Causal Inference

Sid Mangalik

Say the line, Bart!

Correlation does not imply causation.

Correlation

Correlation is a useful measure of the degree to which two variables are related.

Correlation does not tell us about the fundamental mechanisms that underlie the relationship between the two variables.

Correlation **can** coexist with causation, and in fact usually does.

1.00	0.64	0.71	0.65	0.77
0.64	1.00	0.58	0.64	0.62
0.71	0.58	1.00	0.66	0.72
0.65	0.64	0.66	1.00	0.60
0.77	0.62	0.72	0.60	1.00

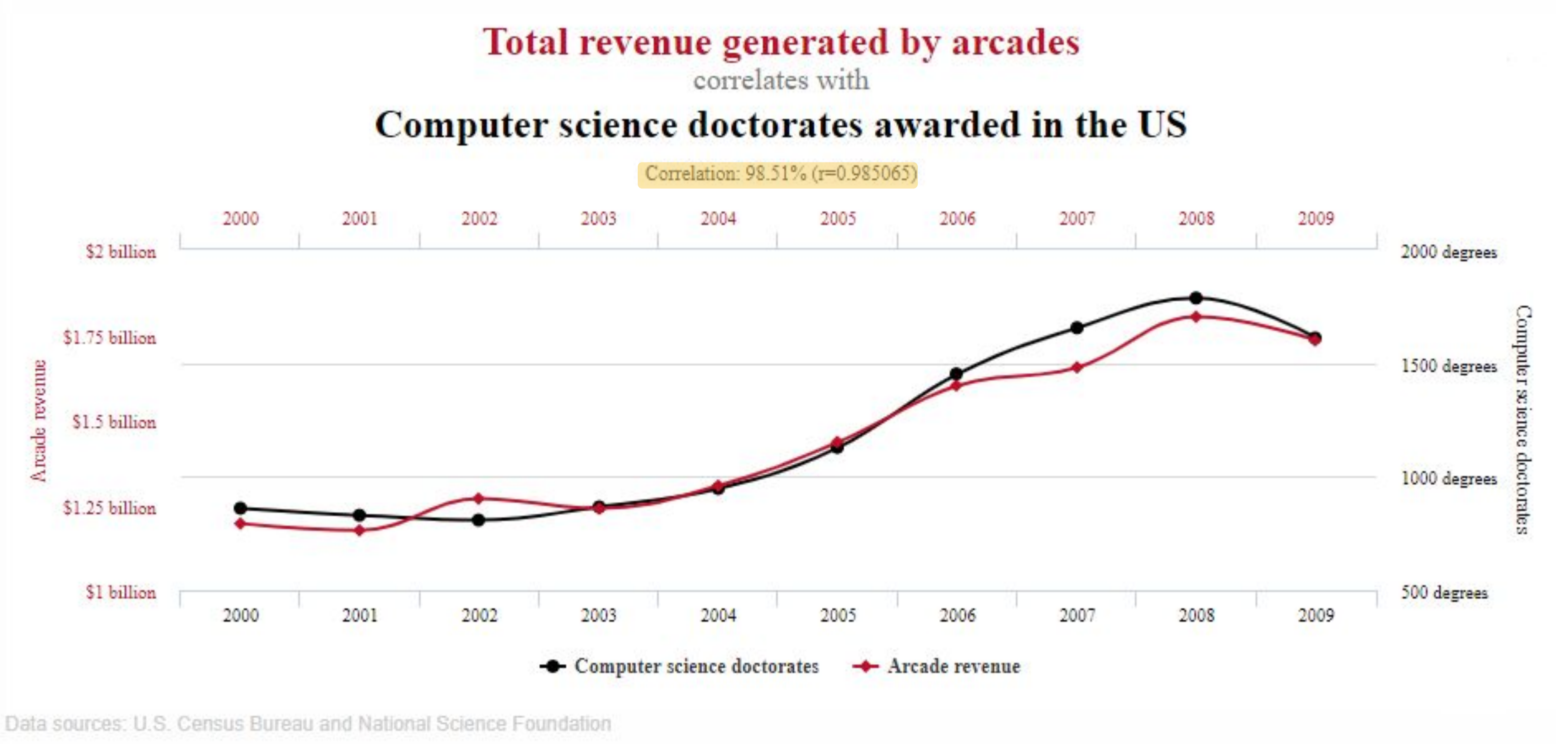

Spurious Correlations

tylervigen.com/spurious-correlations

Causation and Etiology

On a philosophical level, for Causes *A* and Effects *B*:

1. A and B must be contiguous in space and time.
2. A must precede B
3. A and B must be observed to interact (constant conjunction)

Causation?

On a philosophical level, for Causes *A* and Effects *B*:

1. A and B must be contiguous in space and time.
2. A must precede B
3. A and B must be observed to interact (constant conjunction)

Extremely pedantic philosophers, like David Hume, bring the skeptical argument that **all** causality is inferred

Try not to lose too much sleep over this

Types of Causal Relationships

Causal Chain

$$A \rightarrow B \rightarrow C$$

Causal Homeostasis

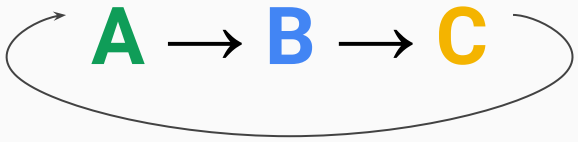

Common-Cause

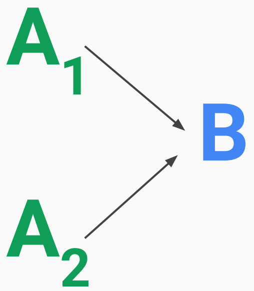

Common-Effect

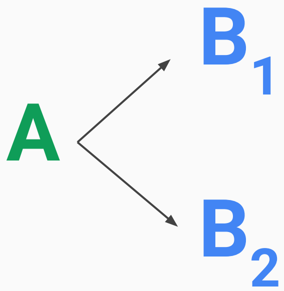

Statistical Inference vs Causal Inference

Statistical

- “X is associated with Y” means that changing X is probable to change the distribution of Y
- X and Y will be associated and vice versa.
- Can be used to generate predictions

Causal

- “X causes Y ” means that changing the value of X will change the distribution of Y
- X and Y will be associated but the reverse is not always true
- Can be used to generate explanations

Binary Causal Effect Size

Treatment (A)	Outcome (B)	Outcome when not treated (C_0)	Outcome when treated (C_1)
0	0	0	Counterfactual
0	0	0	Counterfactual
0	0	0	Counterfactual
0	0	0	Counterfactual
1	1	Counterfactual	1
1	1	Counterfactual	1
1	1	Counterfactual	1
1	1	Counterfactual	1

On the left we see a sample dataset for a treatment A and an outcome B, alongside potential outcomes.

- A = 1 if subject takes Vitamin C = 1, does 0 if not
- B = 1 if subject is healthy = 1, 0 is sick

Note that we do not observe the outcome of the subject if they had the opposite treatment status. We call this a **counterfactual**.

Average causal/treatment effect = θ
$\theta = E(C_1) - E(C_0)$

Association = $\alpha = E(B|A=1) - E(B|A=0)$

Fill in Data

Treatment (A)	Outcome (B)	Outcome when not treated (C_0)	Outcome when treated (C_1)
0	0	0	0*
0	0	0	0*
0	0	0	0*
0	0	0	0*
1	1	1*	1
1	1	1*	1
1	1	1*	1
1	1	1*	1

The data has now been filled in, what can we learn?

We leave a * for the counterfactuals

Association = ***α*** = E(B|A=1) − E(B|A=0)
(1+1+1+1)/4 - (0+0+0+0)/4
= 1

Wow! Vitamin-C consumption has such a great correlation with health, we should tell people about this!

MORE ORANGES!

Treatment (A)	Outcome (B)	Outcome when not treated (C_0)	Outcome when treated (C_1)
0	0	0	0*
1	0	0	0*
1	0	0	0*
1	0	0	0*
1	1	1*	1
1	1	1*	1
1	1	1*	1
1	1	1*	1

Let's tell everyone to take more vitamin-C and then collect some data, shown on the left

Association = α = $E(B|A=1) - E(B|A=0)$

$(0+0+0+1+1+1+1)/7 - (0)/1$

$= 4/7$

What happened??

Average causal/treatment effect = θ

$\theta = E(C_1) - E(C_0)$

$(0+0+0+0+1+1+1+1)/8 - (0+0+0+0+1+1+1+1)/8$

$= 0$

If we had examined the causal effect we would have known better than to trust oranges.

Association != Causation

Treatment (A)	Outcome (B)	Outcome when not treated (C_0)	Outcome when treated (C_1)
0	0	0	0*
0	0	0	0*
0	0	0	0*
0	0	0	0*
1	1	1*	1
1	1	1*	1
1	1	1*	1
1	1	1*	1

Let's examine how association and causal effect compare.

Average causal/treatment effect $= \theta$

$\theta = E(C_1) - E(C_0)$

$(0+0+0+0+1+1+1+1)/8 - (0+0+0+0+1+1+1+1)/8$

$= 0$

Association = $\boldsymbol{\alpha}$ $= E(B|A=1) - E(B|A=0)$

$(1+1+1+1)/4 - (0+0+0+0)/4$

$= 1$

Thus: θ != $\boldsymbol{\alpha}$

Continuous Treatments

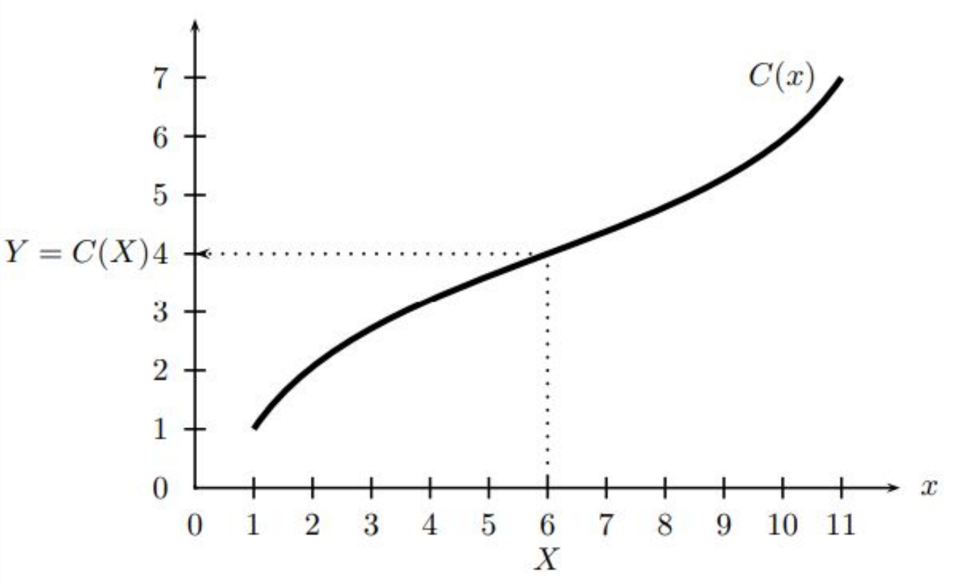

Instead of C_0 and C_1 we have a **counterfactual function** C(x)

Causal Regression Function = θ

$\theta(x) = E(C(x))$

$= \theta$

Association Regression Function = $r(x)$

$r(x) = E(Y|X=x)$

Example

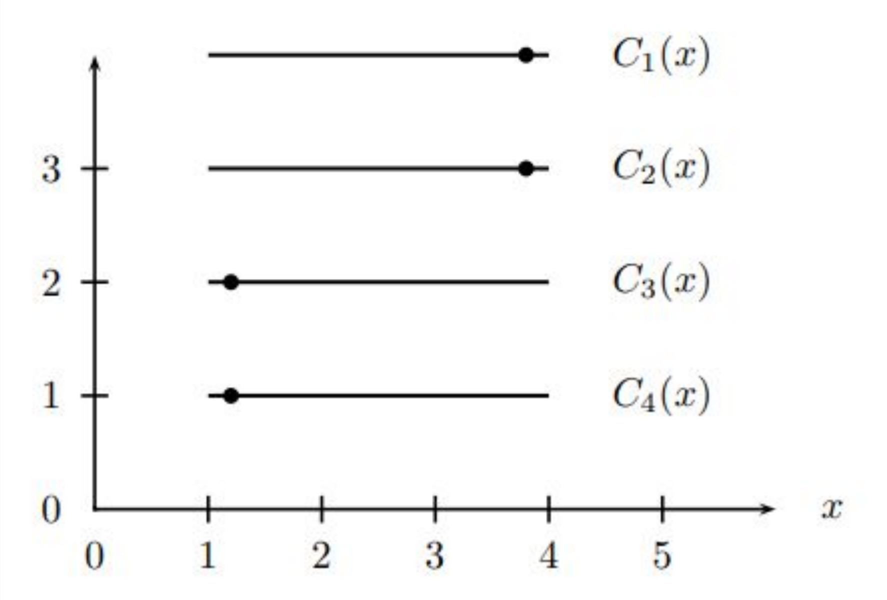

Here we see the counterfactual functions for four subjects.

Here lets treat:
X as the subject’s consumption of vitamin C
Y as # times subject is sick in a year

The dots represent their observed X values:
X_1, X_2, X_3, X_4

Causal Regression

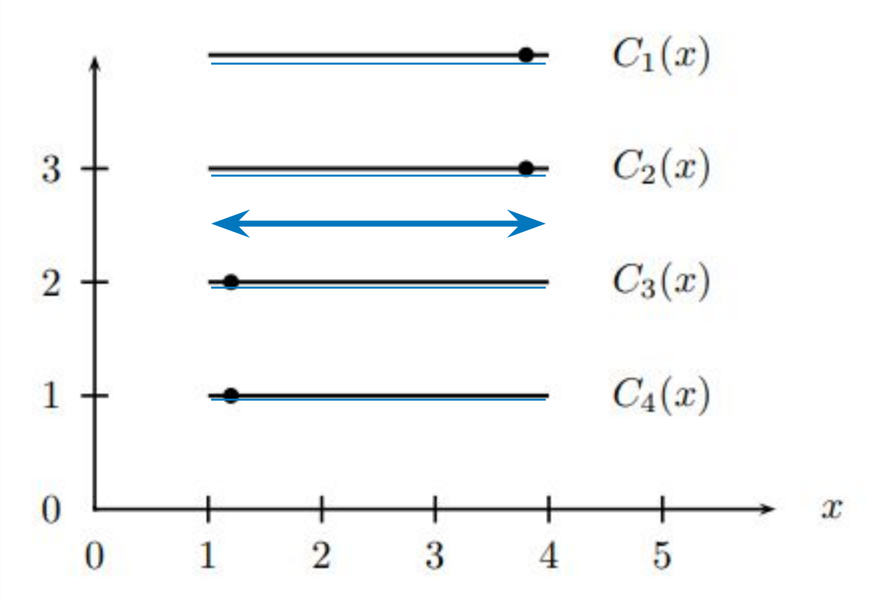

Since $C_i(x)$ is constant over x for all i, there is no causal effect

Causal Regression Function = θ(x)

θ(x) = E(C(x))

$\theta(x) = (C_1(x) + C_2(x) + C_3(x) + C_4(x)) / 4$

Since θ(x) is constant changing the dose x will not change anyone's outcome

Association Regression

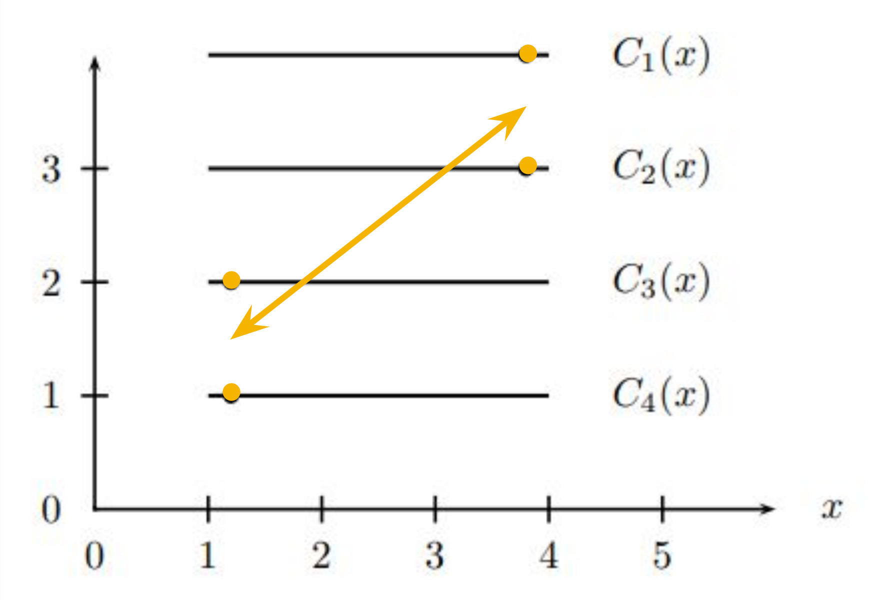

Instead of C_0 and C_1 we have a **counterfactual function** $C(x)$

Association Regression Function = r(x)

$r(x) = E(Y|X=x)$

$Y_1 = C_1(X_1)$

$Y_2 = C_2(X_2)$

$Y_3 = C_3(X_3)$

$Y_4 = C_4(X_4)$

Takeaways

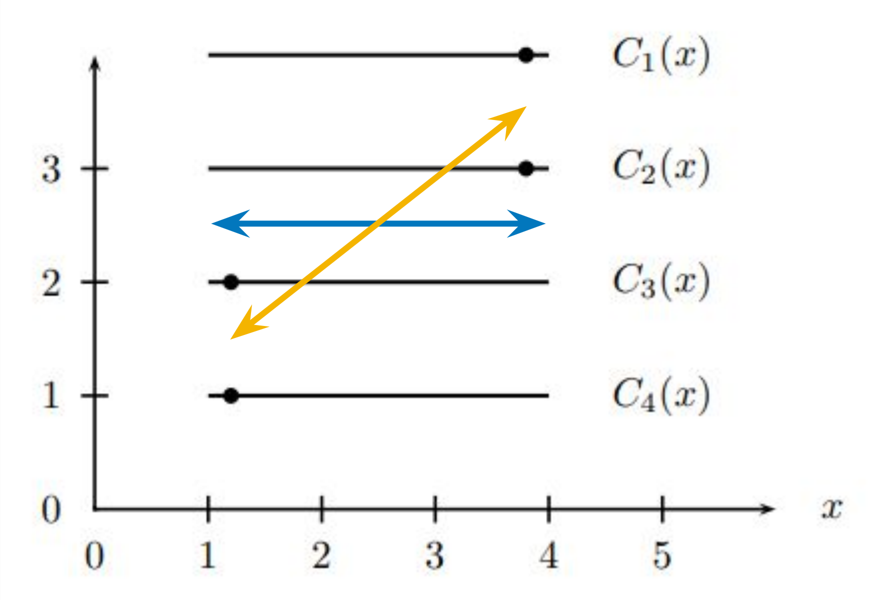

There is no causal effect since $C_i(x)$ is constant for all i.

Although there is no causal effect, there is an association since the regression curve $r(x)$ is not constant.

This was an example where $\theta(x)$ is constant but $r(x)$ is not constant

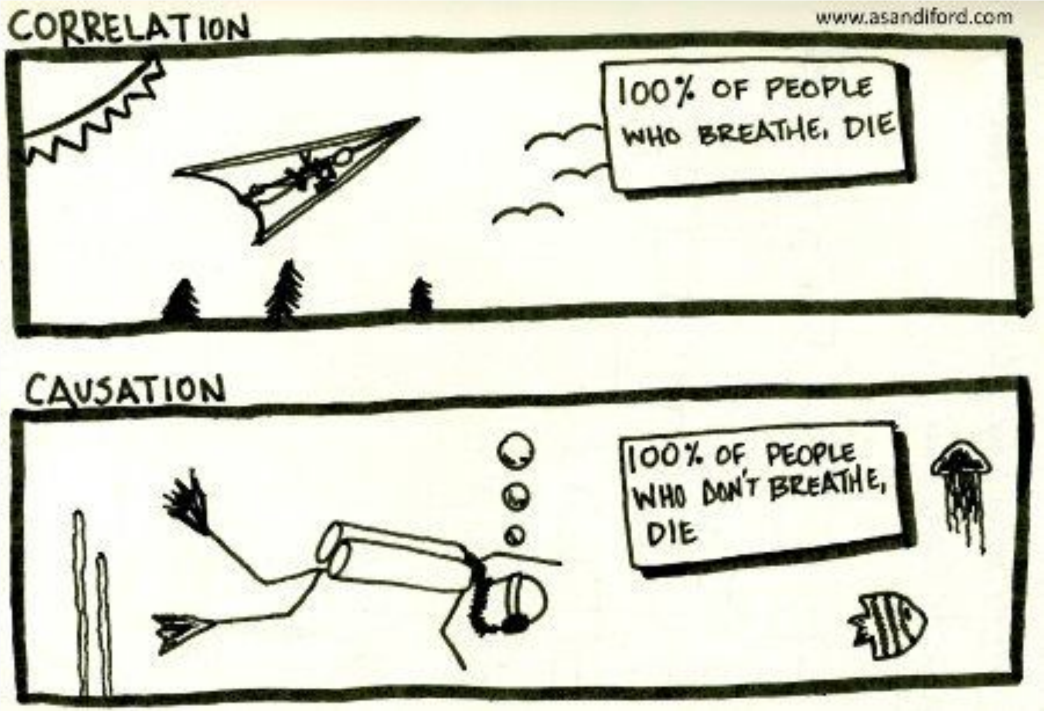

(Worth a small discussion)

Asaniford
asandiford.com/comic/correlation-causation/

Inferring Causal Relationships

A natural way to understand Causal Inference studies is with medical research.

- Experiments can be conducted
- Outcomes are quantifiable
  - Stability of findings can be verified
- Cohorts can be given randomized treatments
  - Ethical restrictions apply
- Participants can be matched on constants
- Confounders can be mitigated and controlled
- Timing of interventions is known

Randomized Control Trial

Sometimes considered a subset of A/B Testing.

The gold standard for experimental design.

Group A (treatment) is as similar as possible to Group B (control). The only difference being a treatment that is randomly assigned to Group A.

As many confounders as possible should be controlled for, such as age/gender/ethnicity.

Working Example - Free Lunch

Bill makers want to propose a policy to offer free lunches to students in US middle schools.

In order to make their case they hire a research team to determine if offering free lunches increases the academic performance of 7th graders.

Randomized Control Trial (RCT)

In order to test the hypothesis that free lunches increase school performance the researchers approach a school to do a study.

Example: Students are **randomly assigned** to two groups to minimize bias. **One group gets a free lunch**. The **other group does not receive free lunch**. The results of the two groups are compared to see if **the group that received lunch experienced an expected effect**.

Concepts Built Into RCT

- **Random assignment** - Ideally the random split balances out the effects of confounding variables
- **Treatment** - Getting a free lunch
- **Treatment group** - The students who received the treatment
- **Control group** - The students who did not receive the treatment
- **Effect size** - Amount of effect attributed to the treatment

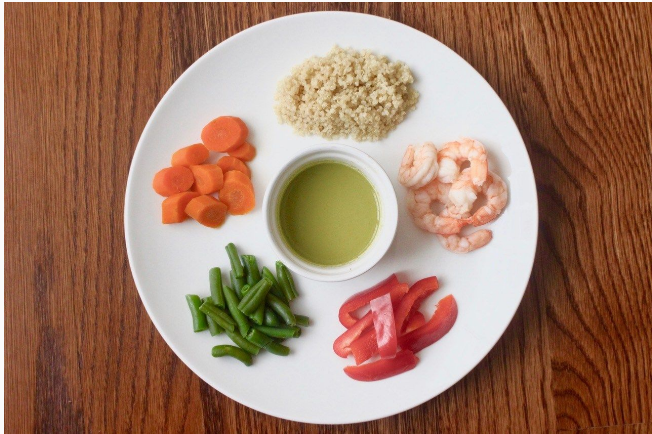

5 Minute Break

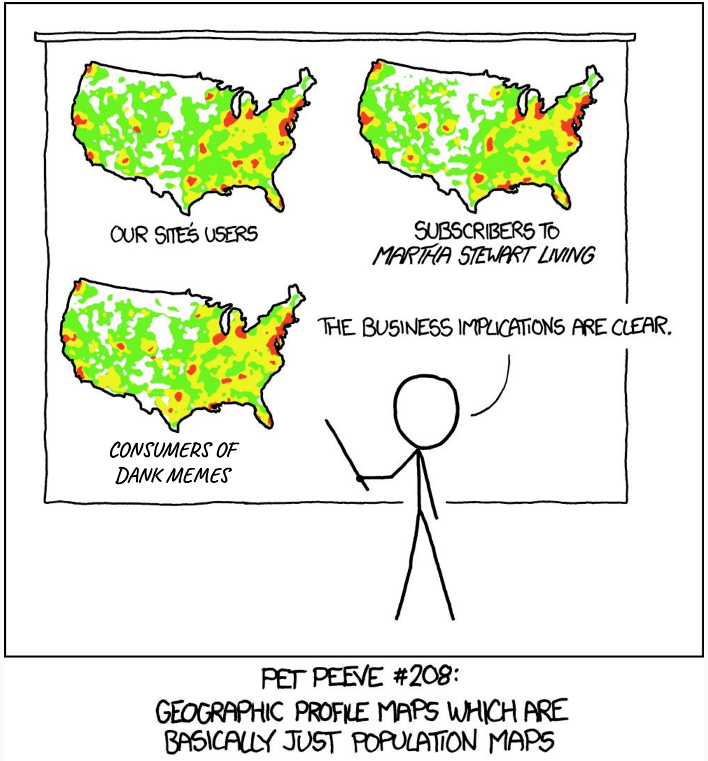

XKCD (1138)

Experimental Designs

Useful **when treatment cannot be withheld** (logistically or ethically)

These designs are popular natural experiment frameworks in econometrics and journalism

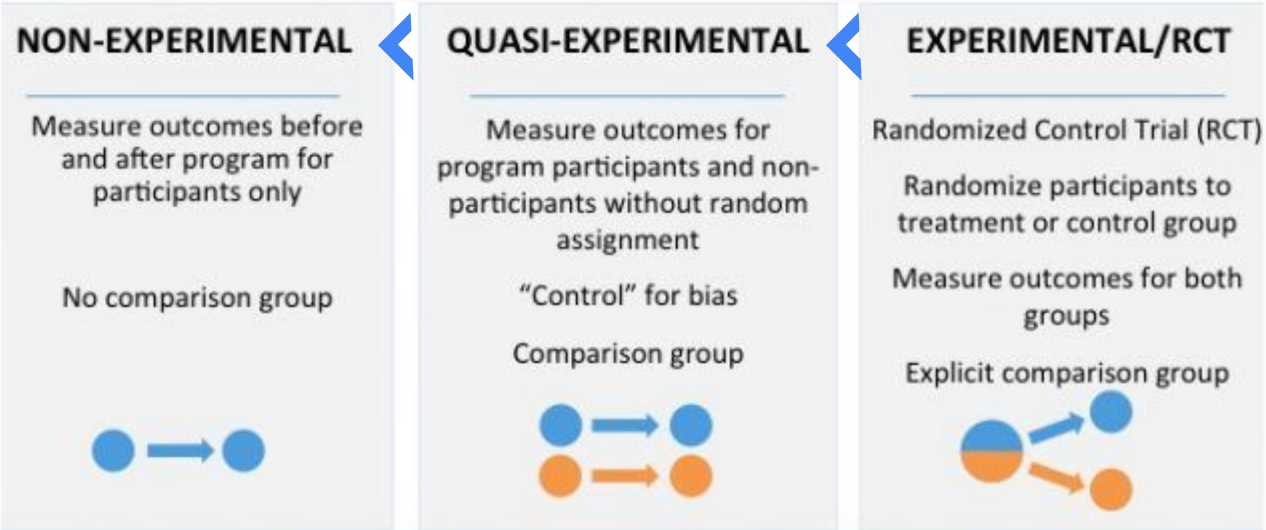

Observational/Non-Experimental Design

In an observational study treatments are **not** randomized.

In these types of studies the amount of treatment is **not independent** of the outcome.

In order to conduct this type of study we will need to be rigorously limit the study to a particular population

- Limit **confounding variables** that affect a subject's likeness to choose the treatment
- Such that choosing the treatment is effectively random within population
- Example: group subjects by age, gender, race, education

Quasi-Experimental Design

Often an RCT is just not possible, so we do our best to simulate an experimental design.

The key missing element is being able to randomize the treatment for the subjects.

These designs are conducted by looking retrospectively at data to tease out causal relations

The Matching Problem

Since quasi-experimental designs do not give us an explicit control and treatment group we must create our own post-hoc.

This is not a simple problem since high quality matches are a large component of the validity of a quasi-experimental causal inference study.

Matching techniques:

- **Propensity Score** - Estimate the effect of a treatment by accounting for the covariates that predict receiving the treatment in the first place.
- **Inverse Propensity Weighting** - Reweight the sample population by 1/propensity for better representativeness of the true population. (Considered a replacement for Propensity Matching)
- **Vector Similarity** - Represent all participants with vectors containing the attributes you wish to match them on, then split the group on clusters defined by any distance metric.

Difference in Difference

Study the effect of an **intervention** on an outcome.

We find the **counterfactual**, which is where we predict the treatment group would have performed without the intervention.

We compare the counterfactual to the actual outcome to determine the size of the intervention effect.

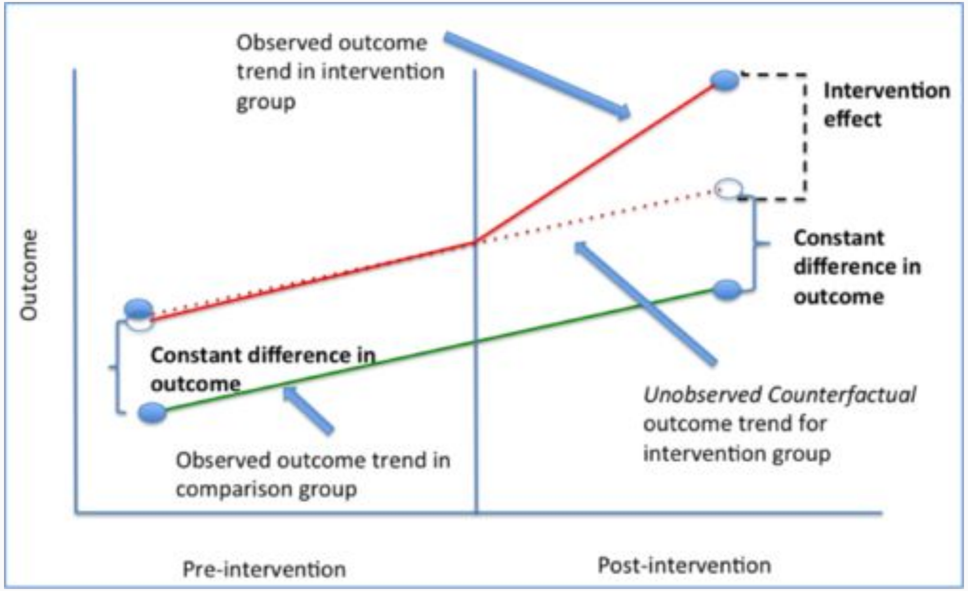

DID / Difference in Difference / Diff in Diff

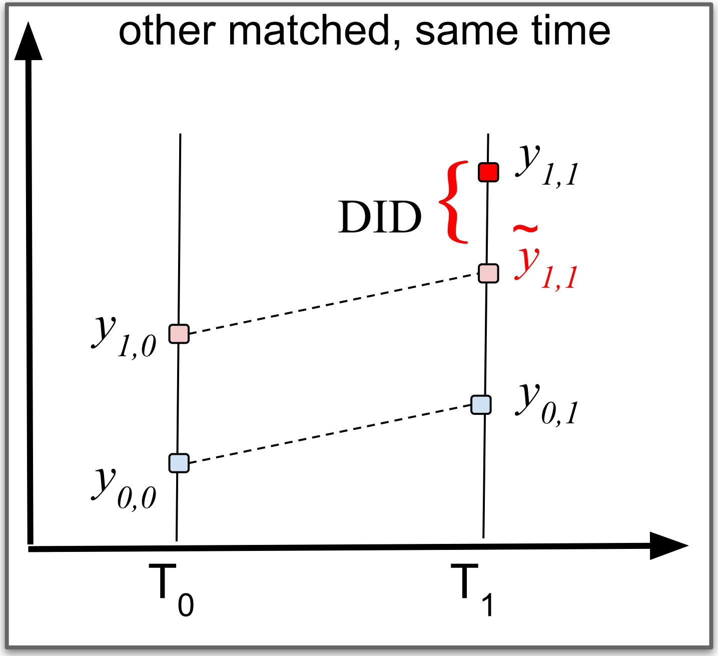

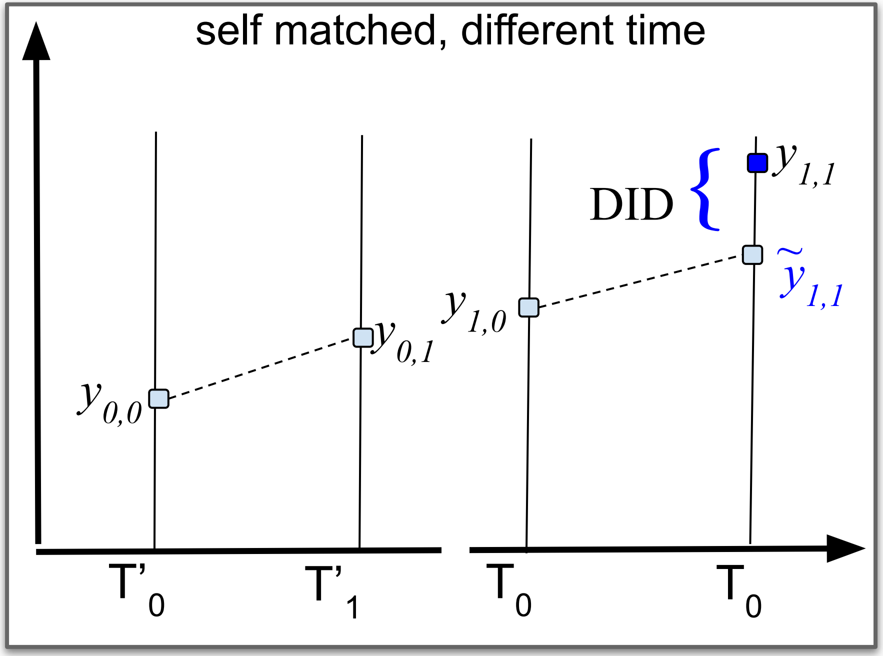

Other matched, Same time

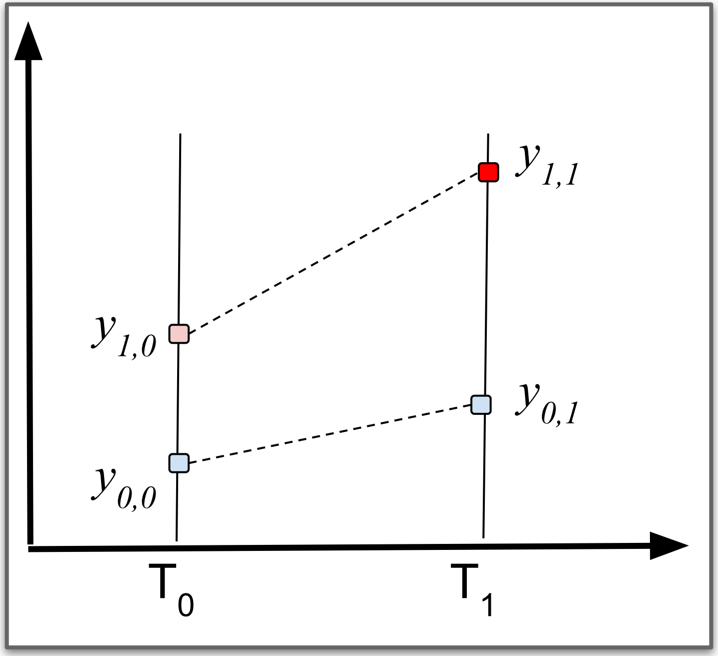

Example:
y_0 and y_1 are comparable school districts

Between T_0 and T_1 students are given free lunches at y_1

Record the average grades of 7th graders before and after the intervention

Other matched, Same time

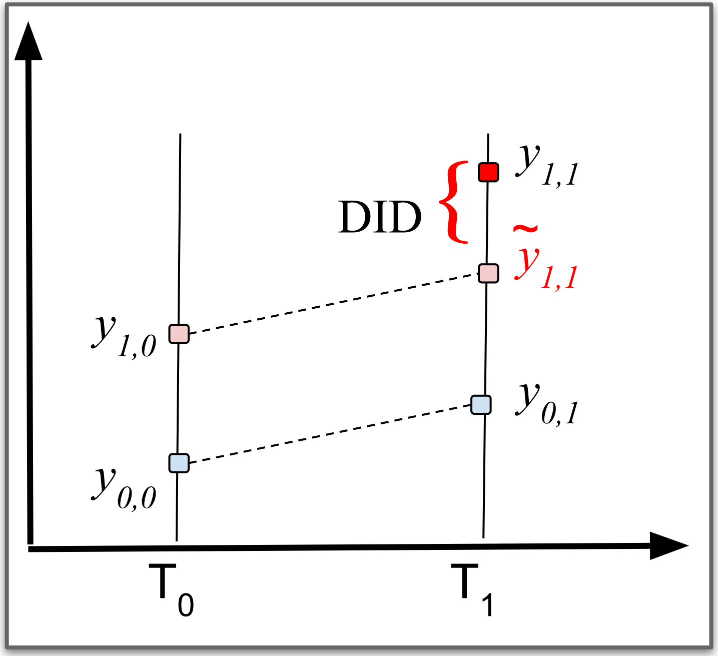

Point $\tilde{y}_{1,1}$ represents the counterfactual or “expected” outcome for y (Grade performance without lunches)

This gap, DID, is the difference in differences or *treatment effect* (Grade increase caused by lunches)

$$\tilde{y}_{1,1} = y_{1,0} + (y_{0,1} - y_{0,0})$$

$$DID = y_{1,1} - \tilde{y}_{1,1}$$

Self matched, Different time

Example:
y_0 and y_1 are both the same school

Between T_0 and T_1 the free lunch policy is enacted

Record the average grades of 7th graders for one time period with and without free lunches

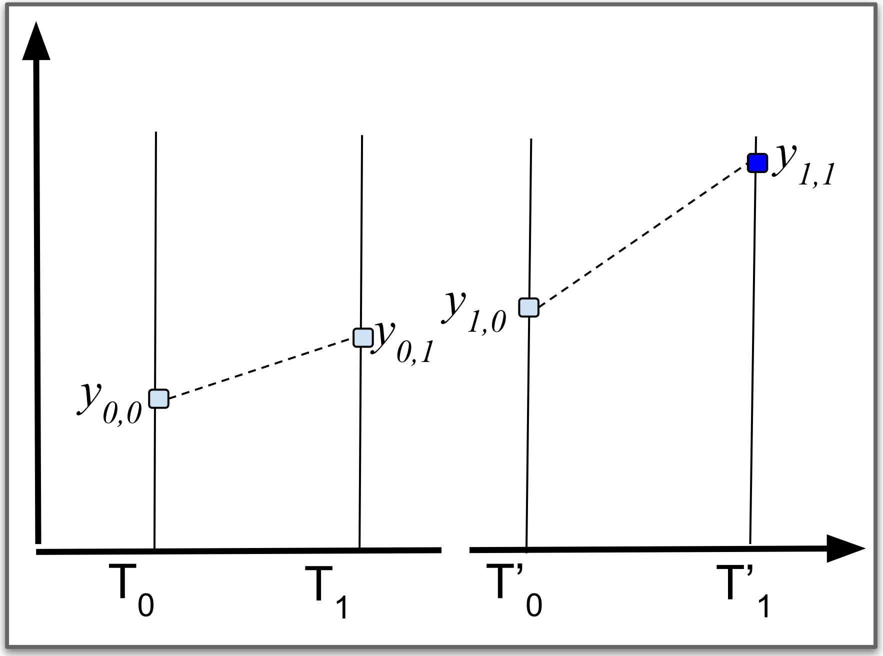

Self matched, Different time

Point $\tilde{y}_{1,1}$ represents the counterfactual or “expected” outcome for y_1 (Grades without free lunch)

This gap, DID, is the difference in differences or *treatment effect* (Grade change caused by lunches)

$$\tilde{y}_{1,1} = y_{1,0} + (y_{0,1} - y_{0,0})$$

$$DID = y_{1,1} - \tilde{y}_{1,1}$$

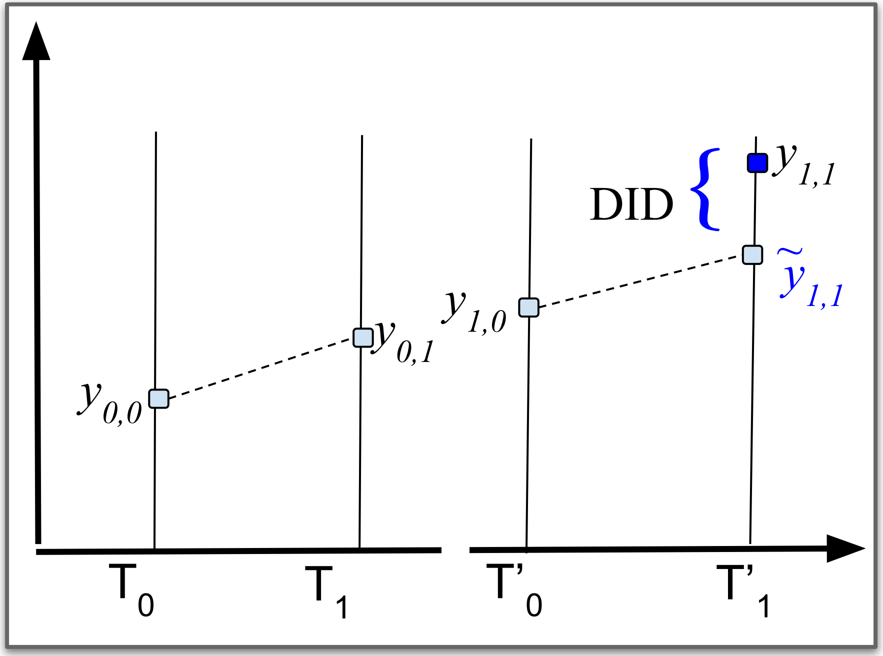

Interrupted Time Series / Regression Discontinuity

Assign a cutoff/threshold above/below which an intervention is defined.

Compare observations lying closely on either side of the threshold to estimate the average treatment effect

Using only this method does not allow for true causal inference, since it does not reject causal effects from potentially confounding variables.

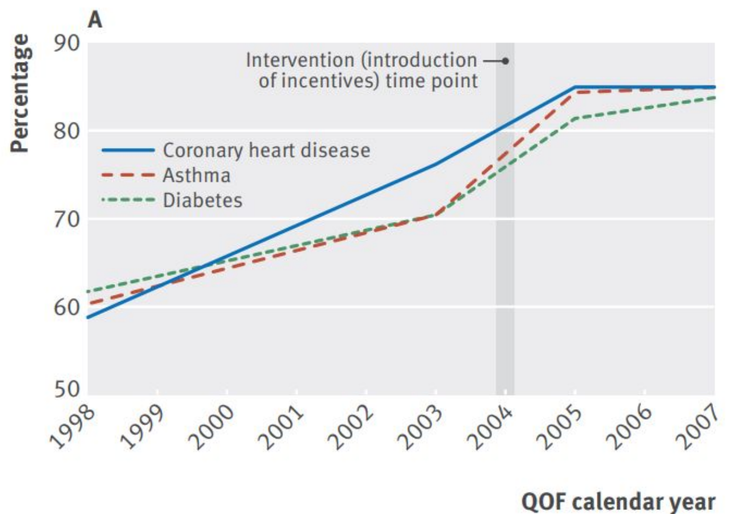

Regression based quasi-experimental approach when randomisation is not an option: interrupted time series analysis

Evangelos Kontopantelis,[1,2] Tim Doran,[3] David A Springate,[2,4] Iain Buchan,[1] David Reeves[2,4]

Google's CausalImpact

Extremely useful for ad performance measurement and pricing experiments.

Uses a structural Bayesian time-series model to estimate how the response metric might have evolved after the intervention if the intervention had not occurred.

Assumes:

- The outcome time-series can be explained in terms of a set of control time series that were themselves not affected by the intervention.
- The relation between treated series and control series is assumed to be stable during the post-intervention period.

github.com/google/CausalImpact

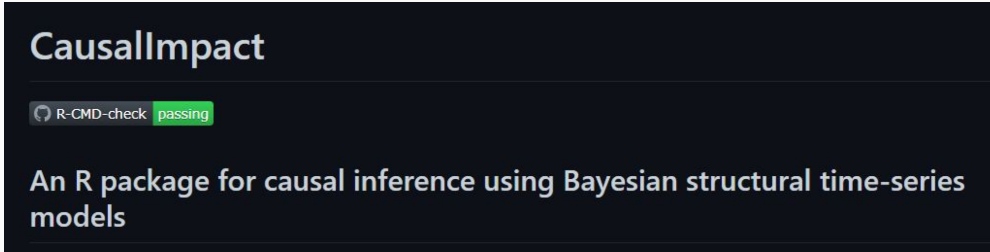

CausalImpact Plotted results

```
plot(impact)
```

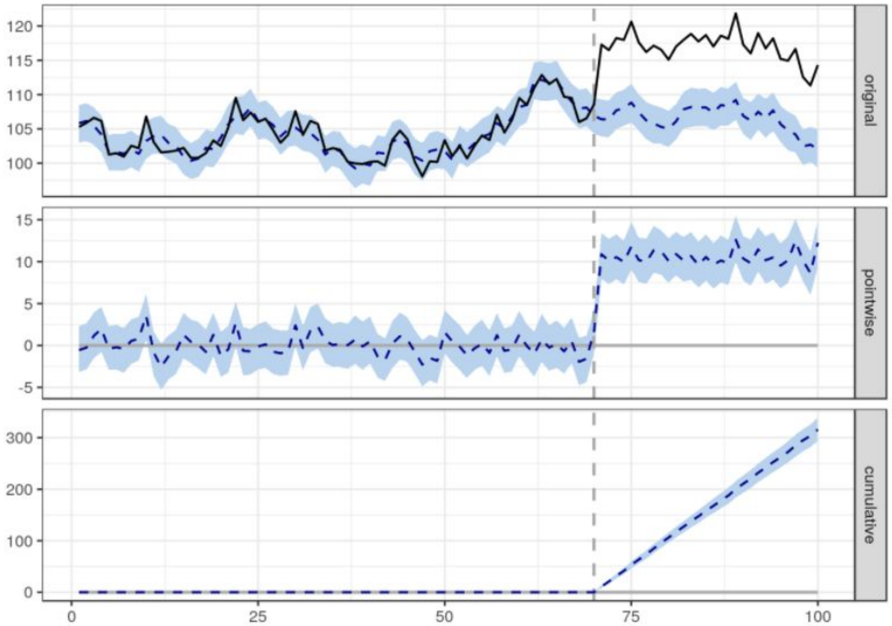

Original: the data and a counterfactual prediction for the post-treatment period.

Pointwise: the difference between observed data and counterfactual predictions. This is the pointwise causal effect, as estimated by the model.

Cumulative: adds up the pointwise contributions from the second panel, resulting in a plot of the cumulative effect of the intervention.

Potential Issues with Causal Inference

- It is a form of inference
- An effect might be better explained by another cause
- Without proper statistical controls, findings may be the result of chance
- Difficult to perform, there is significant debate amongst scientists about proper methodology
- The ghost of Hume still haunts us

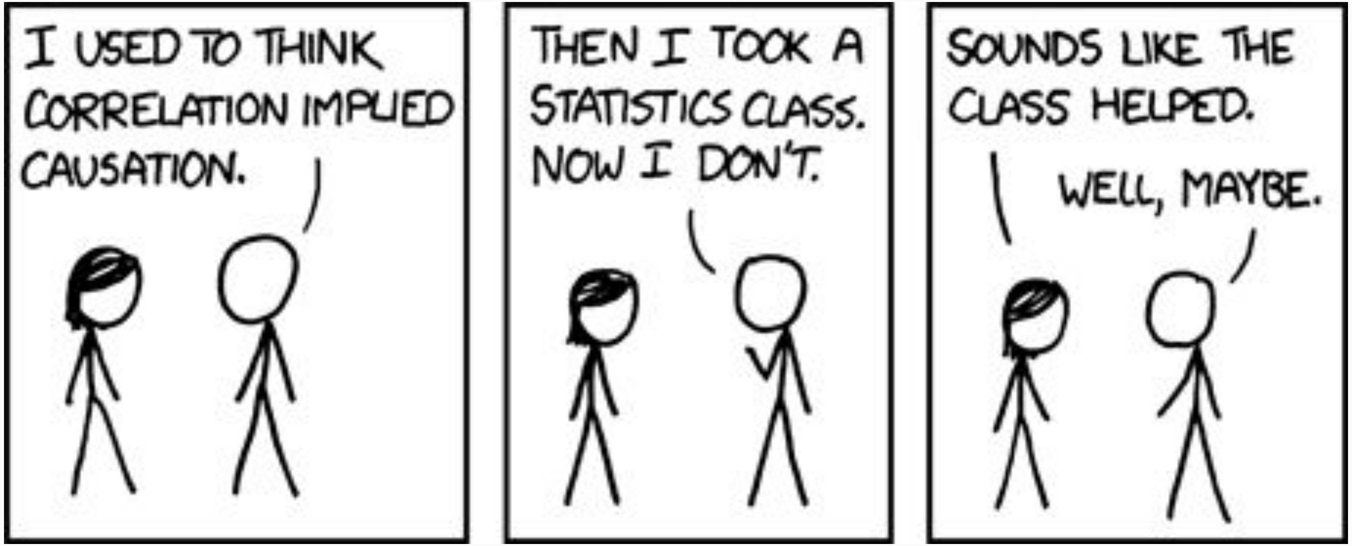

Correlation doesn't imply causation, but it does waggle its eyebrows suggestively and gesture furtively while mouthing “look over there”

XKCD (552)