would be cool as a Quarto interactive blog, but alas that didn't happen

Graph of our R-dgp

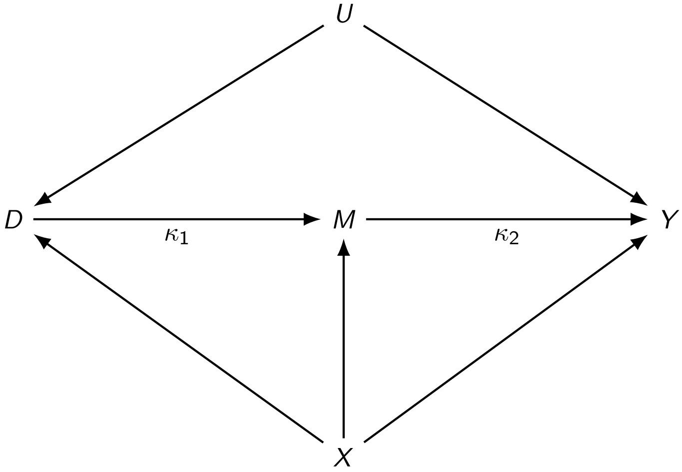

Bibliography I

[BBW19] Marc F Bellemare, Jeffrey R Bloem, and Noah Wexler, *The paper of how: Estimating treatment effects using the front-door criterion*, Tech. report, Working paper, 2019.

[Pea95] Judea Pearl, *Causal diagrams for empirical research*, Biometrika **82** (1995), no. 4, 669–688.

[Pea00] J Pearl, *Causality: Models, reasoning, and inference*, Cambridge University Press, Cambridge, England, 2000.

Thank you