

**The  
Alan Turing  
Institute**

**#TMCFTURING**

Intro to Causal Inference +  
Counterfactual Prediction

[www.turing.ac.uk](http://www.turing.ac.uk)

 turinginst

 theturinginst

**PETER WG TENNANT PHD**

**ROMIN PAJOUHESHNIA PHD**

#TMCFTURING

# Intro to Causal Inference



**Peter WG Tennant**  
University of Leeds, UK



PWGTennant



adhd.scientist

The  
Alan Turing  
Institute

**LEEDS** *Institute for  
Data Analytics*

UNIVERSITY OF LEEDS



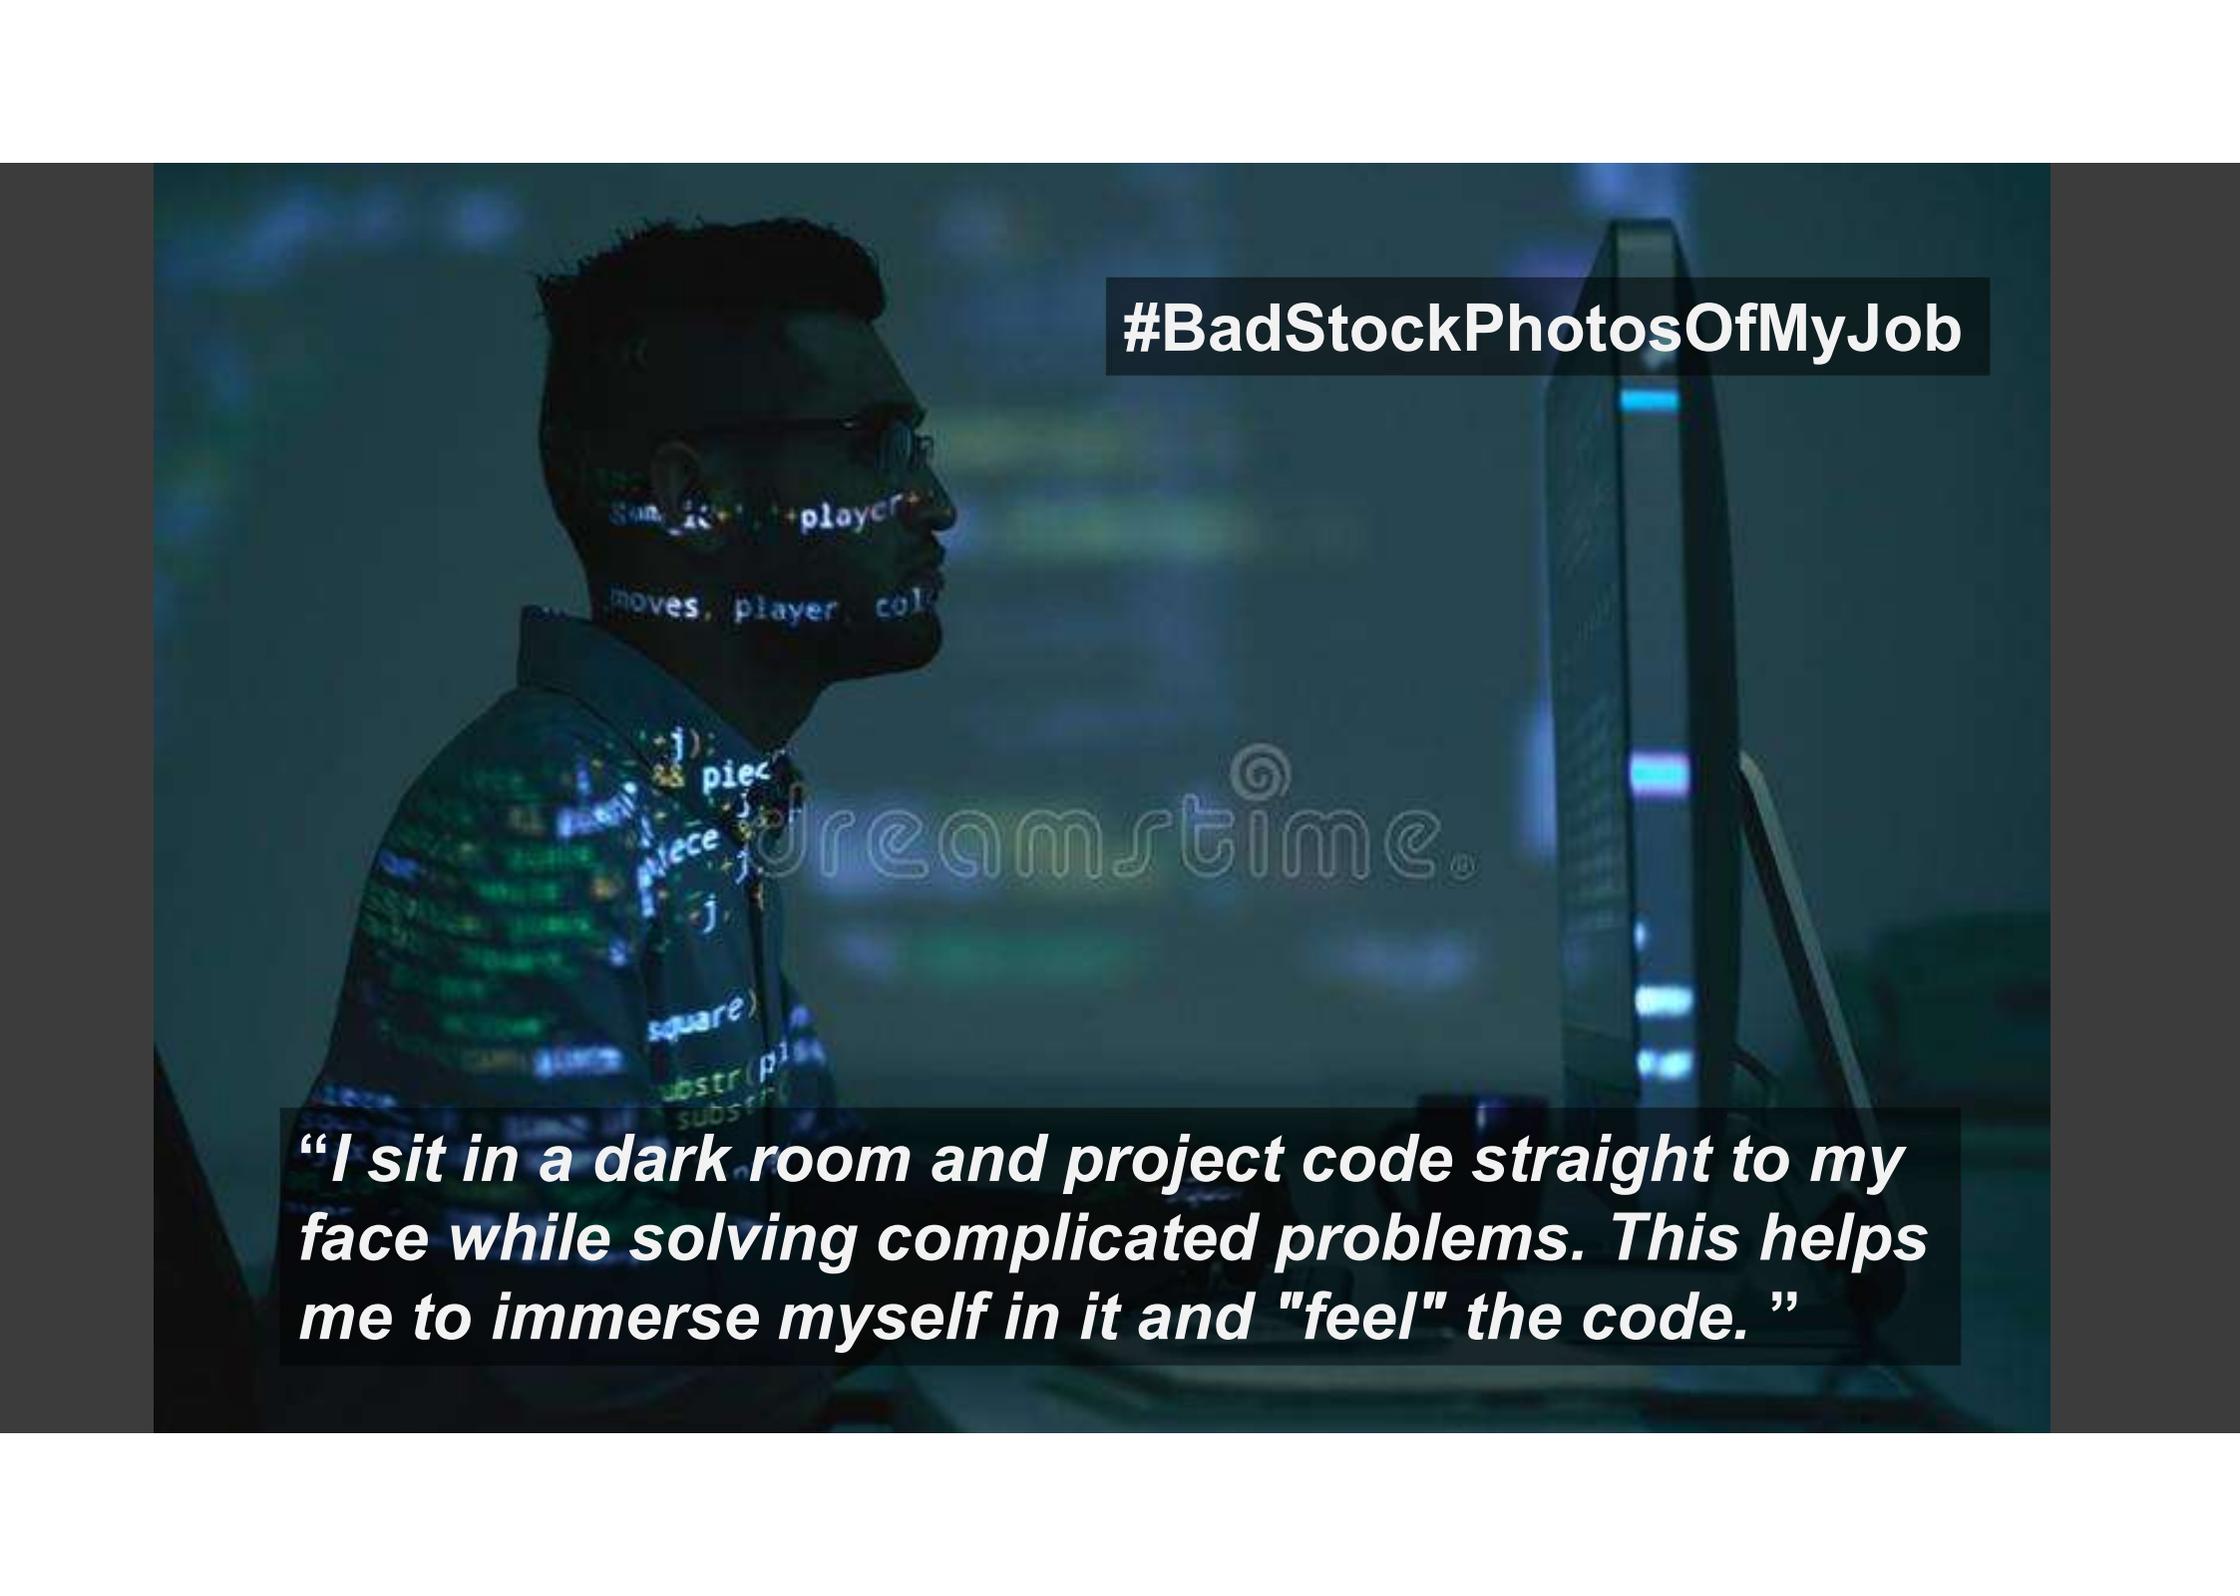
# WHAT IS DATA SCIENCE?

#TMCFTURING



WHAT IS  
DATA  
SCIENCE?



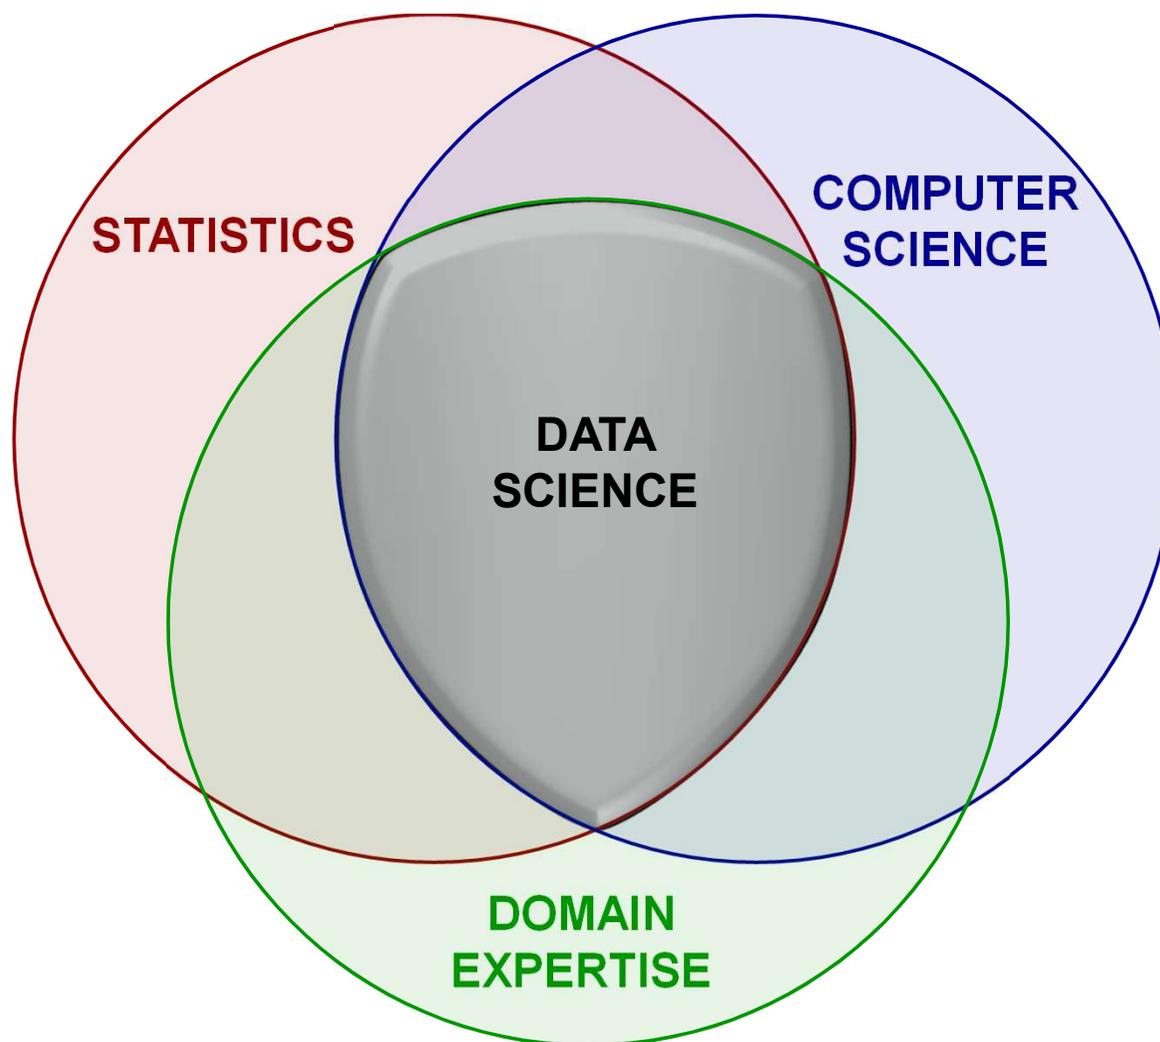
A person is shown in profile, sitting at a computer workstation in a dark room. The scene is dimly lit, with the primary light source being the computer monitor and the code being projected onto the person's face and shirt. The code is displayed in various colors, including green, blue, and yellow. The person's face is illuminated by the code, which appears to be floating in the air around them. The background is dark and out of focus, with some blurred lights from the computer setup. The overall atmosphere is one of intense focus and immersion in the work.

**#BadStockPhotosOfMyJob**

***“I sit in a dark room and project code straight to my face while solving complicated problems. This helps me to immerse myself in it and “feel” the code.”***

# WHAT IS DATA SCIENCE?

#TMCFTURING



- Science is about building knowledge and understanding
- Data science is about '*gaining insights*' and '*extracting meaning*' from data
- Most data science activity can be divided into **three scientific tasks**, each with different methods **and philosophies**



Hernán et al 2019. A Second Chance to Get Causal Inference Right: A Classification of Data Science Tasks, *CHANCE*, 32:1, 42-49, DOI: 10.1080/09332480.2019.1579578



## Description (& visualisation)

- Focussed on **summarising, describing, &/or visualising** features of interest
- Data driven - involves simple calculations & unsupervised learning

## Questions

- What happened?
- Who was affected?
- What was occurrence of **Y** in people with **X**?

## Example

- Occurrence and spread of COVID-19



## Prediction (AKA classification and regression)

- Focussed on **pattern recognition** and **forecasting**
- Data driven – involves statistical modelling and supervised learning

## Questions

- What **will** happen?
- Who **will** be affected?
- Are people with **X** are more likely to have **Y**?

## Examples

- Screening for COVID-19 with symptoms or CT scan
- Predicting prognosis or severity of infection



## Causal inference (AKA counterfactual prediction)

- Focussed on understanding
- **NOT data driven** – involves fusion of external knowledge with statistical modelling and supervised learning

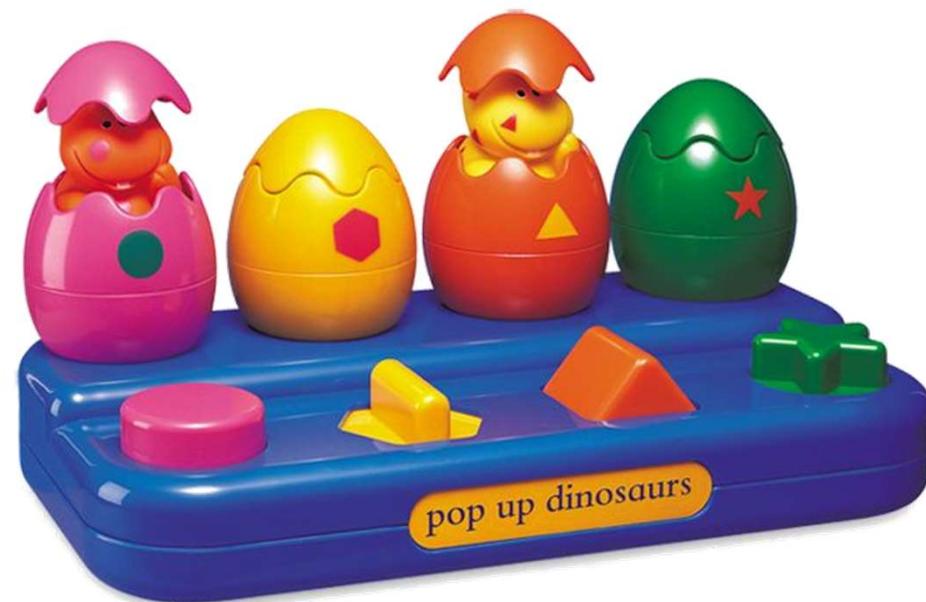
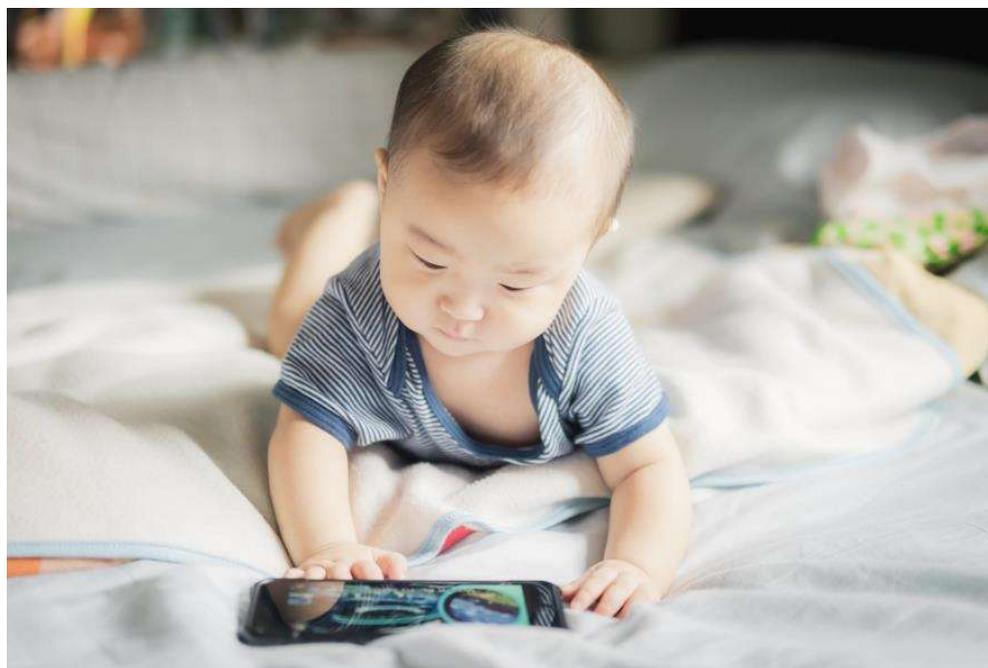
## Questions

- What will happen **if**...?
- **Why** were they affected?
- If we **changed X**, how would it **change Y**?

## Examples

- Effect of opening/closing schools on infection spread
- Risks / benefits of ventilation

- We learn about cause-and-effect from the moment we're born
- From observing how things interact, and what happens when we *do* things

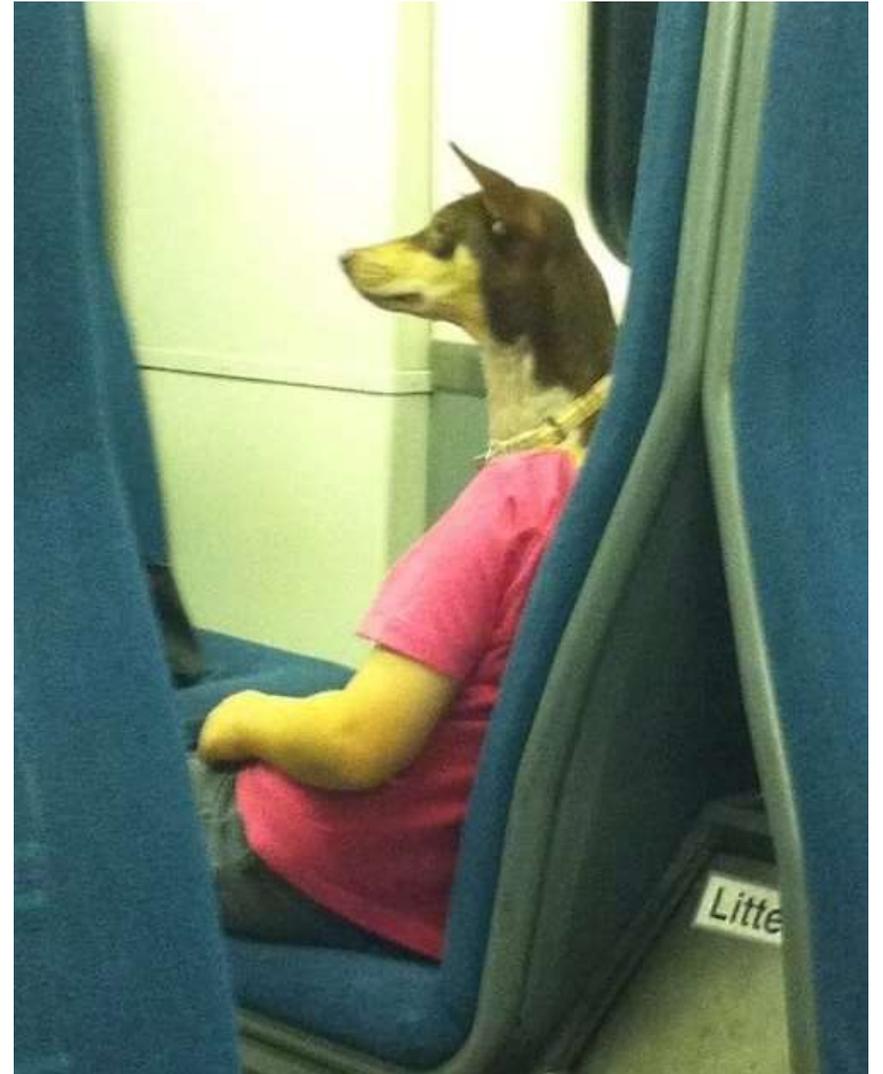
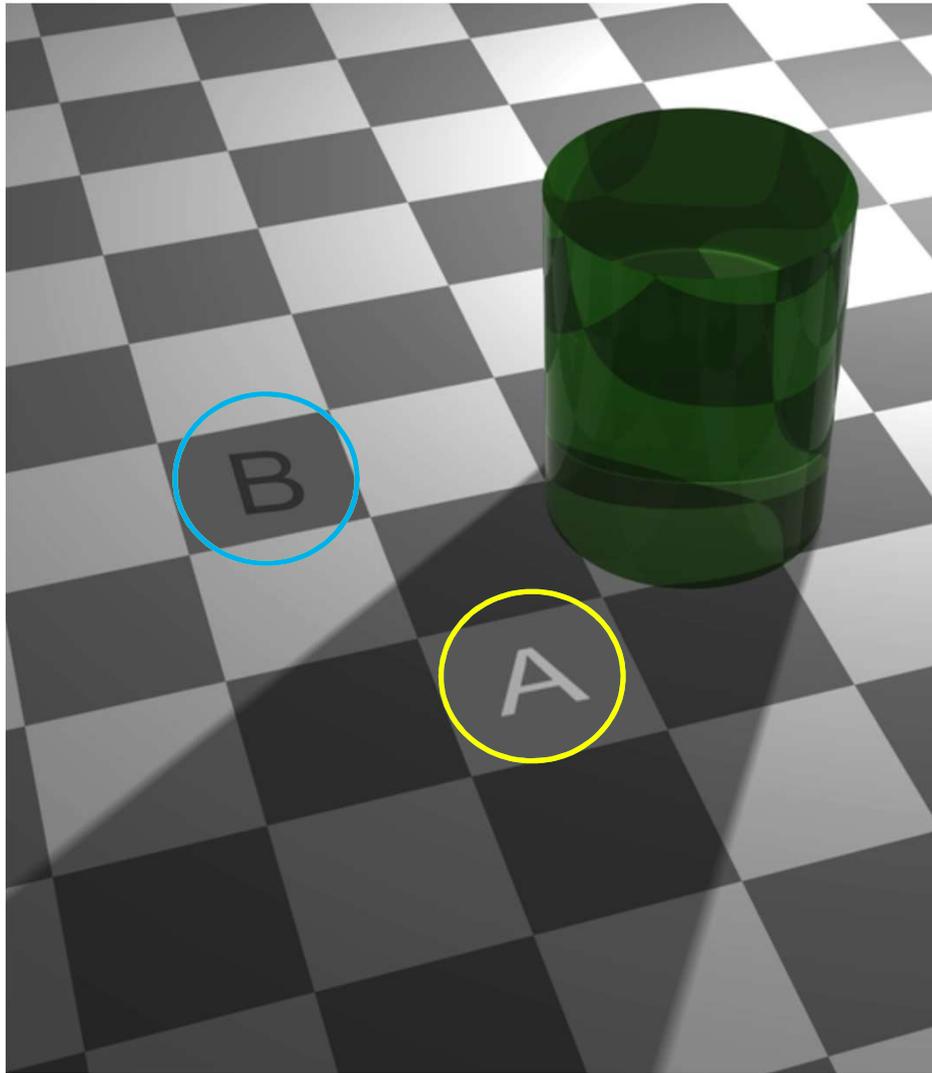


- This extends beyond we can see & do, to whatever we can imagine!
- Our ability to ask '*what might have been?*' and '*how things could be different?*' has led to us changing our world beyond our ancestors imaginations



# EASILY FOOLED

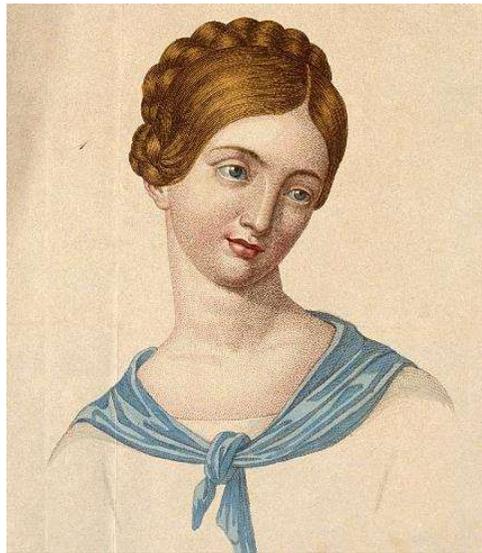
#TMCFTURING



<https://www.deviantart.com/butisit/art/Checker-shadow-illusion-263331875>

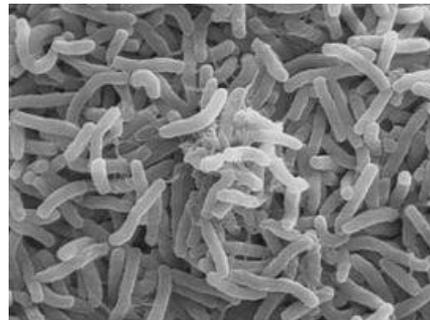
<https://www.newsinq.com/24-photos-need-really-look-understand-v1/>

This is especially true once we're faced with questions and events that don't fit our childhood model of cause-and-effect...



HEALTHY

+



VIBRIO CHOLERAЕ



UNHEALTHY

✓ **Deterministic**

This is especially true once we're faced with questions and events that don't fit our childhood model of cause-and-effect...



HEALTHY

+



SMOKING



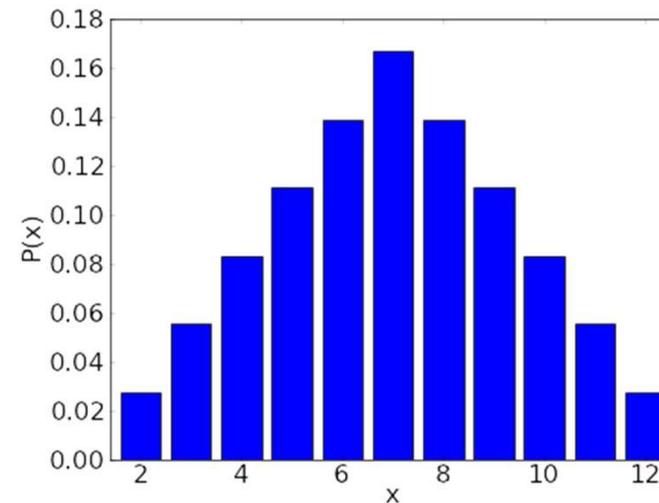
HEALTHY

**× Probabilistic**

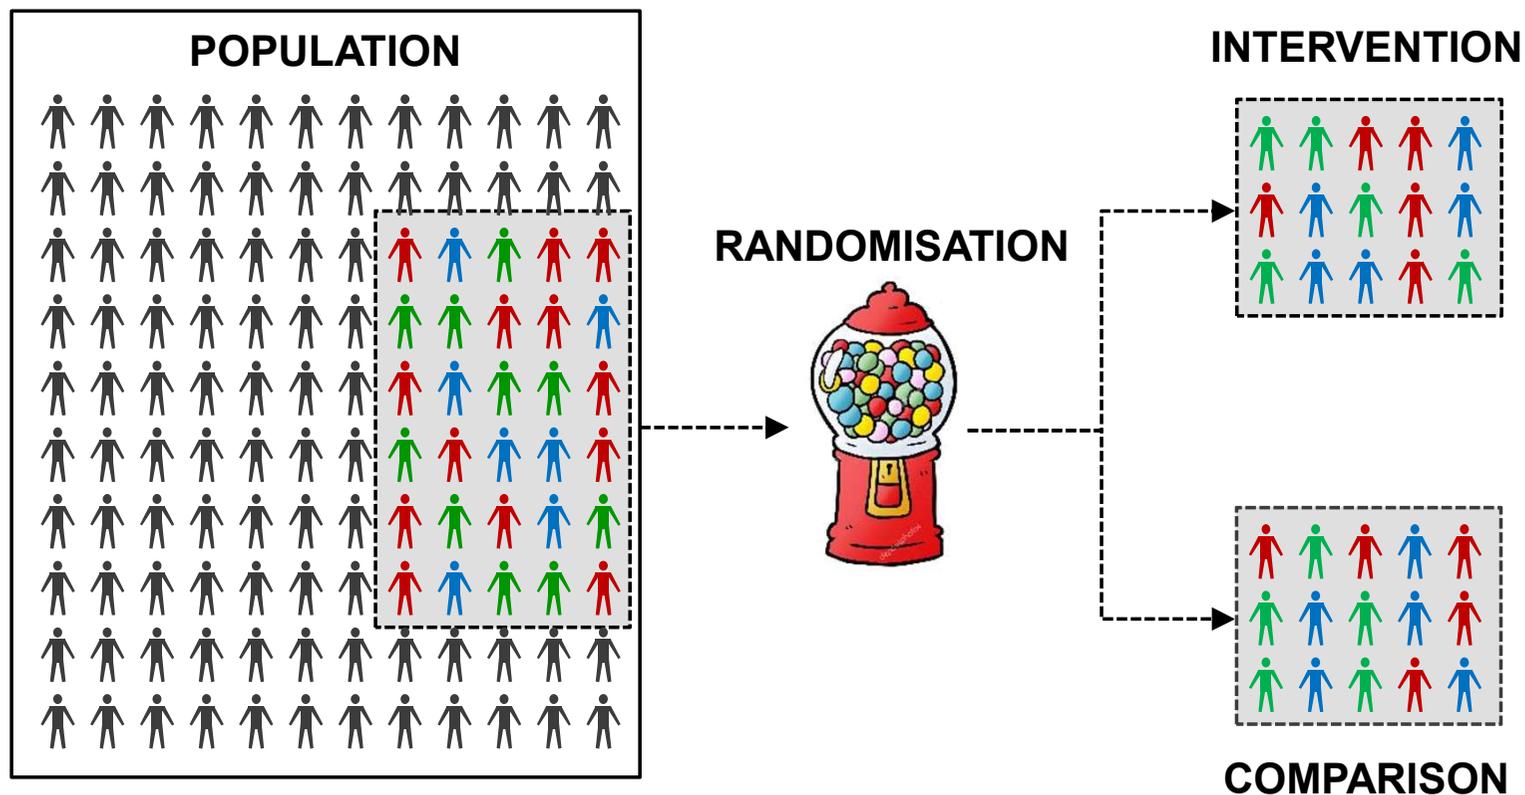
- Statistics & probability have resolved this problem
- What's **unknowable** for an individual... can be **predictable** for a group!



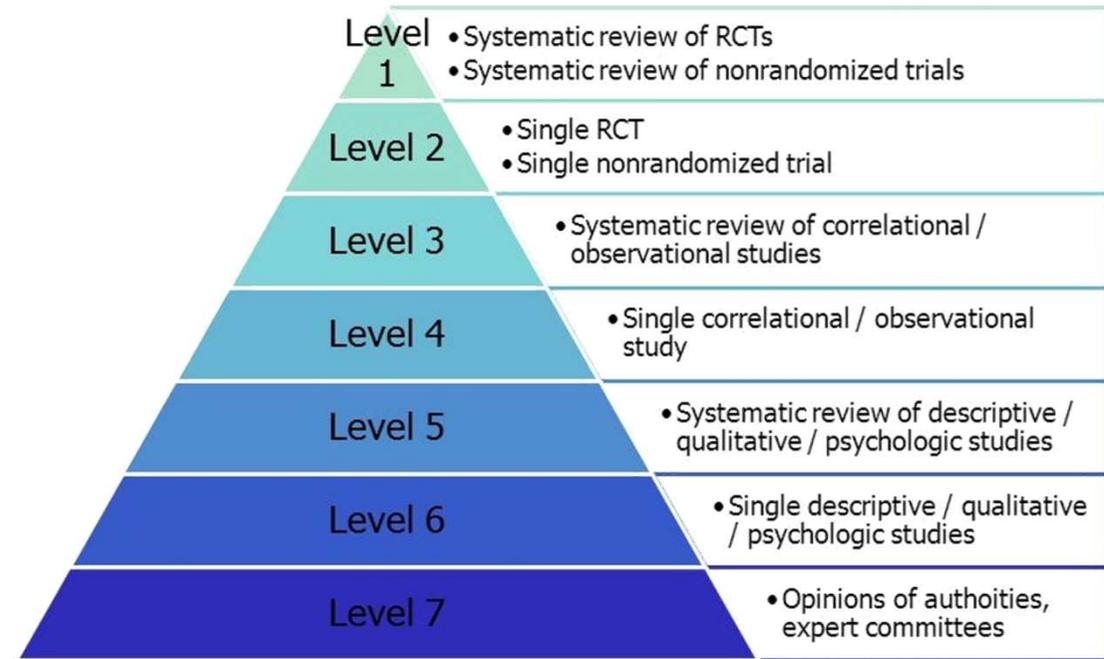
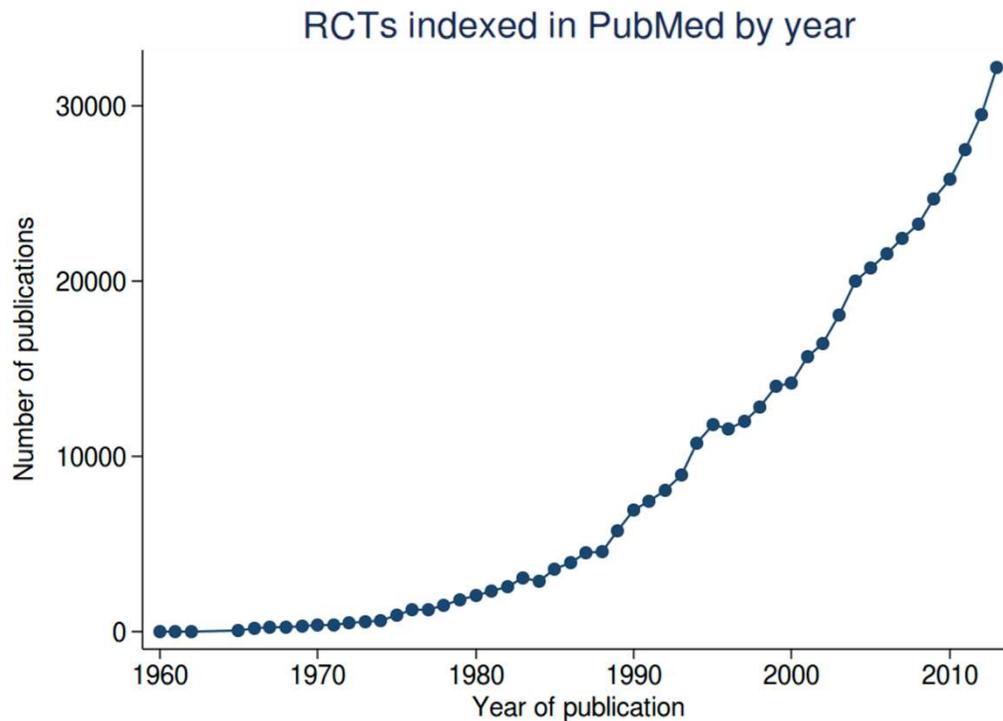
?



With **randomisation**, this provides a potent way to estimate causal effects that we'd have little hope with experience and intuition alone



In health and medical research, **randomised controlled trials** embraced with impassioned fervour, and canonised beyond all other forms of evidence



- Non-experimental studies have a poor record for causal inference
- **12 RCTs** studied **52 claims** from non-experimental data
  - **0** replicated
  - **5** found opposite!

<i>ID no.</i>	<i>Pos.</i>	<i>Neg.</i>	<i>No. of claims</i>	<i>Treatment(s)</i>	<i>Reference</i>
1	0	1	3	Vit E, beta-carotene	<i>NEJM</i> 1994; 330: 1029–1035
2	0	3	4	Hormone Replacement Ther.	<i>JAMA</i> 2003; 289: 2651–2662, 2663–2672, 2673–2684
3	0	1	2	Vit E, beta-carotene	<i>JNCI</i> 2005; 97: 481–488
4	0	0	3	Vit E	<i>JAMA</i> 2005; 293: 1338–1347
5	0	0	3	Low Fat	<i>JAMA</i> . 2006; 295: 655–666
6	0	0	3	Vit D, Calcium	<i>NEJM</i> 2006; 354: 669–683
7	0	0	2	Folic acid, Vit B6, B12	<i>NEJM</i> 2006; 354: 2764–2772
8	0	0	2	Low Fat	<i>JAMA</i> 2007; 298: 289–298
9	0	0	12	Vit C, Vit E, beta-carotene	<i>Arch Intern Med</i> 2007; 167: 1610–1618
10	0	0	12	Vit C, Vit E	<i>JAMA</i> 2008; 300: 2123–2133
11	0	0	3	Vit E, Selenium	<i>JAMA</i> 2009; 301: 39–51
12	0	0	3	HRT + Vitamins	<i>JAMA</i> 2002; 288: 2431–2440
Totals	0	5	52		

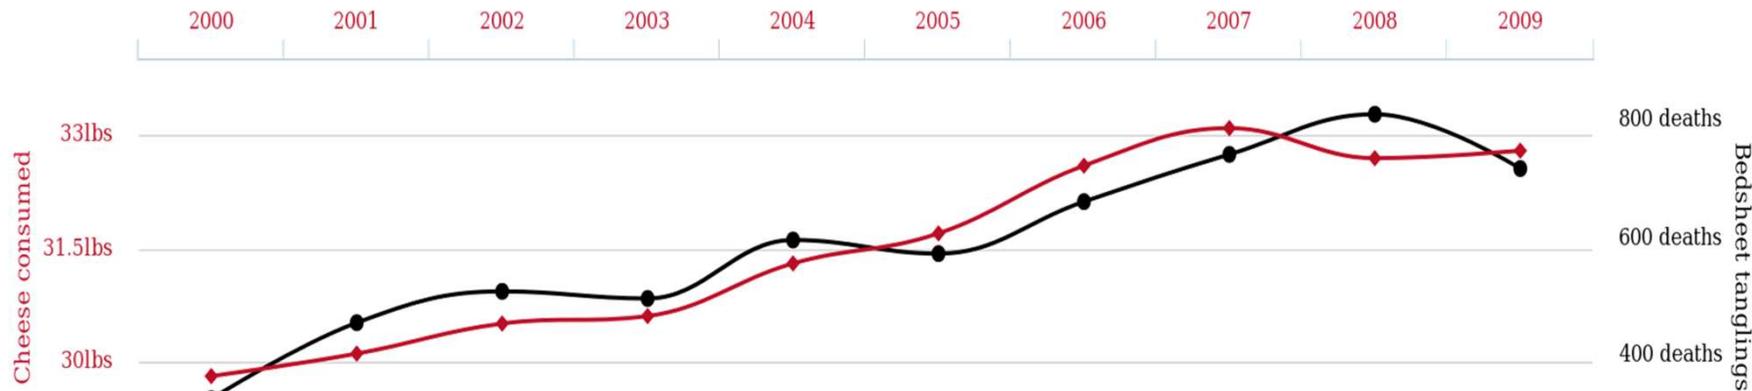
*Young & Karr 2011 Significance, 116:120, 2011*

**Teach:** Correlation  $\neq$  causation

## Per capita cheese consumption

correlates with

## Number of people who died by becoming tangled in their bedsheets



<https://www.tylervigen.com/spurious-correlations>

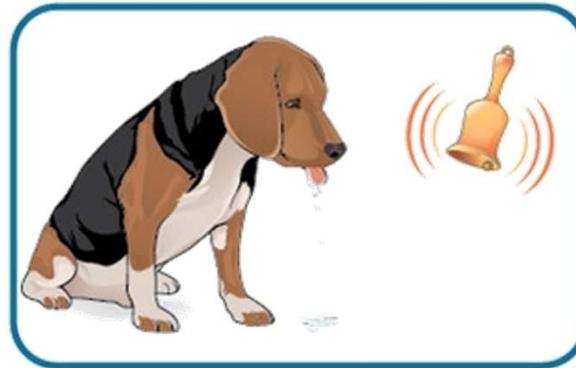
- **Discourage:** Causal inference from non-experimental data

**But:** We are programmed to infer causality...

During Conditioning



After Conditioning



**And:** Spurious correlations have limited scientific and practical interest

**So:** People tend to infer causality regardless

- To infer causality while claiming '*no causal inferences can be made*'

## 5. Discussion

Among 14-year olds living in the UK, we found an association between social media use and depressive symptoms and that this was stronger for girls than for boys. The magnitude of these associations reduced when potential explanatory factors were taken into account, suggesting that experiences of online harassment, poorer sleep quantity and quality, self-esteem and body image largely explain observed associations. There was no evidence of differences for girls and boys in hypothesised pathways between social media use and depressive symptoms. Findings are based largely on cross sectional data and thus causality cannot be inferred.

Our findings add weight to the growing evidence base on the potential pitfalls associated with lengthy time spent engaging on social media. These findings are highly relevant to current policy development on guidelines for the safe use of social media and calls on industry to more tightly regulate hours of social media use for young people [[10], [11]].

Clinical, educational and family settings are all potential points of contact whereby young people could be encouraged to reflect not only on their social media use but also other aspects of their lives including online experiences and their sleep patterns. For instance, in the home setting all family members could reflect on patterns of use and have in place limits for time online, curfews for use and the overnight removal of mobile



BMJ 2015;351:h4596 doi: 10.1136/bmj.h4596 (Published 5 September 2015)

Page 1 of 6

## ANALYSIS



### Increased mortality associated with weekend hospital admission: a case for expanded seven day services?

Nick Freemantle and colleagues discuss the findings of their updated analysis of weekend admissions and the implications for service design

Nick Freemantle *professor of clinical epidemiology and biostatistics*<sup>1,2</sup>, Daniel Ray *professor of health informatics*<sup>2,3,4</sup>, David McNulty *medical statistician*<sup>2,3</sup>, David Rosser *medical director*<sup>5</sup>, Simon Bennett *director, clinical policy and professional standards*<sup>6</sup>, Bruce E Keogh *national medical director*<sup>6</sup>, Domenico Pagano *professor, cardiac surgery*<sup>2,7</sup>

<sup>1</sup>Department of Primary Care and Population Health, University College London, UK; <sup>2</sup>Quality and Outcomes Research Unit, University Hospitals Birmingham NHS Foundation Trust, Birmingham, UK; <sup>3</sup>Department of Informatics, University Hospitals Birmingham NHS Foundation Trust; <sup>4</sup>Farr Institute of Health Informatics Research, University College London; <sup>5</sup>University Hospitals Birmingham NHS Foundation Trust; <sup>6</sup>Medical Directorate, NHS England, London, UK; <sup>7</sup>Department of Cardiothoracic Surgery, University Hospitals Birmingham NHS Foundation Trust

*Freemantle et al 2015  
BMJ 5;351*

The  
Alan Turing  
Institute



home > politics election UK election world sport football opinion cul ≡ all

Health policy Michael White's political briefing

## Quiet hospitals kill, but mindless union bashing can give us a nasty injury

Jeremy Hunt has every right to question NHS consultants' working practices but there are questions they should be allowed to ask him too



Hospital equipment is being underused. Photograph: Justin Paget/Justin Paget/Corbis



home > politics election UK election world sport football opinion cul ≡ all

Health policy

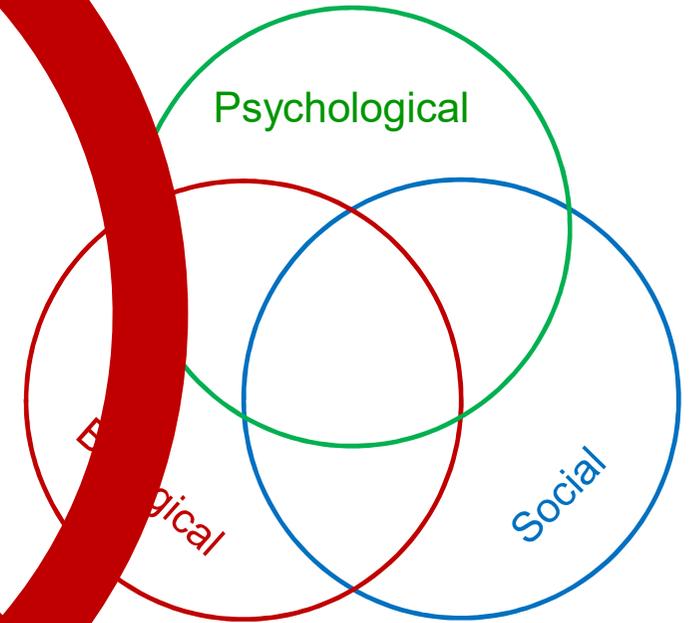
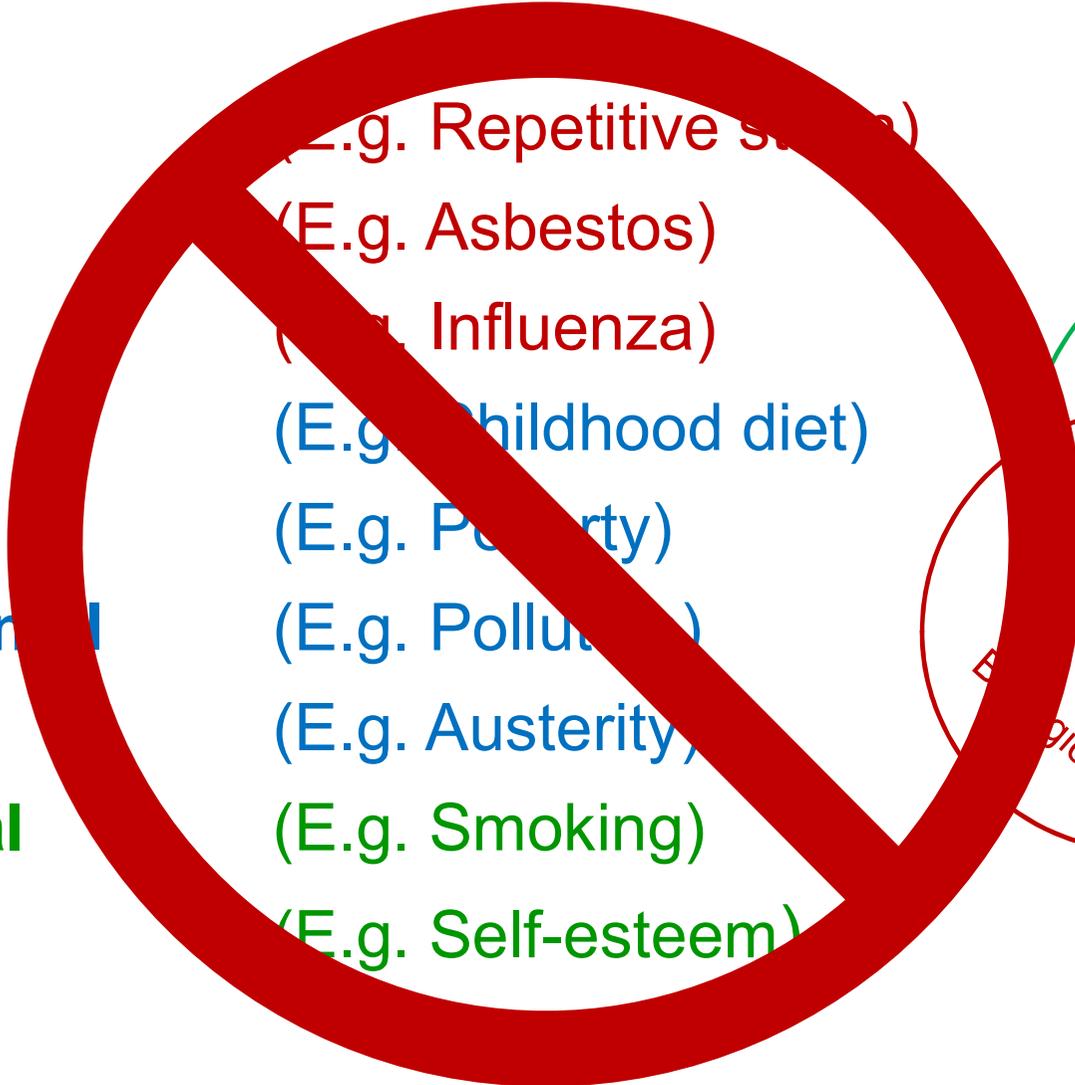
## Two deaths possibly linked to 'Hunt effect', study suggests

Research suggests some patients may be avoiding going to hospital at the weekend because of health secretary's statements about inadequate NHS staffing levels



Almost a third of the 40 patients studied suffered an increased chance of dying as a result of waiting to seek help. Photograph: Chris Radburn/PA

- Physical** (E.g. Repetitive strain)
- Chemical** (E.g. Asbestos)
- Biological** (E.g. Influenza)
- Cultural** (E.g. Childhood diet)
- Economic** (E.g. Poverty)
- Environmental** (E.g. Pollution)
- Political** (E.g. Austerity)
- Behavioural** (E.g. Smoking)
- Personal** (E.g. Self-esteem)



## Accept & admit our causal ambitions

### The C-Word: Scientific Euphemisms Do Not Improve Causal Inference From Observational Data

Causal inference is a core task of science. However, authors and editors often refrain from explicitly acknowledging the causal goal of research projects; they refer to causal effect estimates as a estimates.

This commentary using the term "causal" to improve the observational research.

Miguel A. Hernán, MD, DrPH

See also Galea and Vaughan, p. 602; Begg and March, p. 620; Ahern, p. 621; Chiolero, Glymour and Hamad, p. 623; Jones and Schooling, p. 624; and Hernán, p. 625.

### Data Are Not Enough—Hurray For Causality!

See also Galea and Vaughan, p. 602; Hernán, p. 616; Begg and March, p. 620; Ahern, p. 621; Glymour and Hamad, p. 623; Jones and Schooling, p. 624; and Hernán, p. 625.

Causal inference is of major importance in epidemiology and public health because the determination that an association between an exposure and a health outcome is causal indicates a potential for intervention to improve health. However, causal inference had neither been formalized nor taught in epidemiology. Indeed, training epidemiologists, health scientists, and public health practitioners in causality has often been limited to the study of complex epidemiological and

analyses in complex because the complexity applied to weak as evidence for causal effect findings in public health

### The C-Word: The More We Discuss It, the Less Dirty It Sounds

See also Galea and Vaughan, p. 602; Hernán, p. 616; Begg and March, p. 620; Ahern, p. 621; Chiolero, p. 622; Glymour and Hamad, p. 623; and Jones and Schooling, p. 624.

I thank Chiolero (p. 622), Ahern (p. 621), Glymour and Hamad (p. 623), Jones and Schooling (p. 624), and Begg and

### Start With the "C-Word," Follow the Roadmap for Causal Inference

See also Galea and Vaughan, p. 602; Hernán, p. 616; Begg and March, p. 620; Chiolero, p. 622; Glymour and Hamad, p. 623; Jones and Schooling, p. 624; and Hernán, p. 625.

In this issue of *AJPH*, Hernán (p. 616) argues that when the aim of an investigation is to estimate a causal effect, researchers should be allowed to say so. I agree, and extend this conversation to argue that epidemiology will help

In following the roadmap, the causal question starts the investigation, consistent with Hernán's call. In the subsequent step of identifiability, the specific ways in which a particular application

maximum likelihood estimation (TMLE; [bit.ly/2FVfdjf](http://bit.ly/2FVfdjf)).<sup>4,5</sup> As these newer statistical methods have become more widely used, traditional analysis approaches such as regression have begun to seem that they must fall short when it comes to causality.

This is one common conflation

appropriate incorporation of expert knowledge in the process. Indeed, incorrect causal inferences from observational data are often the result of a flawed emulation of the basic design of the target trial (e.g., choice of time zero and classification of

Again, the process of specifying and emulating a target trial helps by providing a systematic way to explore each type of bias and its potential influence on the effect estimate. The Cochrane tool has adopted this target trial-based approach to assess the risk of bias of nonrandomized studies.<sup>3</sup>

- Upgrade our **epistemological (philosophical) machinery**



"How can we infer causal relations from observations?"



$$p(\mathbf{X} = \mathbf{x}) = \prod_i p(X_i = x_i | \text{PA}_i = \text{pa}_i)$$

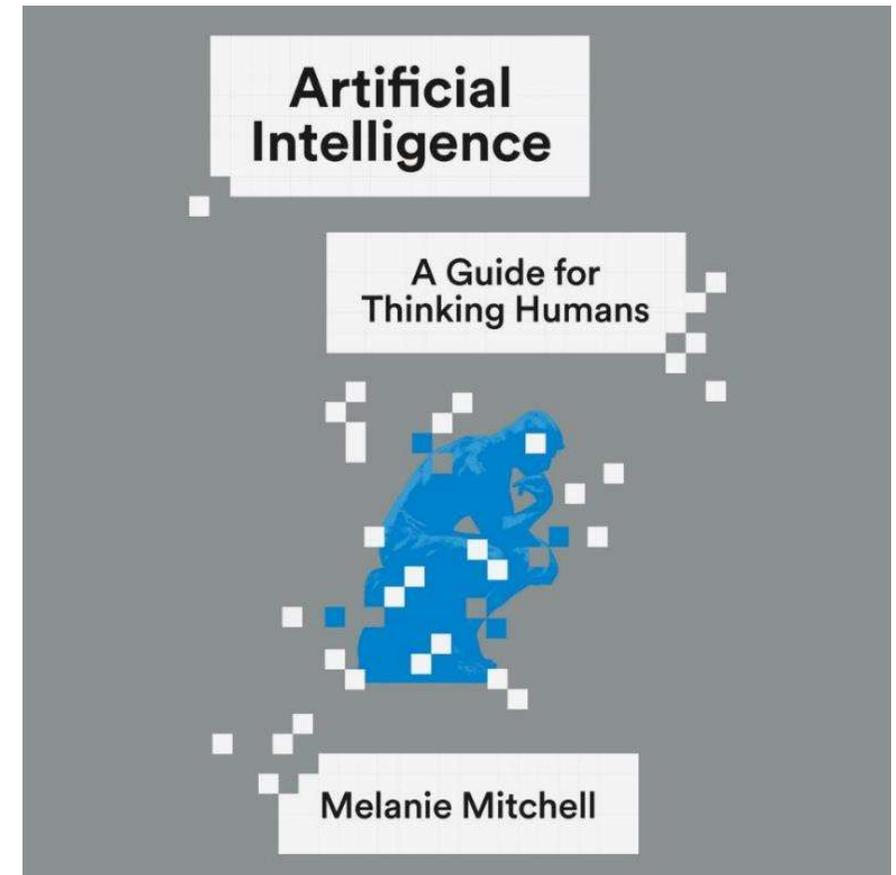


**Philosophy**

**Maths**

**Data science**

- Machine learning is excellent at finding patterns in complex data > **prediction**
- Causal inference requires identifying and estimating **counterfactuals** (i.e. events that did **NOT** happen).
- In non-experimental data, machine learning can't do **causal inference** because you **CAN'T** learn what's not there!
- Need external knowledge of '**data generating process**'; the story behind how the data came into being

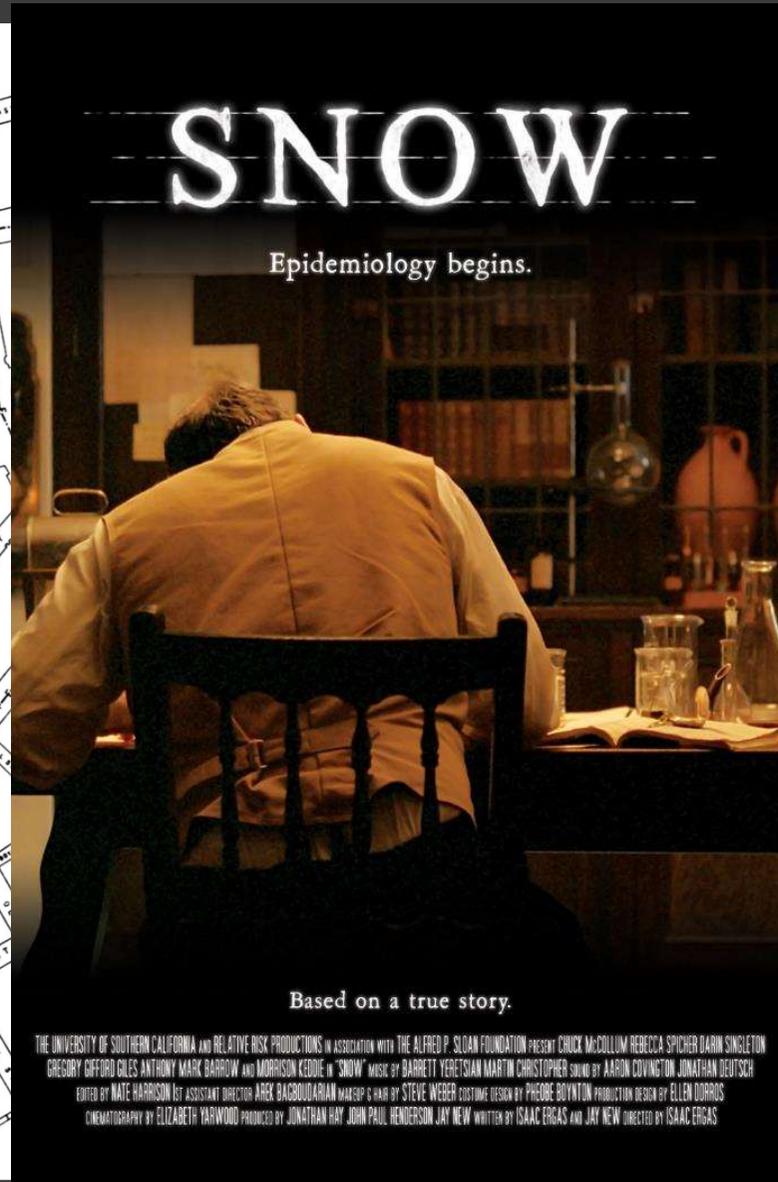


# LEGENDS AND MYTHS

#TMCFTURING



JOHN SNOW, LEGEND,  
DATA SCIENTIST



'BROAD STREET' PUMP

The  
Alan Turing  
Institute

1854

31 Aug Cholera outbreak in Soho

5 Sep Snow man where deaths were occurring

6 Sep 83% of people drank from Broad Street pump

7 Sep Snow melts, parish guardian to argue for pump closure

8 Sep Snow removed, cholera outbreak ends

**FAKE**

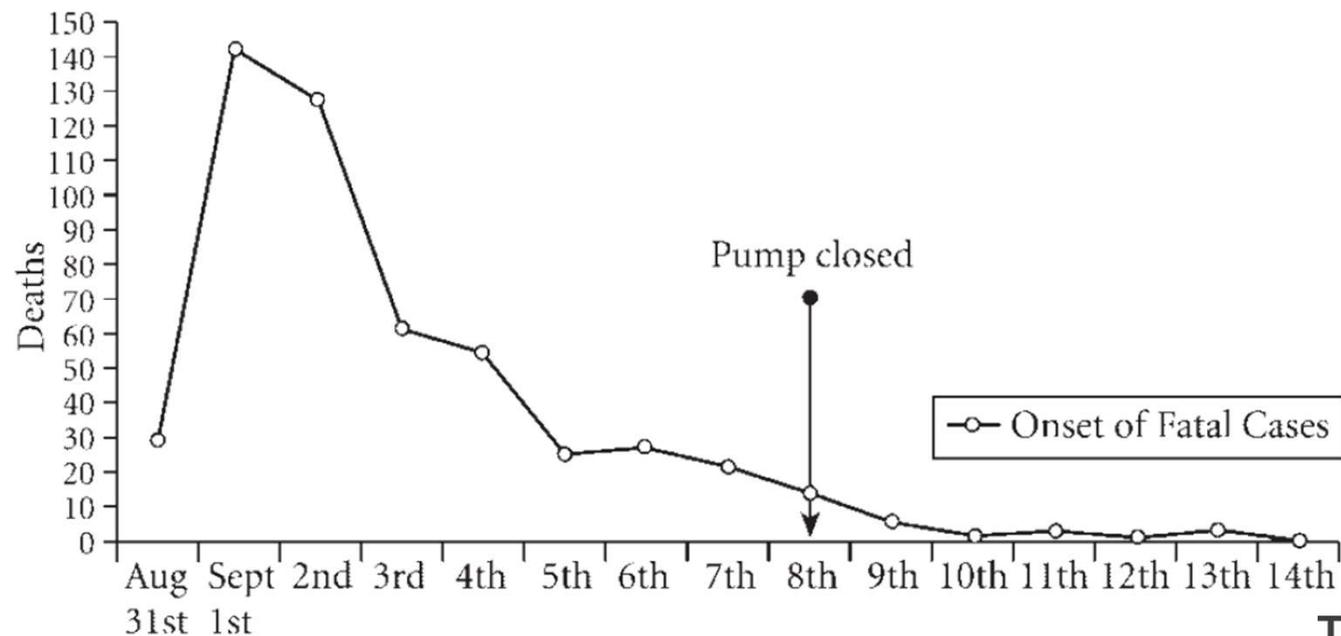
**NEWS**





**JOHN SNOW, LEGEND,  
DATA SCIENTIST**

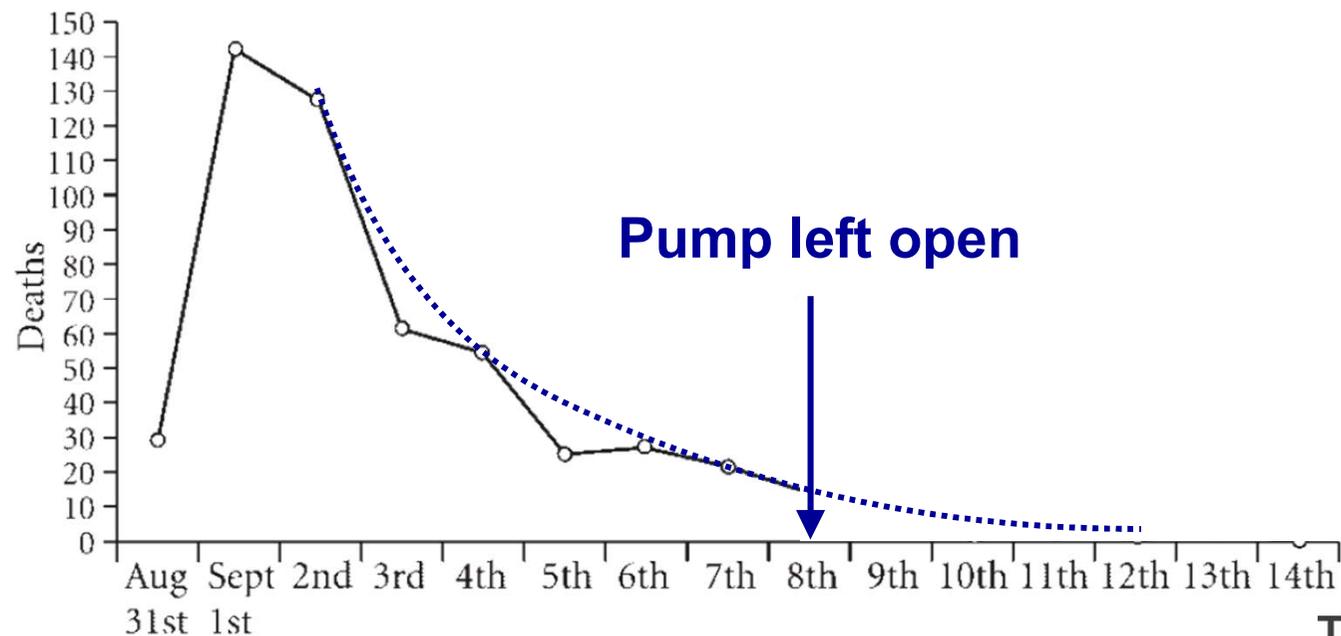
Snow observed the **outcome** of what happened when the pump was 'closed' (**factual**)





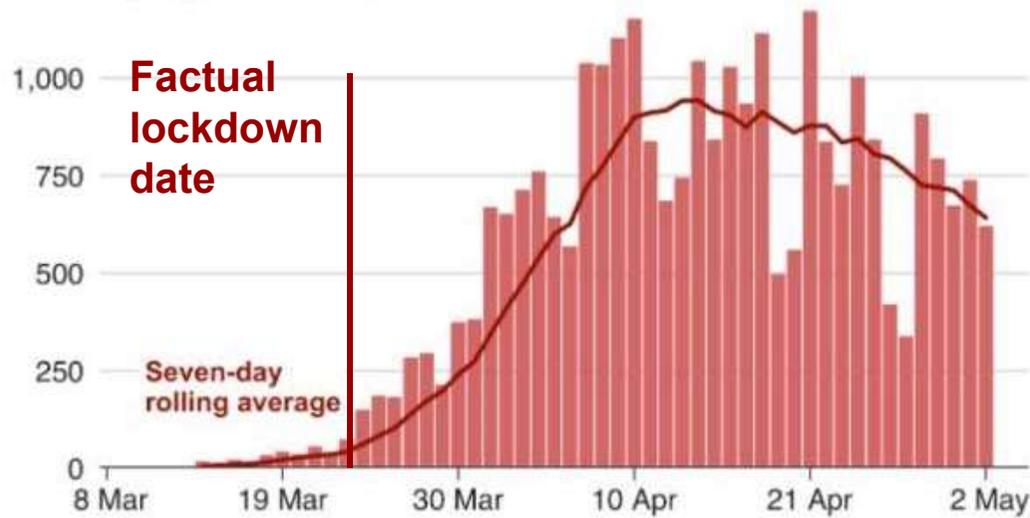
JOHN SNOW, LEGEND,  
DATA SCIENTIST

Didn't observe the **potential outcome** that *would have happened* if the pump had been – counter to fact - left 'open' (**counterfactual**)



- We can **never** know the **potential outcome** for a **counterfactual exposure!**
- For each 'unit of analysis' we can only observe one potential outcome
- This is known as the **fundamental problem of causal inference**

UK daily reported deaths with coronavirus



UK daily reported deaths with coronavirus

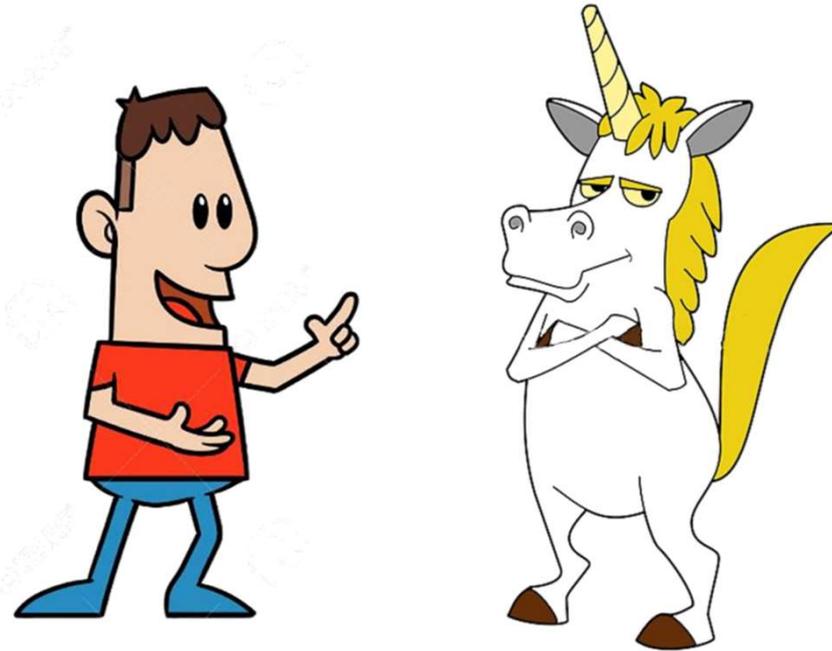


Instead we must *estimate* the **potential outcome** for the **counterfactual exposure** from **exchangeable units of analysis**



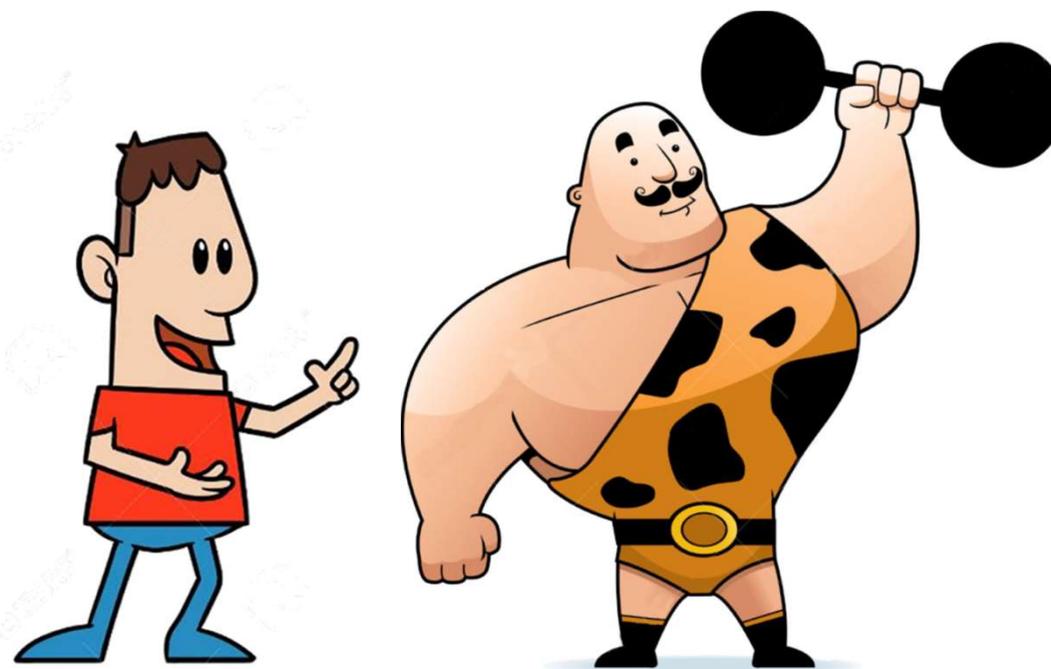
GIVE **MAGIC BEANS**





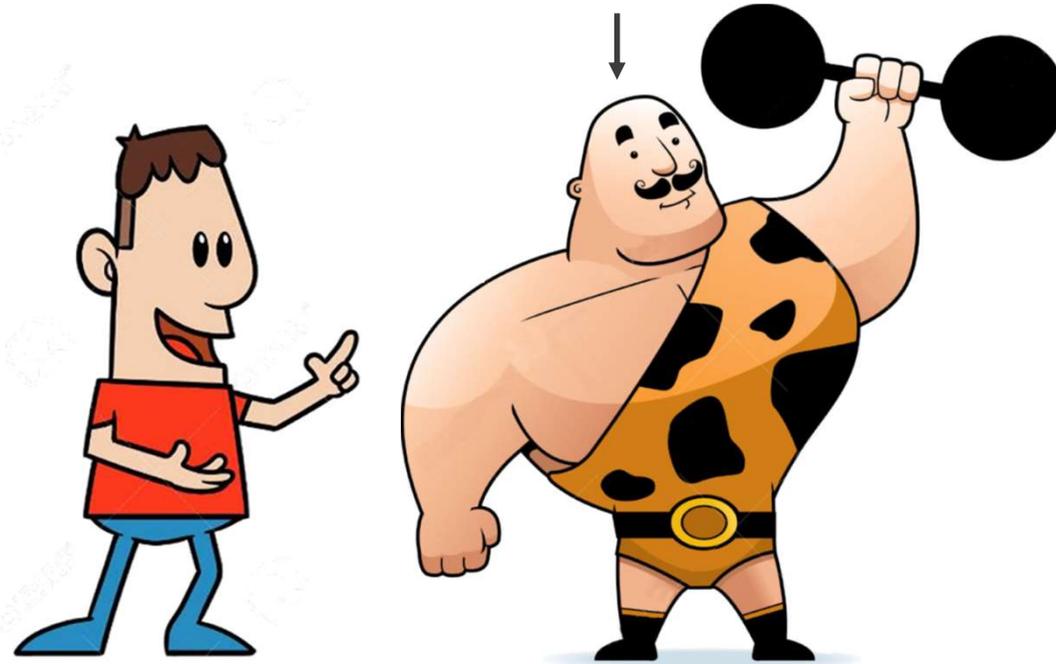
## Problem:

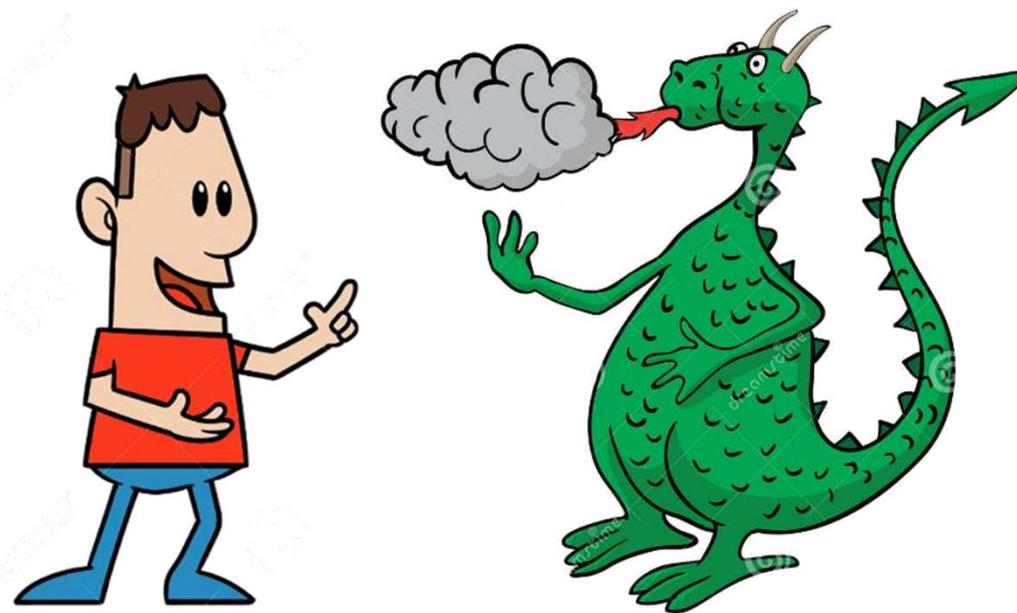
- Units (e.g. people) are **very different**
- Even the **same units** can respond **differently** at **different times**



Not **exchangeable** units of analysis!

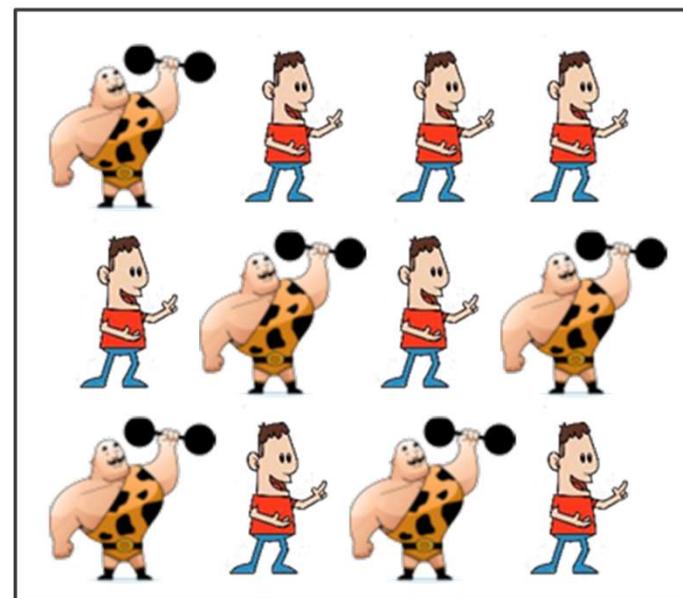
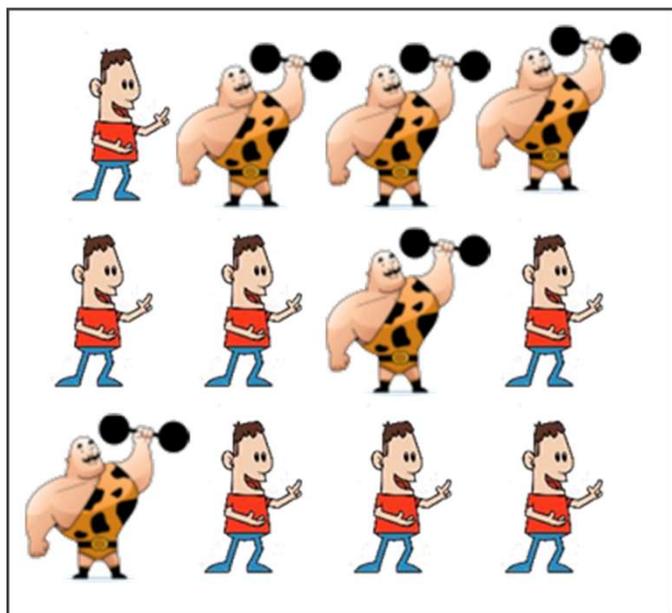
GIVE **MAGIC BEANS**

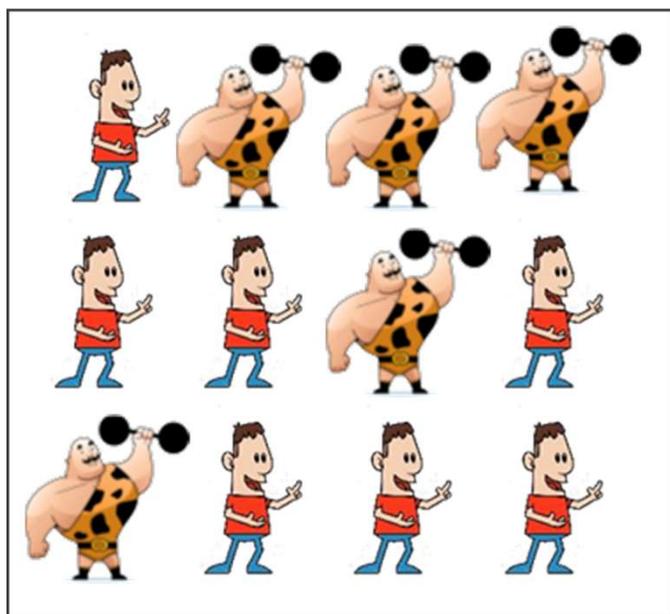




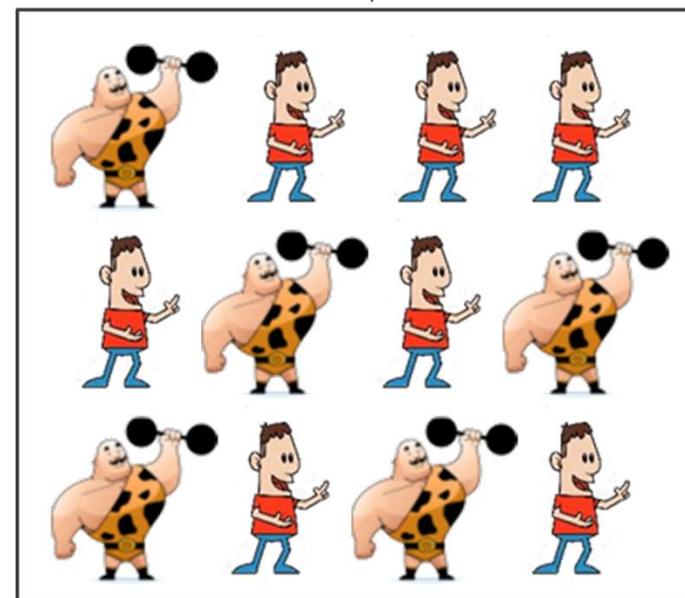
Poor estimate of **potential outcome!**

- We therefore have to **identify** (*sub*)*groups* of units that were **exchangeable** at the time of exposure
- We can then estimate the **average causal effect** by **comparing the outcomes** between these (sub)groups



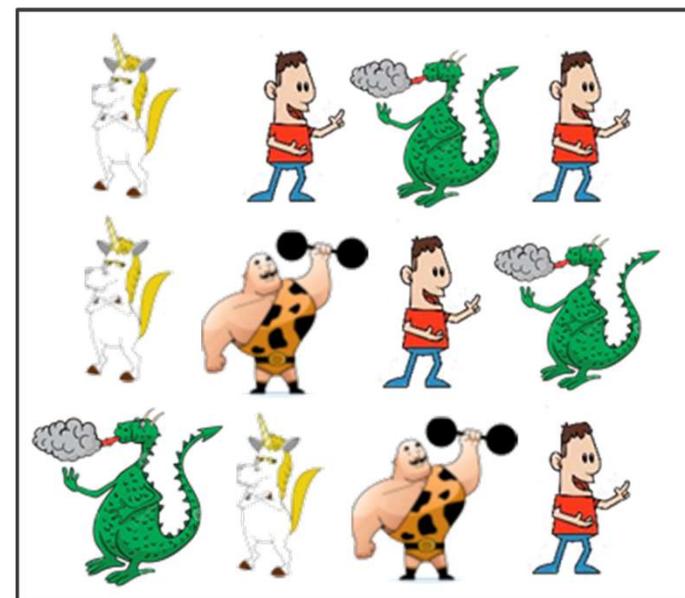
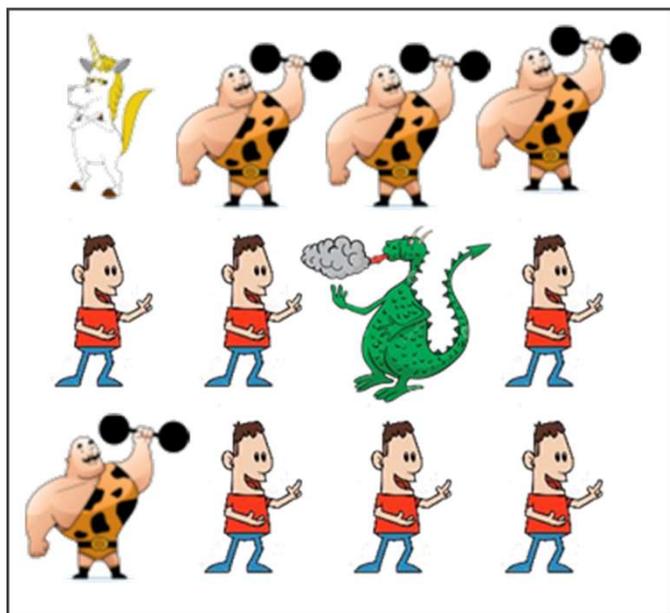


GIVE **MAGIC BEANS**

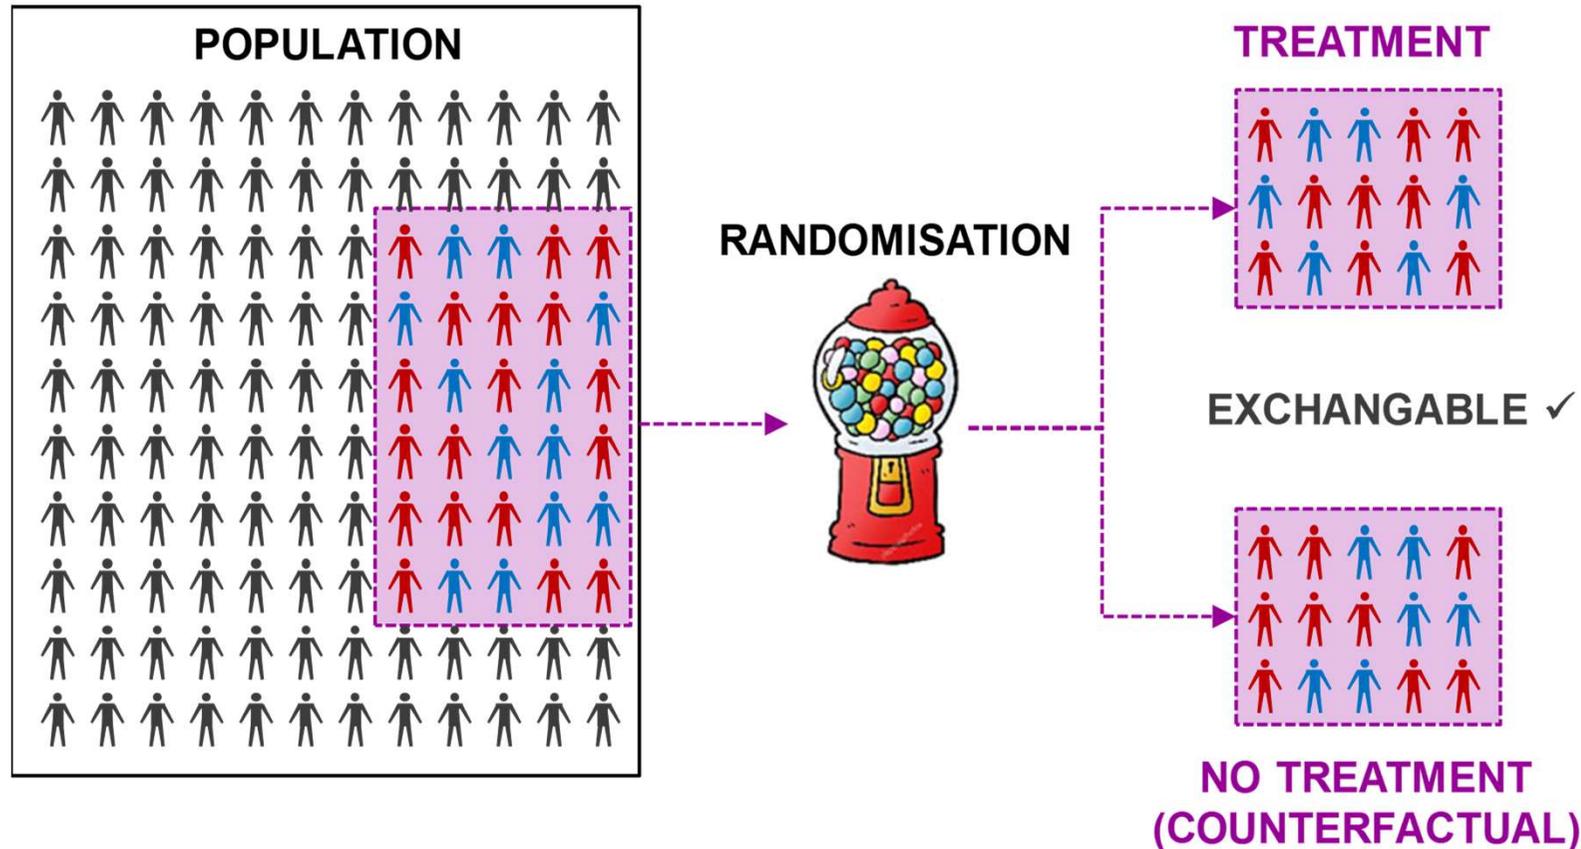


# AVERAGE CAUSAL EFFECTS

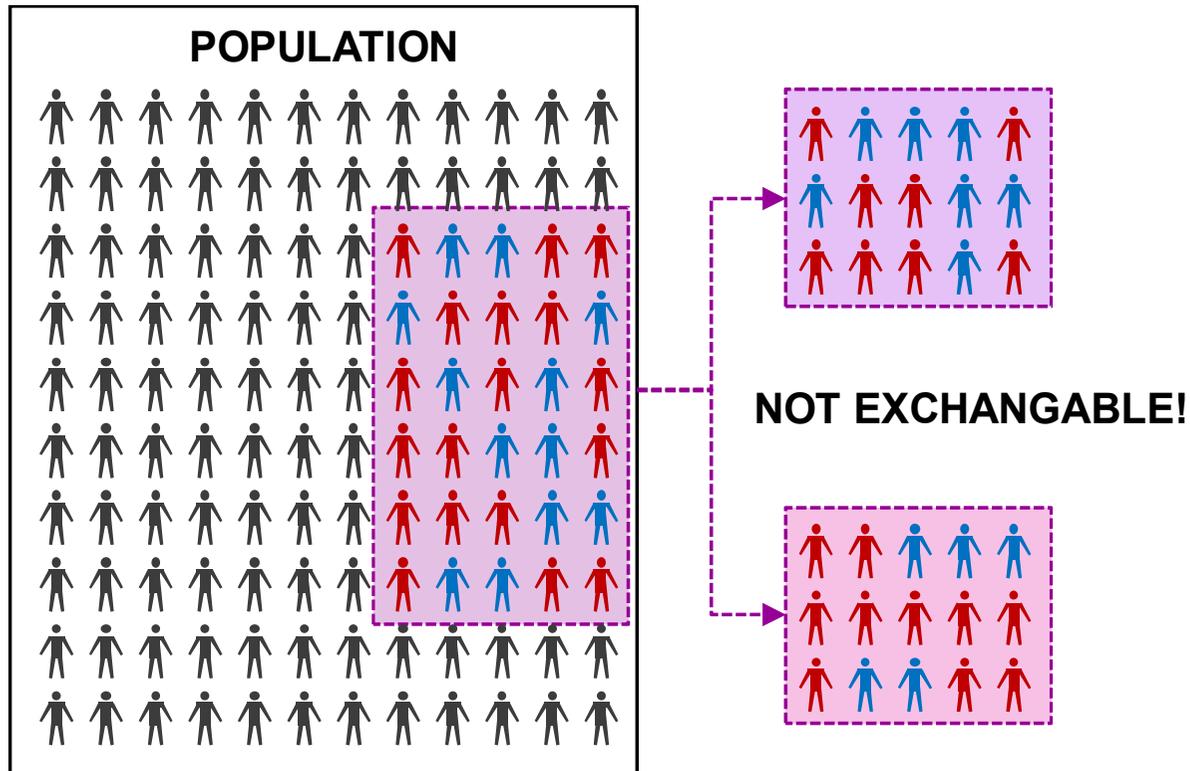
#TMCFTURING



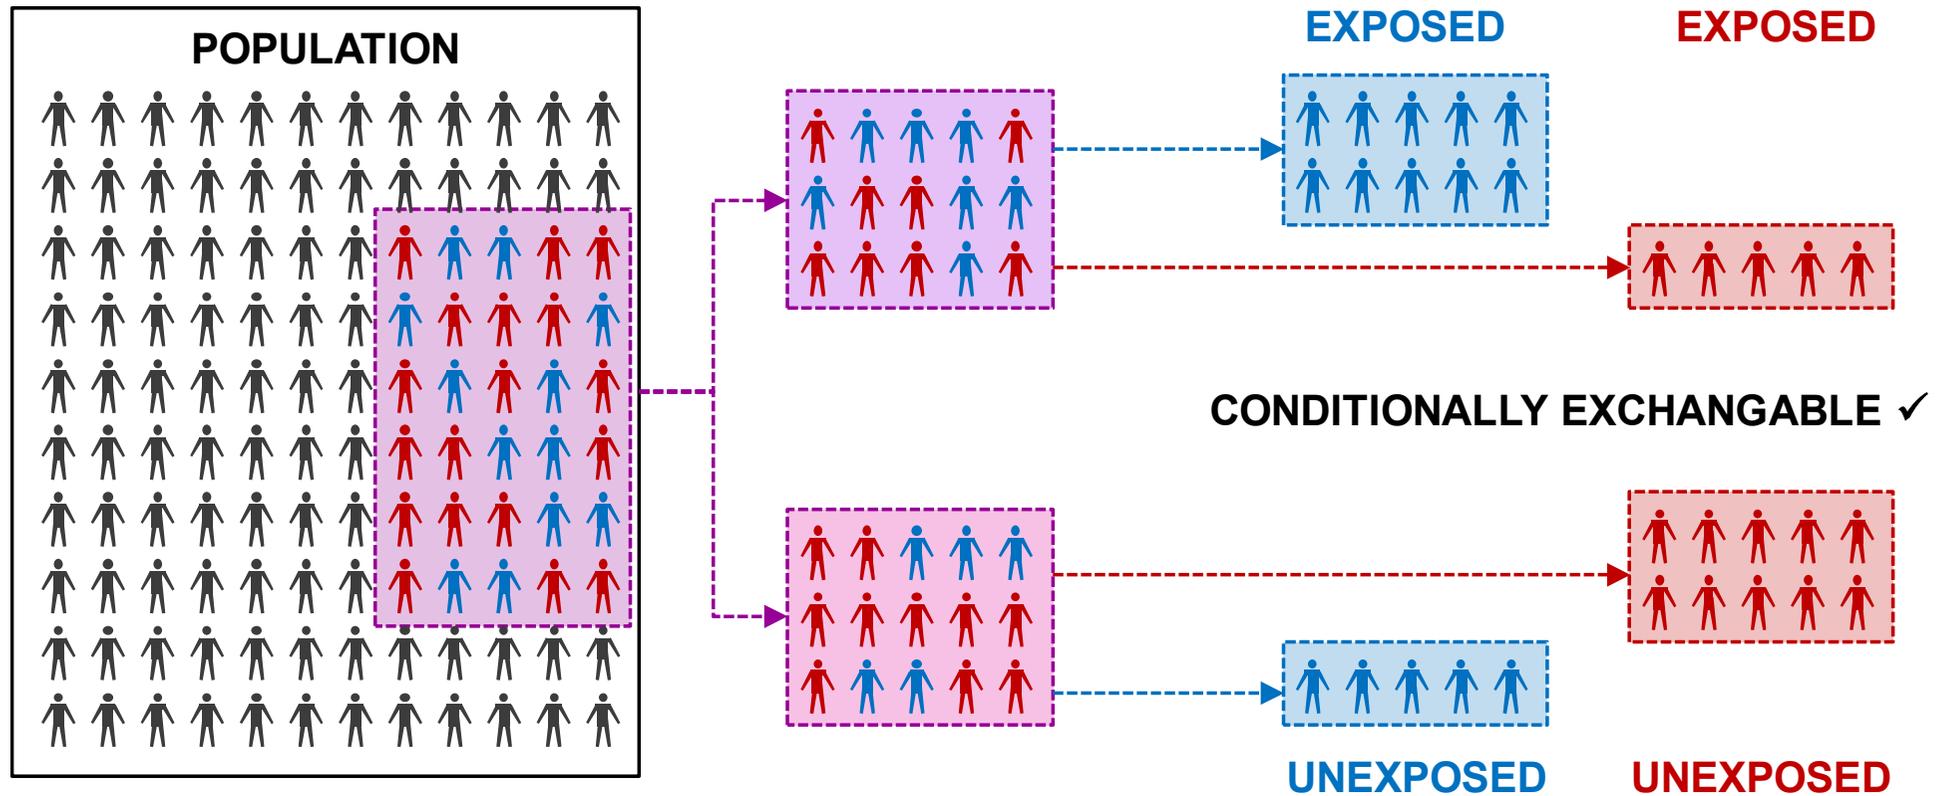
- The easiest way to do this is through randomisation



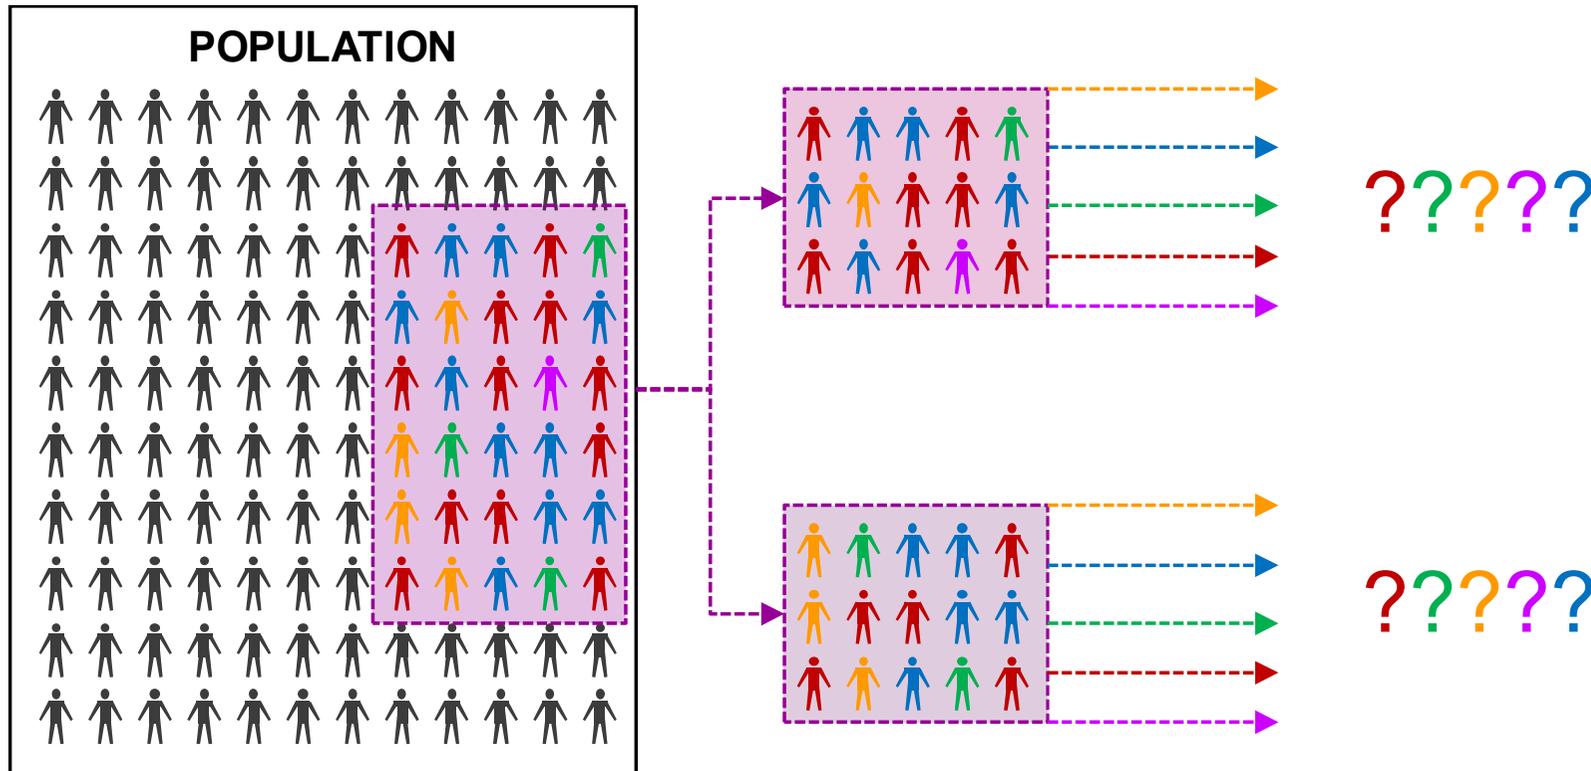
- Without randomisation, you have to identify and compare like-for-like subgroups



- Without randomisation, you have to identify and compare like-for-like subgroups



- Without randomisation, you have to identify and compare like-for-like subgroups

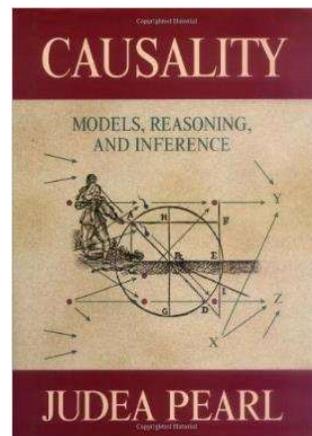


- Need to identify, measure, & control for every variable that (potentially) *causes* both the exposure and outcome (**confounder**)

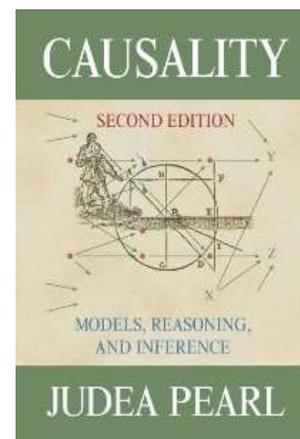


Judea Pearl, Computer scientist, philosopher, Turing Prize winner

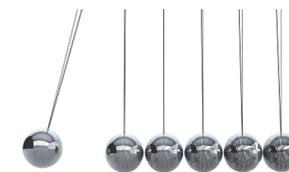
- **Causal inference methods** – like Judea Pearl’s ‘**Structural Causal Model**’ – provide a mathematical and philosophical framework for considering causal effects drawing on:
  - **Probability theory**
  - **Counterfactual reasoning**
  - **Graphical model theory**



2000



2009



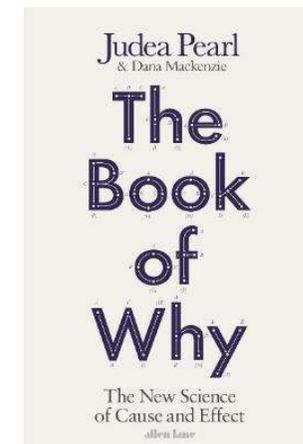
CAUSAL INFERENCE  
IN STATISTICS

A Primer

Judea Pearl  
Madelyn Glymour  
Nicholas P. Jewell

WILEY

2016



2018

- **Key feature:** Formally **identifying** your effect before analysis!

**ESTIMAND**  
What you seek



E.g. The true difference in Y  
due to exposure

**ESTIMATOR**  
How you will get there

### Method

1. Preheat your oven to 190°C /170°F / Gas Mark 5. Grease and line the base of 2 cake tins, one 8 inch/20cm and one 6 inch/15cm
2. Cream together the butter and caster sugar until light and fluffy.
3. Add the eggs one at a time with a spoonful of flour and blend in well.
4. Sift in the flour and baking powder and gently fold in. Finally add the milk and mix until you have a smooth batter.
5. Pour 1/3 of the batter into the small tin and 2/3 into the large tin.
6. Bake on the same shelf in the preheated oven, the smaller tin at the front.
7. Check the smaller cake after 20 minutes. When it is cooked remove from the oven, leaving the larger one still baking. The large cake should be done by 30 minutes.
8. Leave the cakes for 5 minutes in the tins, then turn out onto a rack to cool completely.
9. To make the icing, beat together the butter and icing sugar, add the vanilla and then the milk. Whisk the icing hard using an electric stand mixer if you can. Whisk it for 5 minutes and it will become really pale and light.

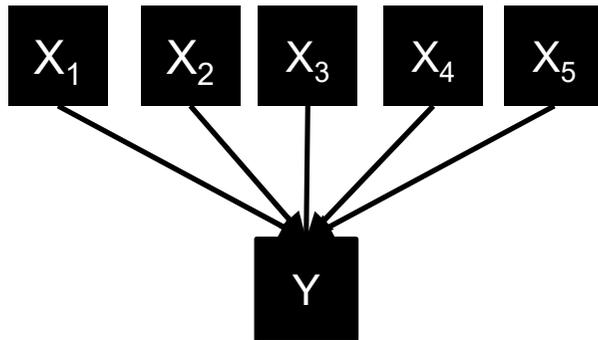
E.g. Your regression  
model

**ESTIMATE**  
What you get

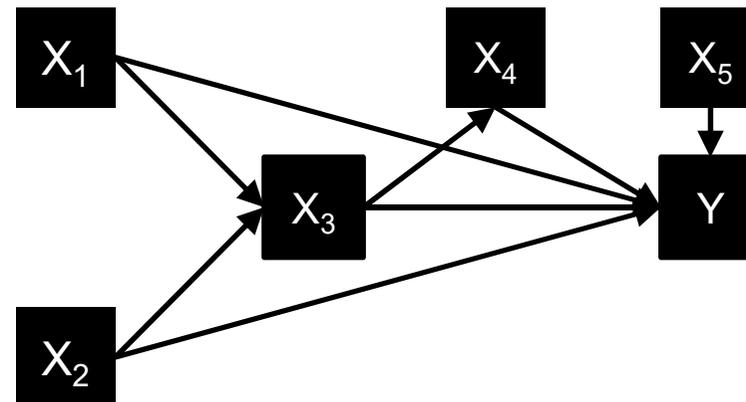


E.g. the estimated difference  
in Y from model coefficient

- **Causal diagrams** – such as **directed acyclic graphs** – encode counterfactual & probabilistic reasoning into diagrammatic form
- **First benefit:** helps to identify which variables are **confounders** and need controlling to provide exchangeable comparisons
- We must do this with our **outside knowledge**, the machine/software cannot!

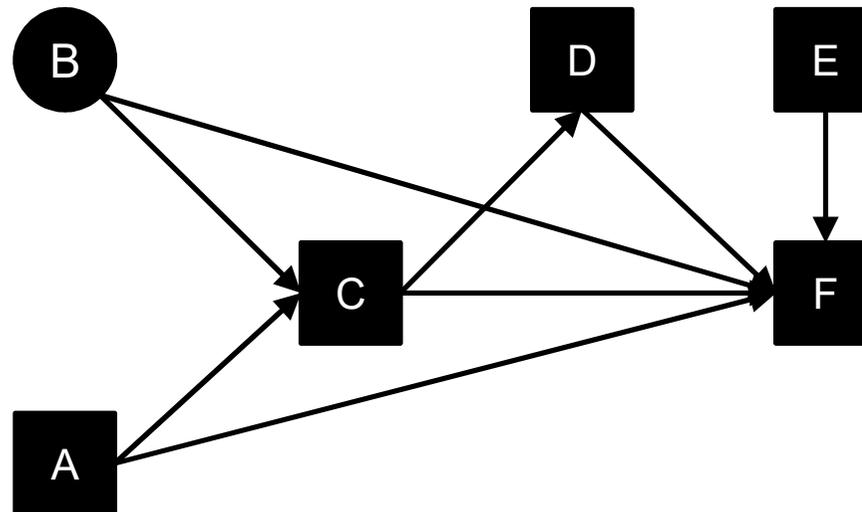


How the machine/software sees it!

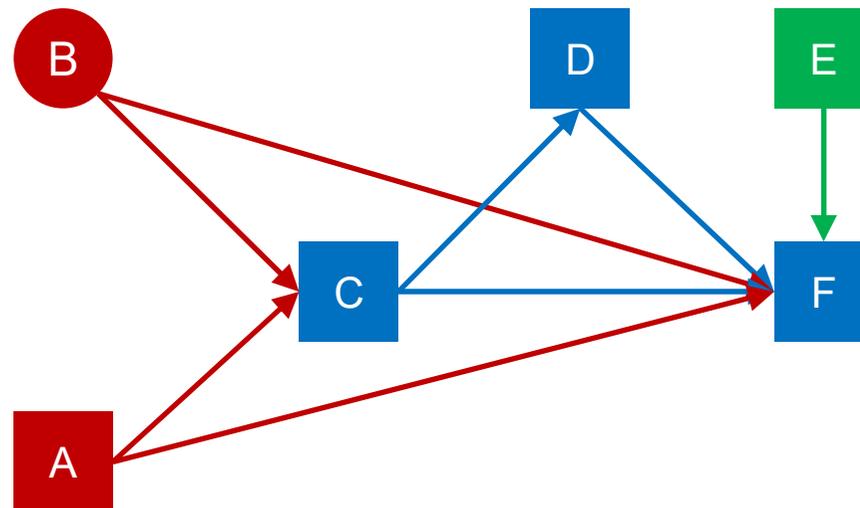


How nature created it!

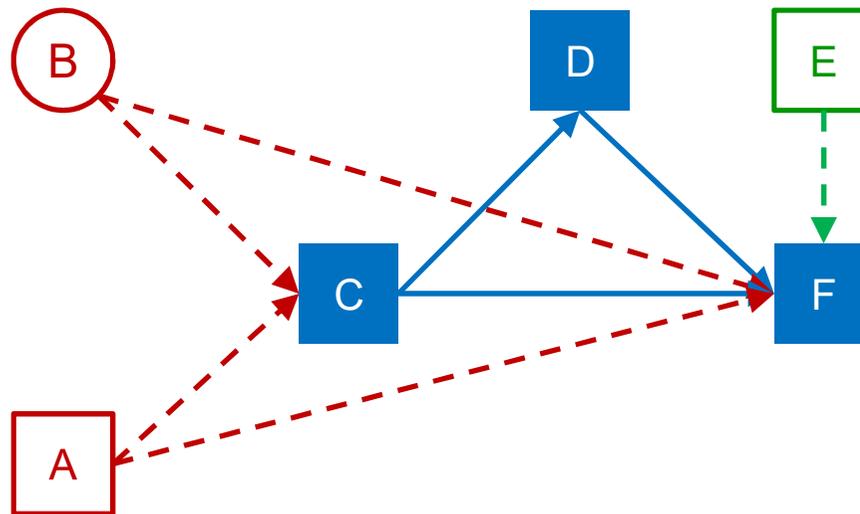
- **Directed Acyclic Graphs (DAGs)** are nonparametric representations of the (hypothesised) causal relationships between variables
  - Relationships between variables ('**nodes**') are represented by arrows ('**arcs**') creating **paths** between them
  - Paths can be **causal** or **confounding**; **open** or **closed**.
  - Open paths transmit correlations, closed paths do not



- To estimate the causal effect of **C** on **F** (our '**focal relationship**'):
  - We want all **causal paths** open
  - And all **confounded paths** closed

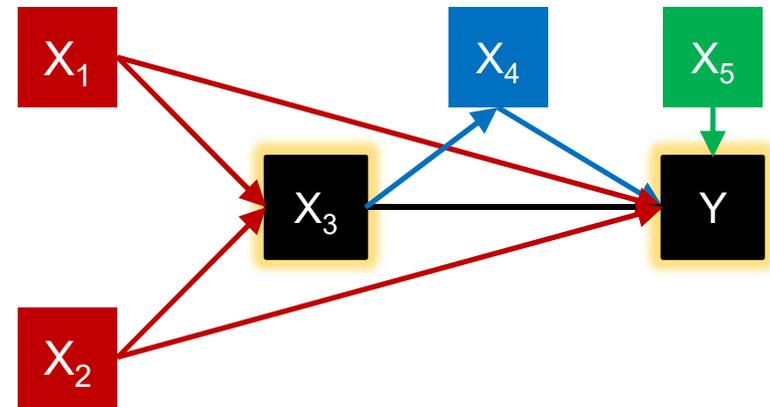
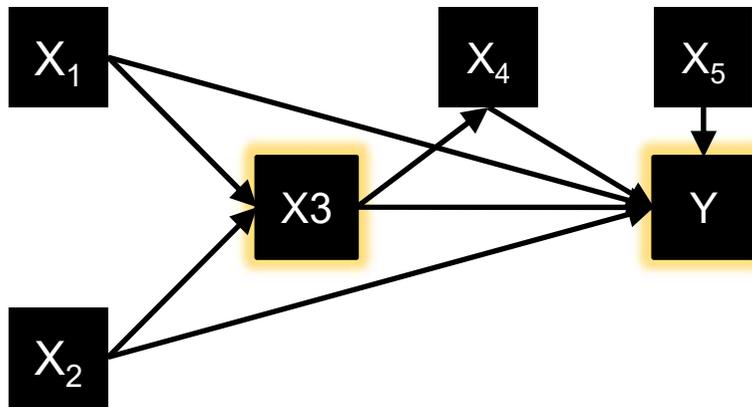


- To estimate the causal effect of **C** on **F** (the '**focal relationship**'):
  - We want all **causal paths** open
  - And all **confounded paths** closed
  - This means **controlling for all confounders** but no **mediators**



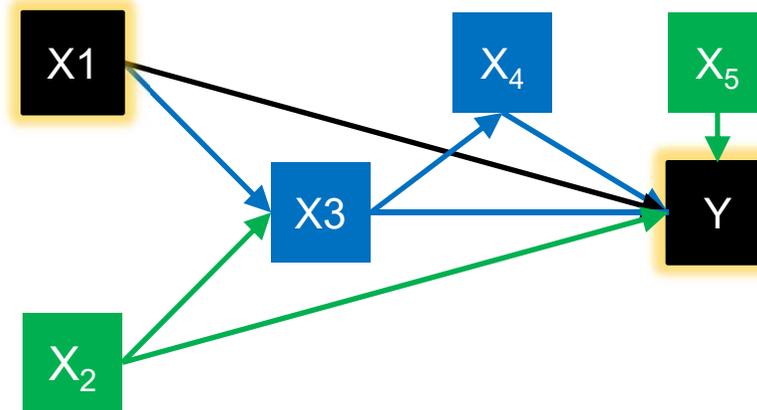
**Example:** Total causal effect of  $X_3$  on  $Y$ :

- Model should include **confounders** ( $X_1$ ,  $X_2$ ) [and **competing exposures** ( $X_5$ )]
- $Y \sim X_3 + X_1 + X_2 [+ X_5]$



- This model was constructed to estimate the effect of  $X_3$  on  $Y$
- We should not interpret the other covariates ( $X_1$ ,  $X_2$ ,  $X_5$ ), because they will not be appropriately controlled
- E.g. The coefficient on  $X_1$  is **NOT** the total causal effect of  $X_1$  on  $Y$ , because the inclusion of  $X_3$  means the causal path  $X_1 \rightarrow X_3 \rightarrow Y$  will be closed.

$$Y \sim X_3 + X_1 + X_2 + X_5$$



- The tradition of including all ‘predictors’ of our outcome (Y) in a **single model**, and interpreting the coefficients ( $X_1, X_2, X_3, X_4, X_5$ ) as ‘**mutually adjusted**’ effects has been dubbed the ‘**Table 2 Fallacy**’

---

## The Table 2 Fallacy: Presenting and Interpreting Confounder and Modifier Coefficients

Daniel Westreich\* and Sander Greenland

\* Correspondence to Dr. Daniel Westreich, Department of Obstetrics and Gynecology, Duke Global Health Institute, Duke University, DUMC 3967, Durham, NC 27710 (e-mail: daniel.westreich@duke.edu).

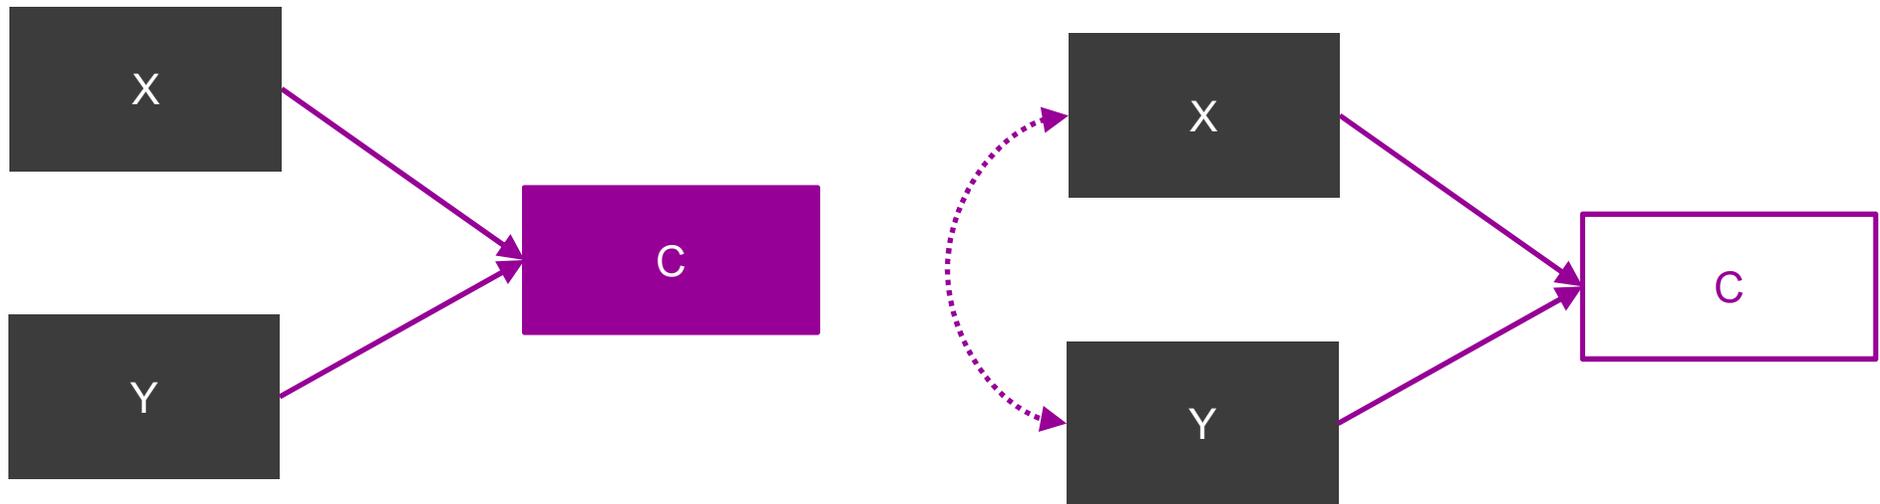
*Initially submitted January 13, 2012; accepted for publication October 11, 2012.*

---

It is common to present multiple adjusted effect estimates from a single model in a single table. For example, a table might show odds ratios for one or more exposures and also for several confounders from a single logistic regression. This can lead to mistaken interpretations of these estimates. We use causal diagrams to display the sources of the problems. Presentation of exposure and confounder effect estimates from a single model may lead to several interpretative difficulties, inviting confusion of direct-effect estimates with total-effect estimates for covariates in the model. These effect estimates may also be confounded even though the effect estimate for the main exposure is not confounded. Interpretation of these effect estimates is further complicated by heterogeneity (variation, modification) of the exposure effect measure across covariate levels. We offer suggestions to limit potential misunderstandings when multiple effect estimates are presented, including precise distinction between total and direct effect measures from a single model, and use of multiple models tailored to yield total-effect estimates for covariates.

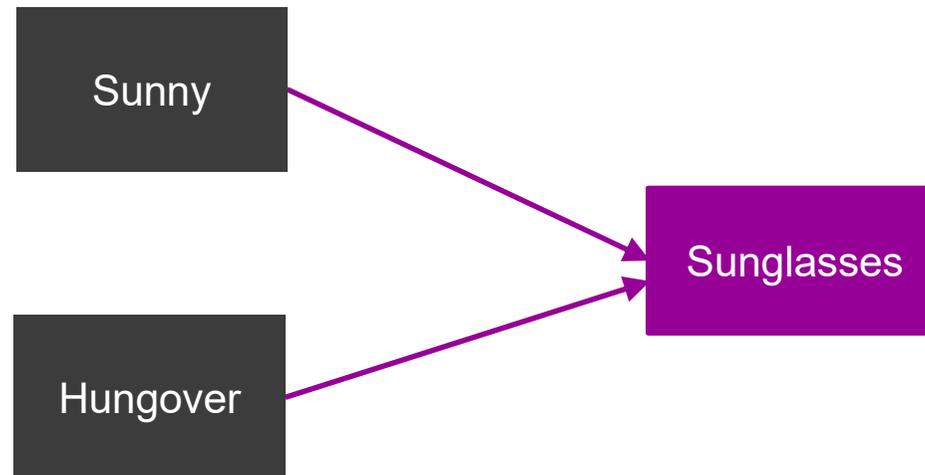
Westreich & Greenland. *Am J Epidemiol.* 2013;177(4):292–298.

- DAGs have also been revolutionary in helping us understand **conditional dependencies** – non-causal associations introduced by conditioning



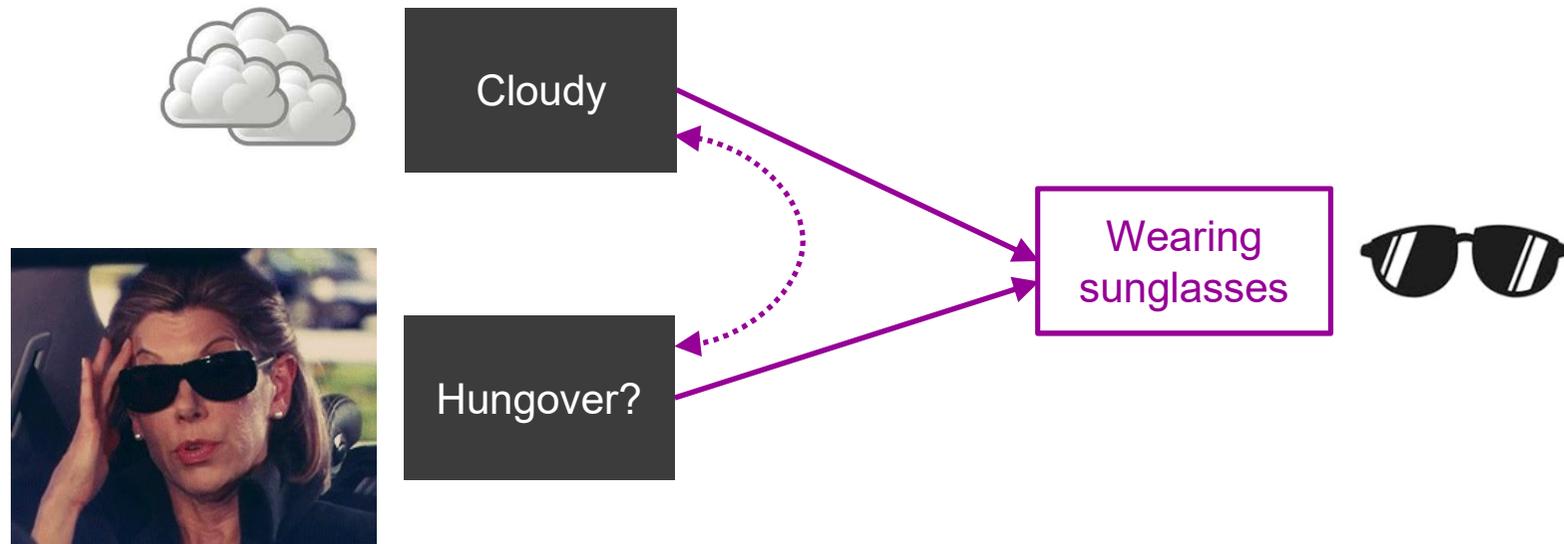
- Conditioning on **C** would transmit a **non-causal association** between **X** and **Y**

- **Example:** The Sunglasses Revelation
  - British people typically wear sunglasses when the **sun is shining**.
  - They sometimes also wear them when they are **very hungover**



- On average, suppose hangovers occur equally all year round, regardless of weather

- **Example:** The Sunglasses Revelation
  - You see someone **wearing sunglasses** on a **cloudy day**



- What does this tell you about their likelihood of being **hungover**?



American Journal of Epidemiology  
Copyright © 2006 by the Johns Hopkins Bloomberg School of Public Health  
All rights reserved; printed in U.S.A.

Vol. 164, No. 11  
DOI: 10.1093/aje/kwj275  
Advance Access publication August 24, 2006

## Practice of Epidemiology

### The Birth Weight “Paradox” Uncovered?

**Sonia Hernández-Díaz<sup>1,2</sup>, Enrique F. Schisterman<sup>3</sup>, and Miguel A. Hernán<sup>1</sup>**

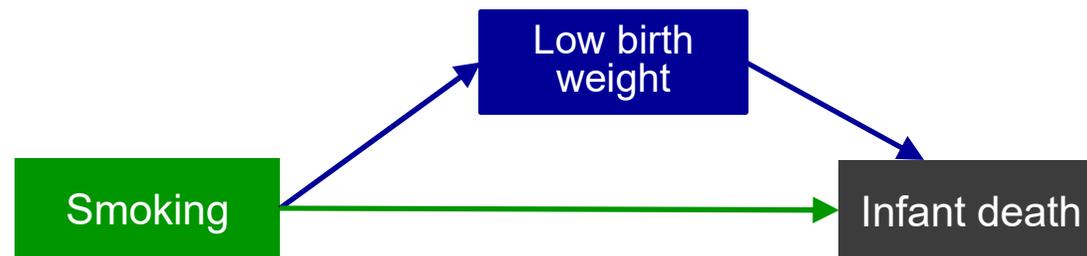
<sup>1</sup> Department of Epidemiology, Harvard School of Public Health, Boston, MA.

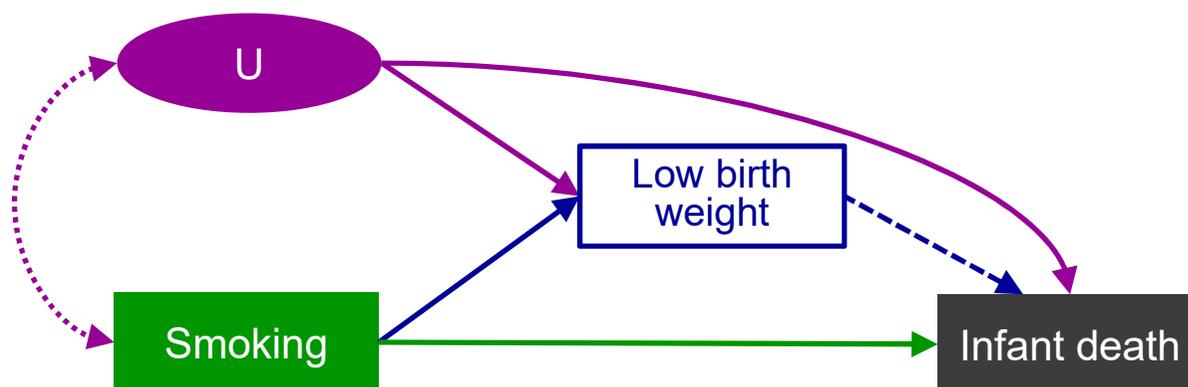
<sup>2</sup> Slone Epidemiology Center, Boston University, Boston, MA.

<sup>3</sup> Epidemiology Branch, National Institute of Child Health and Human Development, Bethesda, MD.

*Received for publication February 7, 2005; accepted for publication January 23, 2006.*

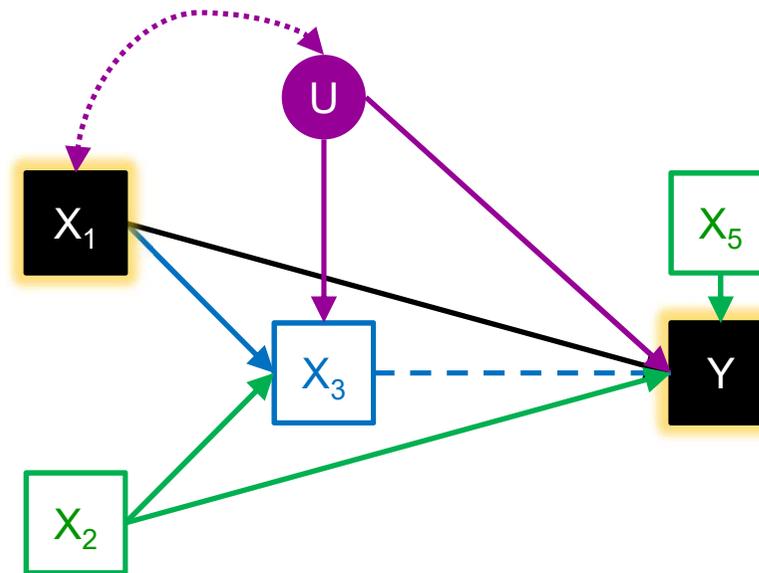
*Hernández-Díaz et al. Am J Epidemiol. 2006;164(11):1115-20*





- In our model, we recognised the coefficient on  $X_1$  was **NOT** the total causal effect of  $X_1$  on  $Y$ , but we might think it gave a meaningful (direct) causal effect.
- In fact, the apparent effect will be **collider-biased** by all other mutual causes of  $X_3$  and  $Y$  (**mediator-outcome confounders**)

$$Y \sim X_3 + X_1 + X_2 + X_5$$

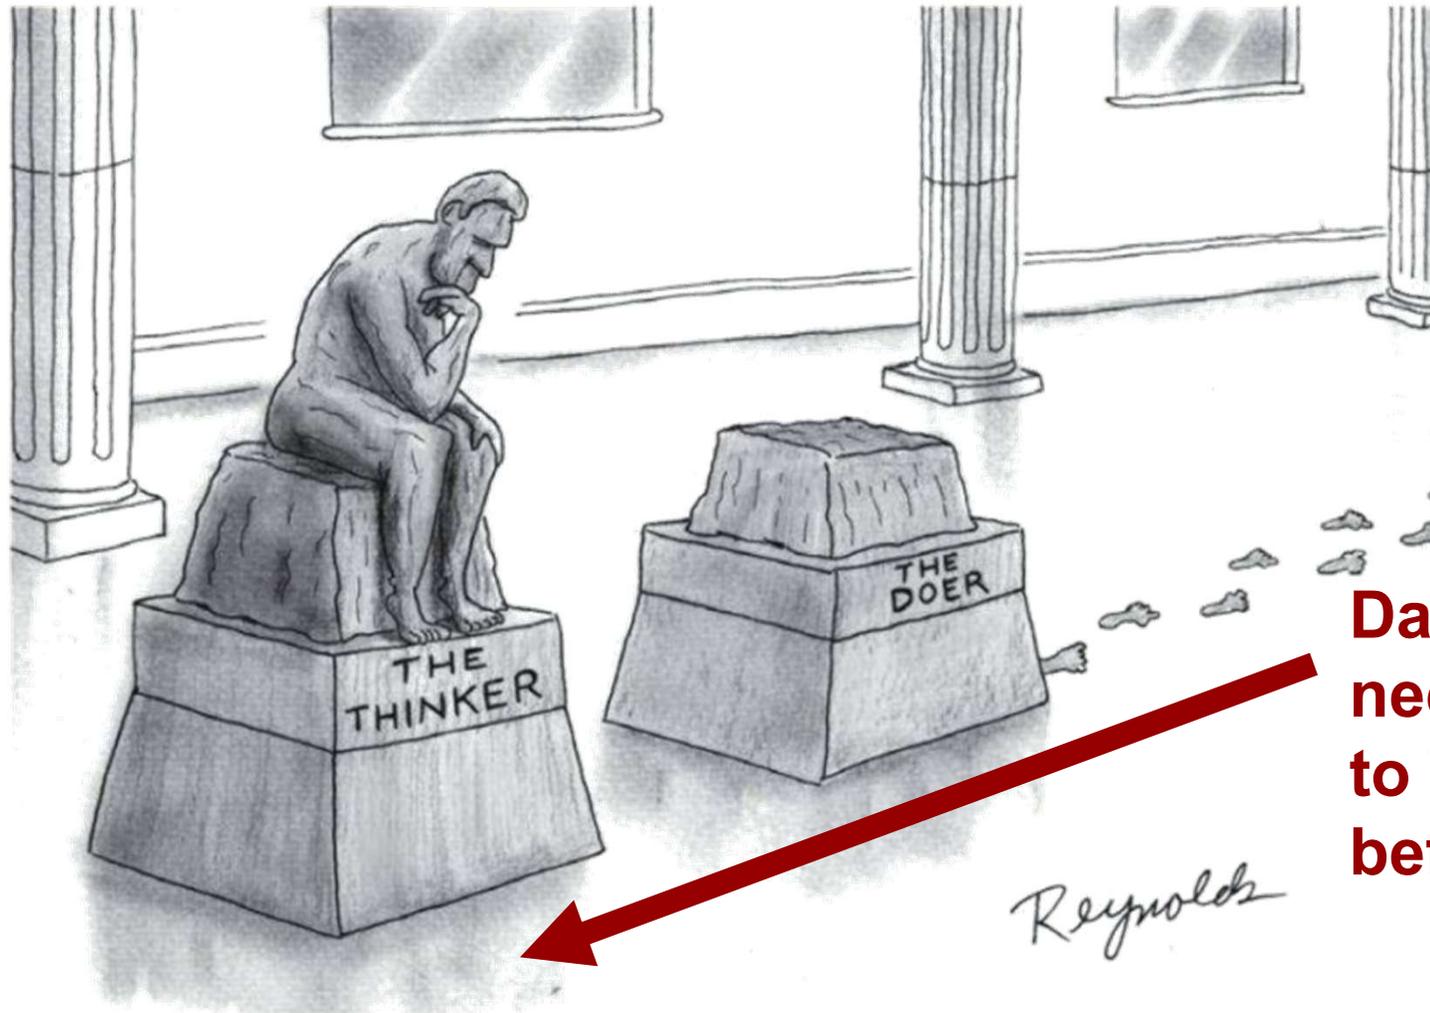


- This can **reverse** the apparent direction of effect (**reversal paradox**)!

- **Causal inference** is one of the three core tasks of data science
- Unlike **description** and **prediction**, it requires external contextual knowledge
- Traditional data-driven approaches have a poor record at causal inference
- Joining the **causal revolution** means embracing our **causal ambitions**, & recognising the need for a completely different approach
- This starts with explicitly **identifying** what we want to know before **estimating**
- **Causal diagrams** are a simple & transparent aid to causal reasoning, & are helpful for understanding many issues in non-experimental data



- Causal inference requires external knowledge → careful thought



**Data scientists  
need more time  
to THINK  
before they DO!**

### ▪ **Differences between prediction and causal inference**

- Arnold KF *et al.* Generalised linear models for prognosis and intervention: Theory, practice, and implications for machine learning. ArXiv 2019. <https://arxiv.org/abs/1906.01461v2>

### ▪ **Using DAGs in your research**

- Tennant PWG *et al.* Use of directed acyclic graphs (DAGs) in applied health research: review and recommendations. MedArXiv 2019. <https://doi.org/10.1101/2019.12.20.19015511>

### ▪ **Causal Inference Book**

- Hernán MA & Robins JM. Causal Inference: What If. Boca Raton: Chapman & Hall/CRC, 2020. <https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/>