Machine learning as a tool for causal inference

Bill Simpson-Young

(Based on the work of Dan Steinberg, Finn Lattimore and others)

GRADIENT
INSTITUTE

Overview

- About Gradient Institute (https://gradientinstitute.org/)
  - Independent, non-profit research institute
  - Building ethics, accountability and transparency into AI systems
  - Technical research, consulting, teaching and AI system assessment

Gradient Institute focus areas

GRADIENT INSTITUTE

1. Machine learning research:

- Algorithmic decision-making and bias
- Causal inference using machine learning
- Algorithmic fairness
- Generative AI (and containing some of the harms)

2. Responsible AI training:

- Technical training (data scientists, software engineers)
- Leader training (people responsible for AI systems, board/exec training)

3. AI system/process/practices design and assessment:

- Responsible AI system design
- AI system/process/practices/risk assessment (technical and governance)
- Responsible AI software tool development

4. Responsible AI policy:

- Input into federal, state and global AI policy (incl OECD)

Talk overview

- Using machine learning as a tool for causal inference
- Case studies:
    - student well-being
    - (school leadership)
    - (early childhood care)

Policy aims to change the world

GRADIENT INSTITUTE

An aim of policy is (usually) to improve outcomes for people by **changing something** about our society/the world.

We have to take care when using data and models to inform policy:

- the data we have is of the world as it is now
- we want to know how a policy would lead to a **different world**

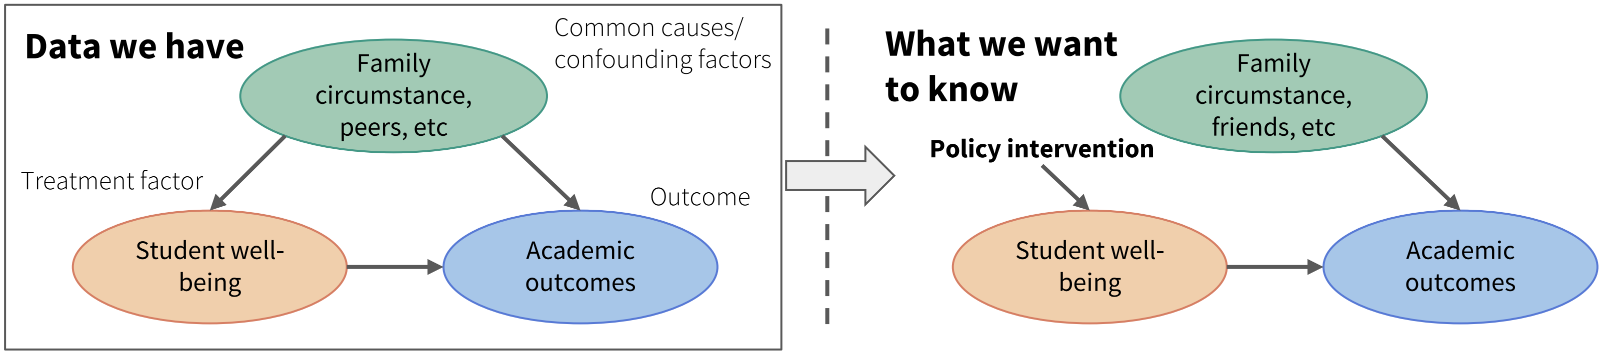

Designing policy requires causal inference

Statistics/prediction (correlation) is no longer enough, we need **causal inference**.

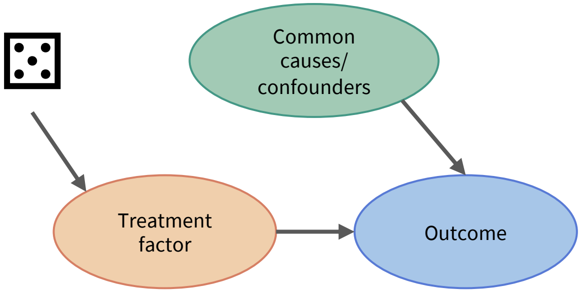

Experimentation (randomisation)**:** changes the world

- Removes influence of confounding factors
- But can be expensive/ infeasible/ unethical

Modelling: “simulates” changing the world

- But have to **capture all confounding factors** and
- **Model the world accurately**

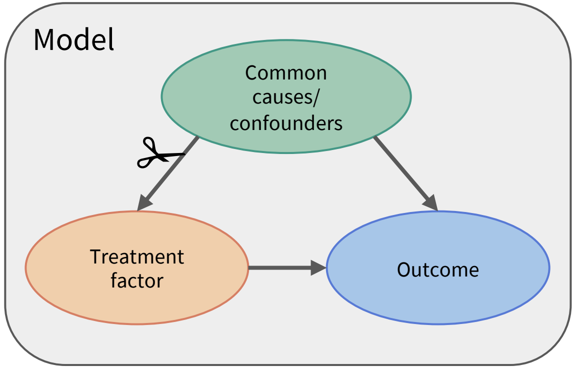

Linked datasets can be helpful for modelling

Many, varied datasets are **more likely to contain factors** we need for the modelling approach to succeed.

But with these larger datasets come **new challenges**:

- They contain **complex relationships** - making it hard to construct accurate models
- There are many **related factors** - can make traditional statistical models unhappy

This is where **machine learning** (ML) can help.

Why and how: using ML for causal inference

Where ML adds value

- Econometric/statistical models have typically been used for causal inference
- We are now dealing with **larger** and more **complex** datasets
- **ML designed for discovering relationships in large and complex data**
- But - need to keep formal statistical theory and assumptions

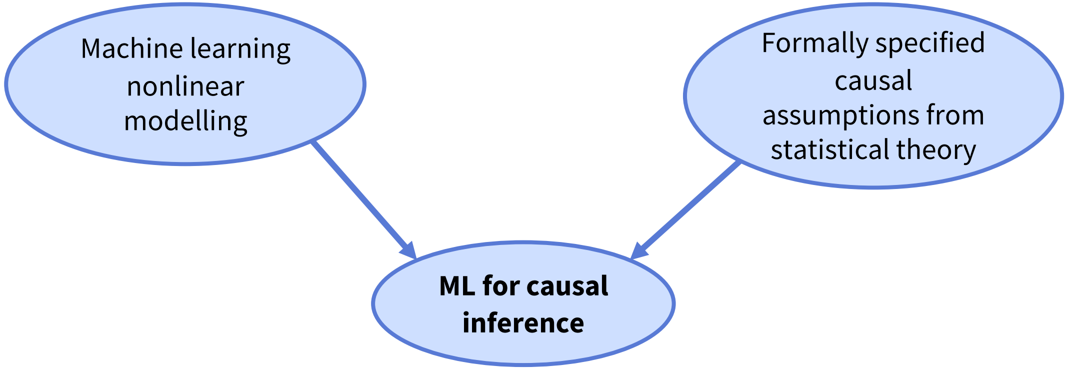

ML model: high no. parameters relative to no. samples
Stat model: low no. parameters relative to no. samples

Complex: ML can model complex relationships

We only care about the **treatment → outcome** relationship

But, have to **model every relationship accurately** to estimate causal effect

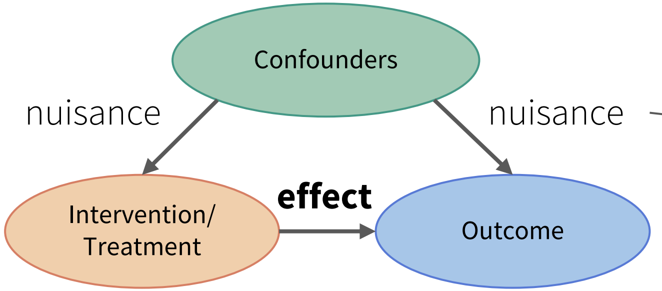

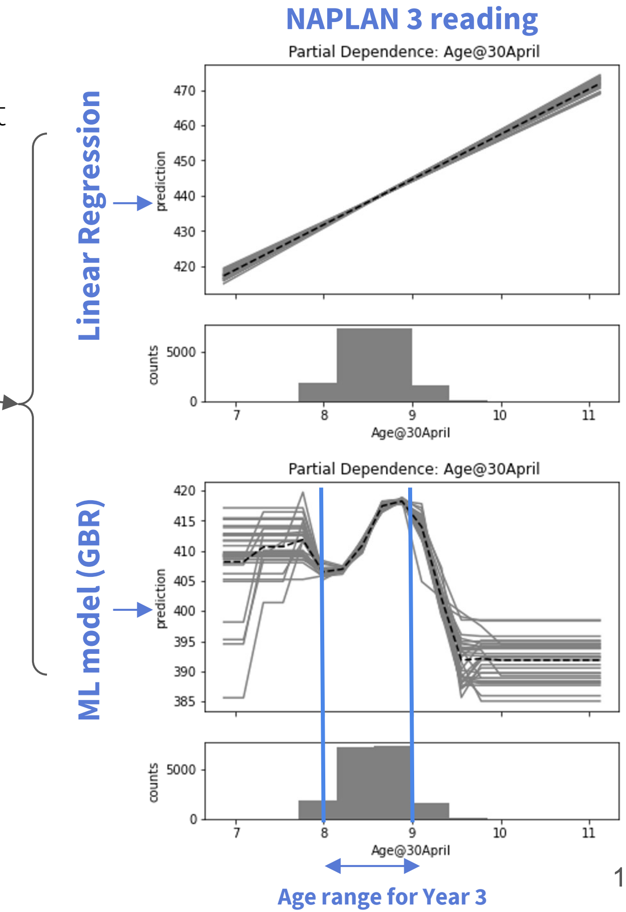

- Let ML take care of nuisance relationships (it’s more likely to succeed)
- Carefully model the effect relationship we care about

The typical analysis process

1. Formulate assumptions about causal relationships [1], see if you can **identify** the causal effect

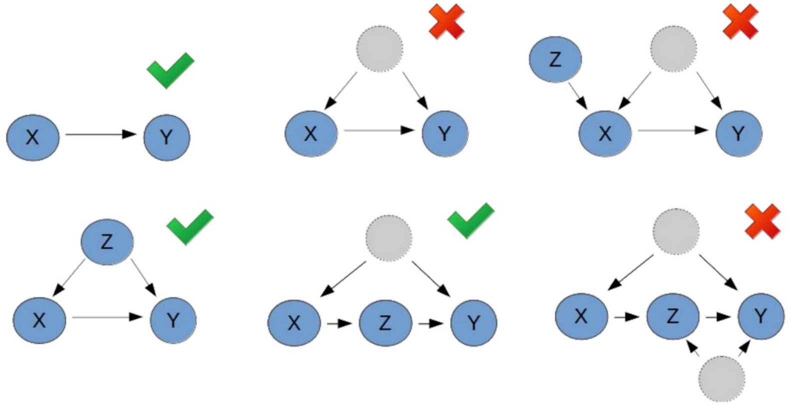

1. Construct a **model** of the system
    - Direct Regression
    - Doubly Robust
    - Double ML

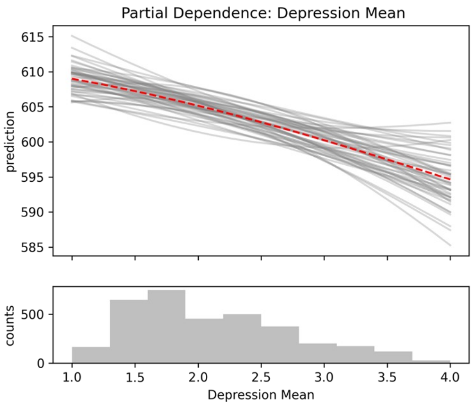

1. Run “**interventions**” on your model - report (visualise) causal effect [2, 3]

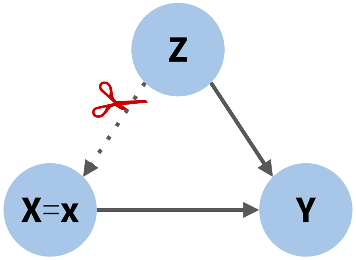

[1] Judea Pearl. "Causal inference in statistics: An overview." Statistics Surveys, 3 96-146 2009. https://doi.org/10.1214/09-SS057
[2] Cook, T.R., Gupton, G., Modig, Z., Palmer, N.M., 2021. Explaining Machine Learning by Bootstrapping Partial Dependence Functions and Shapley Values. RWP. https://doi.org/10.18651/RWP2021-12
[3] Zhao, Q., Hastie, T., 2021. Causal Interpretations of Black-Box Models. Journal of Business & Economic Statistics 39, 272–281. https://doi.org/10.1080/07350015.2019.1624293

Case study

Data available to us - ACT public schools

Survey data[1]:

- Student climate
  - Wellbeing
  - School factors
- Staff climate*
- Parent climate*

Administrative data:

- Public school census
- Teacher length of service and employment type*
- Leadership changes*
- Casual teacher utilisation*
- ICSEA*
- Early childhood care

Exam data:

- NAPLAN 3-9
- BASE/PIPS start and end
- HSC ATAR

[1] Reynolds, K.J., Subasic, E., Bromhead, D., & Lee, E. (2017). The school as a group system: School climate, school identity and school outcomes. In K. Mavor, M. J. Platow & B. Bizumic (Eds). The self, social identity and education. London, UK: Psychology Press.

* These datasets have been aggregated to the school-year level to link with the individual student data

Case study - student well-being

- Does **self reported well-being affect academic outcomes** (year 7 to 9 NAPLAN)?
- Well-being composed of three survey constructs:
  - Depression
  - Anxiety
  - Positive affect
- Controlling for a large array of individual and school factors
- Relationships were quite complex, but can be simplified
- ~3400 students

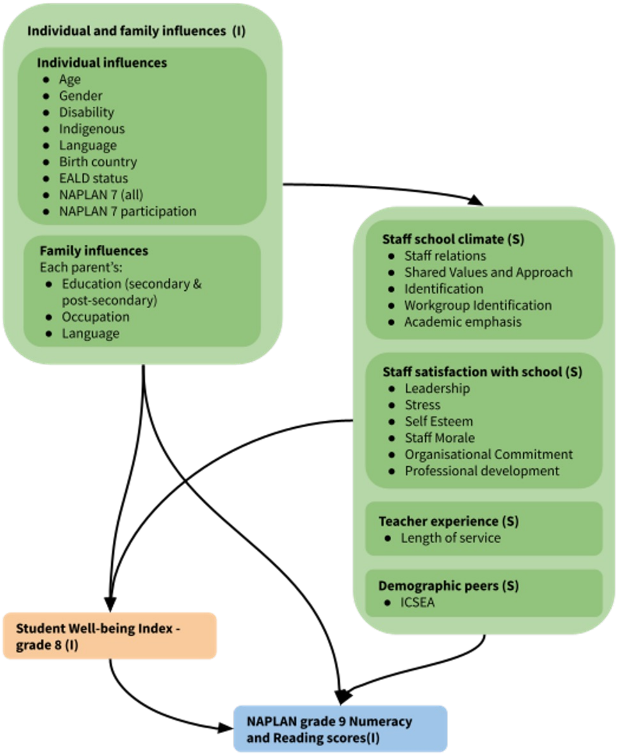

Simplified Graph

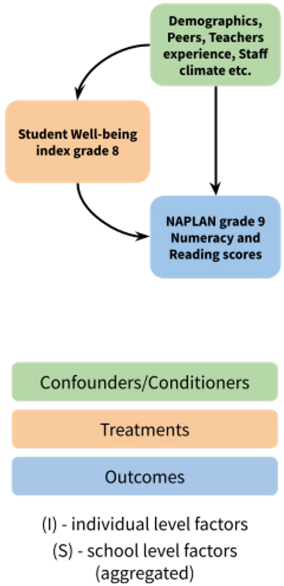

Cárdenas, D., Lattimore, F., Steinberg, D. et al. **Youth well-being predicts later academic success**. Scientific Reports 12, 2134 (2022). https://doi.org/10.1038/s41598-022-05780-0

Case study - student well-being

GRADIENT INSTITUTE

scientific reports

Explore content ⌄ About the journal ⌄ Publish with us ⌄

nature › scientific reports › articles › article

Article | Open Access | Published: 08 February 2022

Youth well-being predicts later academic success

Diana Cárdenas, Finnian Lattimore, Daniel Steinberg & Katherine J. Reynolds

Scientific Reports **12**, Article number: 2134 (2022) | Cite this article

Abstract

Young people worldwide face new challenges as climate change and complex family structures disrupt societies. These challenges impact on youth's subjective well-being, with evidence of decline across many countries. While the burden of negative well-being on productivity is widely examined amongst adults, its cost among youth remains understudied. The current research comprehensively investigates the relationship between youth subjective well-being and standardized academic test scores. We use highly controlled machine learning models on a moderately-sized high-school student sample (*N* ~ 3400), with a composite subjective well-being index (composed of depression, anxiety and positive affect), to show that students with

Case study

- **Self-reported depression** gave us the most significant relationship
- 1 std. dev. improvement in depression → 3 additional NAPLAN points (~7% improvement from year 7)
- High agreement amongst models

Nonlinear treatment models

NAPLAN Numeracy Yr 9

Linear Regression (Ridge)

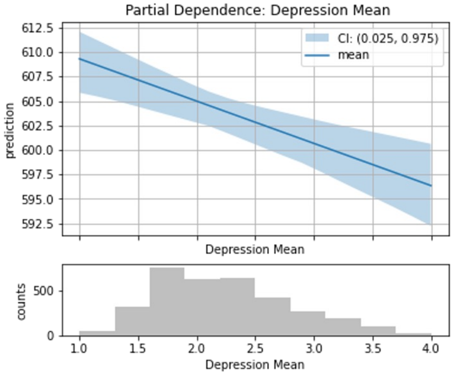

Kernel Regression

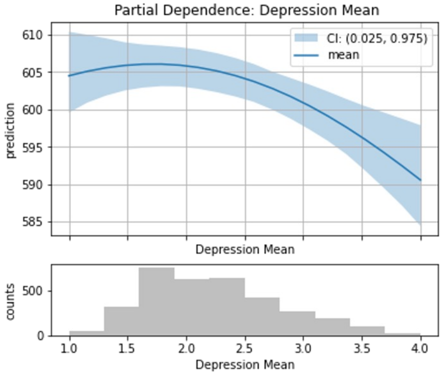

Gradient Boosted Trees

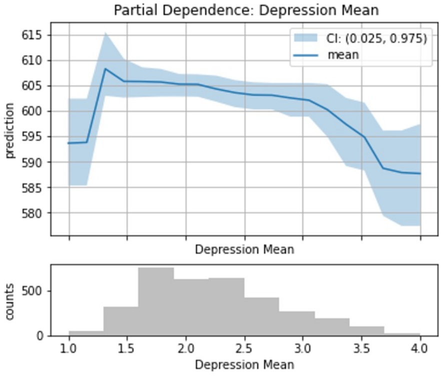

Student self reported depression

Open source software

GRADIENT INSTITUTE

Python:

- EconML (Microsoft Research) - https://github.com/microsoft/EconML
- CausalML (Uber) - https://github.com/uber/causalml
- *Causal Inspection (Gradient Institute)* - https://github.com/gradientinstitute/causal-inspection
- *TwoStageRidge (Gradient Institute)* - https://github.com/gradientinstitute/twostageridge
- Accelerated bayesian causal forests - https://github.com/socket778/XBCF

R:

- Generalised random forests - https://cran.r-project.org/web/packages/grf/
- bartMachine - https://cran.r-project.org/web/packages/bartMachine/
- Bayesian causal forests - https://github.com/jaredsmurray/bcf
- Others?

Literature

Books:

- *Elements of Causal Inference*, Peters et. al. 2017 (ML heavy!)
- *Counterfactuals and causal inference*, Morgan et. al. 2015 (Estimation)
- *Causality*, Pearl 2009 (identification, do calculus)

Papers:

- DML, DRLearners look at: https://econml.azurewebsites.net/spec/references.html
- Interpretation/Intervention
  - Cook et. al. 2021. Explaining Machine Learning by Bootstrapping Partial Dependence Functions and Shapley Values
  - Zhao et.al. 2021. Causal Interpretations of Black-Box Models
- Others
  - Hill, 2011. Bayesian Nonparametric Modeling for Causal Inference.
  - Hahn et. al. 2020. Bayesian Regression Tree Models for Causal Inference: Regularization, Confounding, and Heterogeneous Effects.

In summary...

- Econometric/statistical models have typically been used for causal inference
- We are now dealing with **larger** and more **complex** linked datasets
- **ML designed for discovering relationships in large and complex data**
- But - need to keep formal statistical theory and assumptions

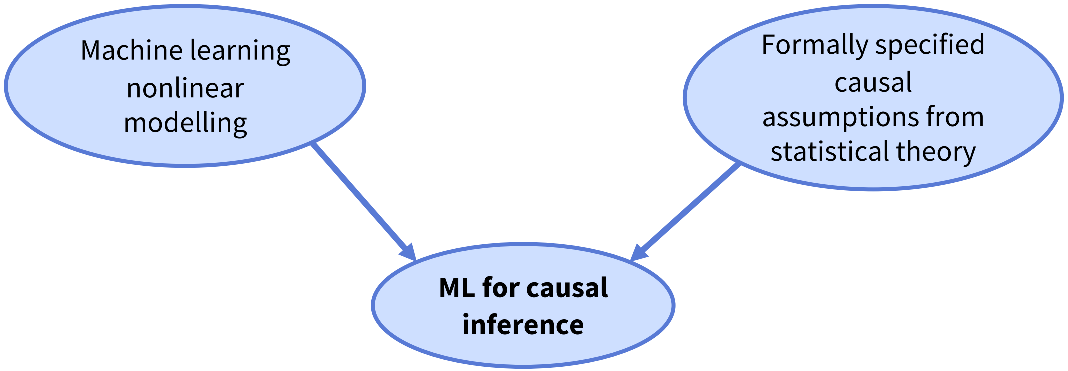

References:

- *Elements of Causal Inference*, Peters et. al. 2017 (ML theory)
- *Counterfactuals and causal inference*, Morgan et. al. 2015 (Estimation)
- *Causality*, Pearl 2009 (Identification, do calculus)

Thank you.

https://gradientinstitute.org/

Bill Simpson-Young
bill@gradientinstitute.org

GRADIENT INSTITUTE

Large: machine learning can help with control

We have to capture **all confounding factors** to compute a causal effect (**there are exceptions, e.g. IV)

These can be **many** and **highly related** (collinear), statistical methods can have numerical issues

We have a choice:

- **a.** **Select confounding factors** to leave out - potentially making results overconfident (or re-confounding the results)
- b. Use **machine learning** (regularisation) - potentially biasing results

New specialised ML estimators *also* reduce bias

Bias - Variance Tradeoff

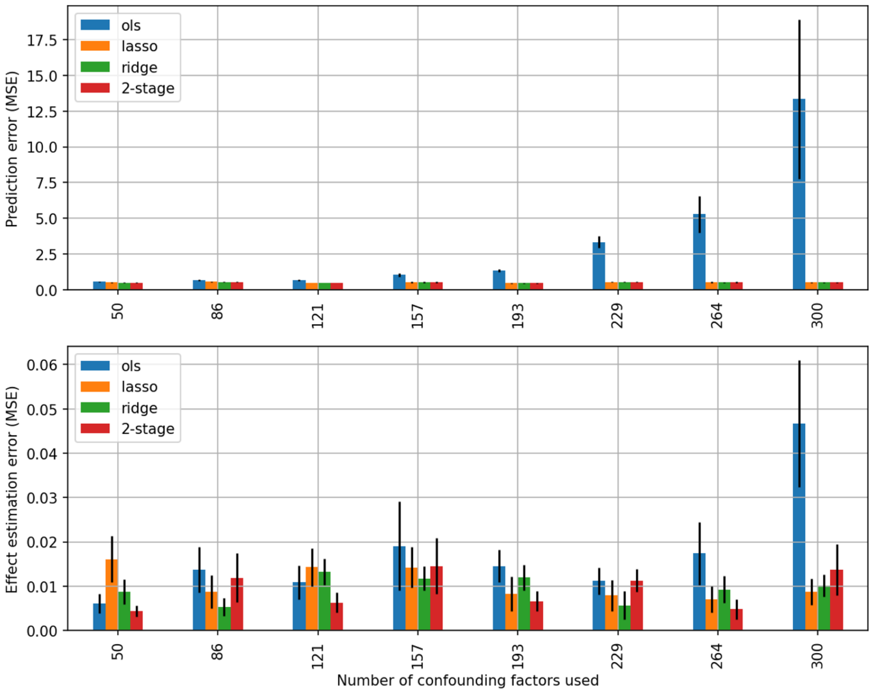

https://gradientinstitute.org/machine-learning-as-a-tool-for-evidence-based-policy/

Three important graphs

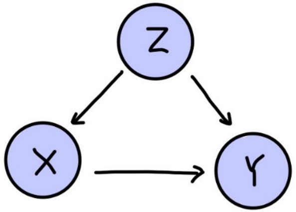

Adjusting for confounders - Ignorability conditional on Z

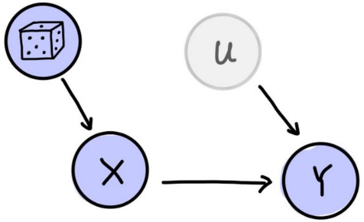

Randomised experiments

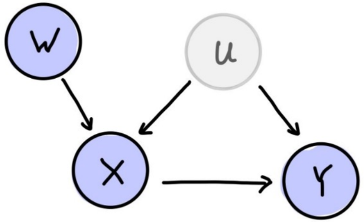

Instrumental variables

X - intervention/policy
Z - observed confounder
u - Unobserved confounder
Y - outcomes
W - IV’s

Should we adjust for this variable?

GRADIENT INSTITUTE

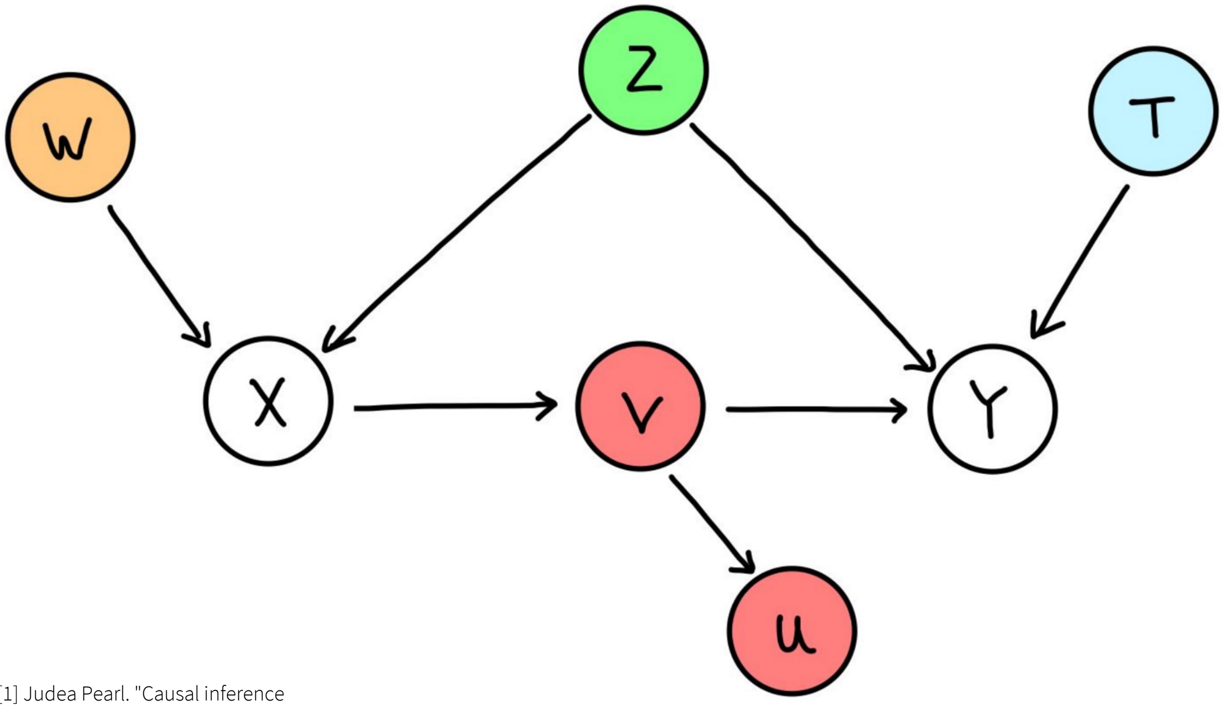

Yes

Probably

Probably not

No

W - we should use IV estimation

V - we should use front door estimation

[1] Judea Pearl. "Causal inference in statistics: An overview." Statistics Surveys, 3 96-146 2009. https://doi.org/10.1214/09-SS057