Section 5

Lecture 2

Main message from Lecture 1

- The interpretation of (conditional) associations is subtle (Think about our death penalty example and GRE example).
- We introduced counterfactual random variables: say, Y^a, the outcome if we intervened to set A to a.
- We saw that individual level effects ($Y_i^{a=1} - Y_i^{a=0}$) are essentially never observable.

The previous session was fairly conceptual.
Now we are going to be more explicit.

Plan for today

- Clarify causal questions
- Identification
  - Consistency, Exchangeability, Positivity
- Do identification proofs (Motivate estimators)
- Consider a conditional randomized trial
- Observational study
- Effect modification
- Interaction

...

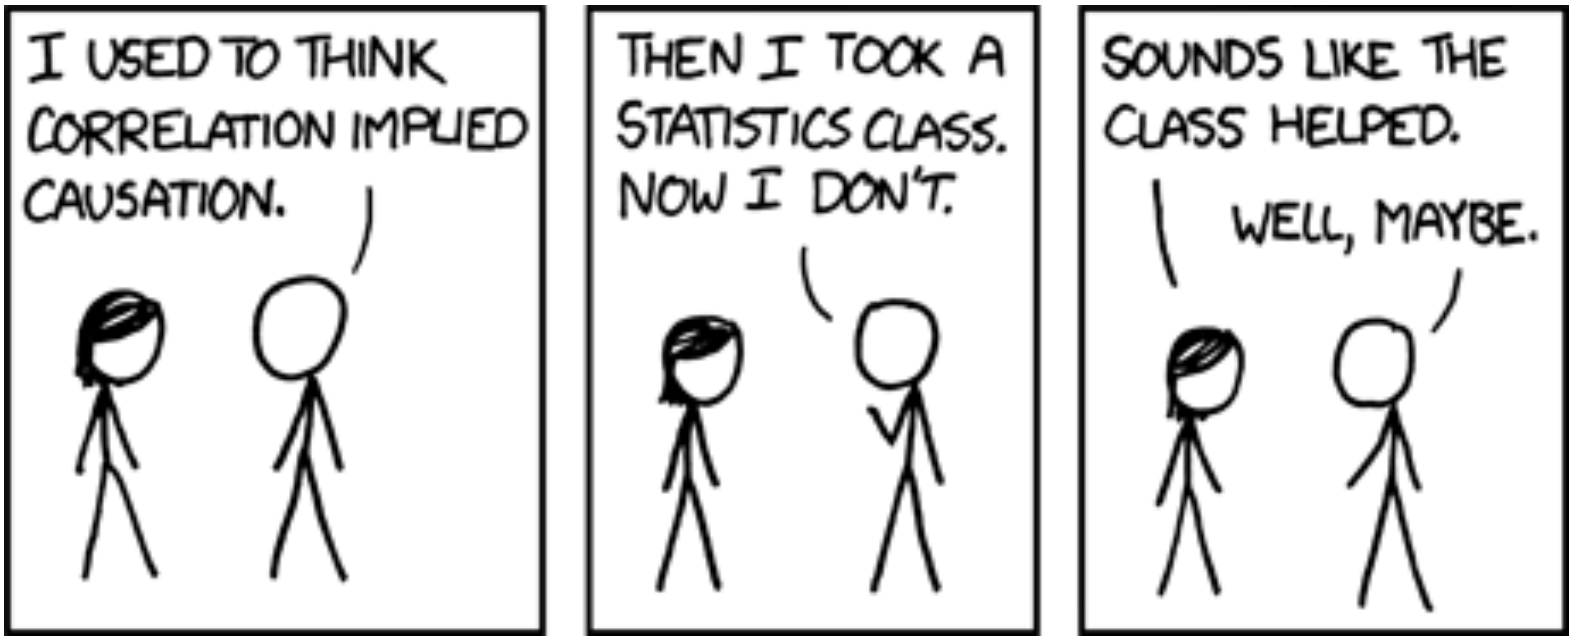

Instead of saying association is not causation, we will formalize *when* an association can be interpreted causally.

Questions

The following slides just give some examples on questions of different types.

- Descriptive / predictive:
  - "Is this patient at high risk of developing complications during surgery?"

  Causal:
  - "Which type of anaesthetic should this patient receive to reduce the risk of complications during surgery?"
  - "How does the amount of anaesthetic affect the risk of complications during surgery?"
  - "What can be done to reduce the risk of complications during surgery for an average / a particular type of patient?"

Questions

- Descriptive / predictive:
  - "Which type of client will buy which kind of product?"

  Causal:
  - "Should advert be at the top or bottom of website to increase the probability of viewing product?"
  - "How does the size of advert affect the probability of viewing product?"
  - "How can I get a client to buy my product?"

Questions

- Descriptive / predictive:
  - "Who is most likely to become long-term unemployed?"

  Causal:
  - "Will a minimum wage legislation increase the unemployment rate of a country?"
  - "How does the size of advert affect the probability of viewing product?"
  - "What can be done to prevent someone from becoming unemployed?"

Effect contrasts

- Additive effect: $\mathbb{E}(Y^{a=1}) - \mathbb{E}(Y^{a=0}) = \mathbb{E}(Y^{a=1} - Y^{a=0})$.
  The additive effect is an average over individual level causal effects. These are marginal quantities.
- Relative effect: $\frac{\mathbb{E}(Y^{a=1})}{\mathbb{E}(Y^{a=0})} \neq \mathbb{E}\left(\frac{Y^{a=1}}{Y^{a=0}}\right)$.
  The relative effect is not an average over individual level causal effects.

Causal effects in the population

More generally, we can consider population causal effects[8]:

Definition (Population causal effect)

A population causal effect can be defined as a contrast of any functional of the distributions of counterfactual outcomes in the same (sub)population under different interventions.

- For example $\text{VAR}(Y^{a=1}) - \text{VAR}(Y^{a=0})$.
  Remember that we cannot identify $\text{VAR}(Y^{a=1} - Y^{a=0})$.
- I will often say *causal effect* when I talk about *average causal effect*.

[8] Hernan and Robins, *Causal inference: What if?*

Section 6

Randomisation

Example conditions that ensure identification of causal effects

Suppose that the following 3 conditions hold:

1. $Y^a \perp\!\!\!\perp A, \forall a \in \{0,1\}$ (exchangeability[9]).
2. $P(A = a) > 0 \ \forall a \in \{0,1\}$ (positivity[10]).
3. $Y^a = Y$ for every unit with $A = a$ (consistency[11]).
   that is, $Y = I(A = 0)Y^{a=0} + I(A = 1)Y^{a=1}$.

From (1)-(3), $\mathbb{E}(Y^a) = \mathbb{E}(Y \mid A = a)$.
That is, we have *identified* $\mathbb{E}(Y^a)$ as a functional of observed data.
Assumptions (1)-(3) are external to the data, but – importantly – they hold by design in a perfectly executed experiment.
Just to be clear: the counterfactual independence $Y^a \perp\!\!\!\perp A, \forall a \in \{0,1\}$ does NOT imply the factual independence $Y \perp\!\!\!\perp A$.

[9] Also called ignorability.
[10] Also called overlap. Note that this is a feature of the distribution, not the sample.
[11] Similar to the condition SUTVA: Stable Unit Treatment Value Assumption.

A simple example of estimation of causal effects

Because $\mathbb{E}(Y^a) = \mathbb{E}(Y \mid A = a)$, the simple difference-in-means estimator,

$$\hat{\delta} = \frac{1}{n_1} \sum_{A_i=1} Y_i - \frac{1}{n_0} \sum_{A_i=0} Y_i, \;\; n_a = \sum_{i=1}^{n} I(A = a),$$

is an unbiased estimator of the average (additive) causal effect of A in a randomised experiment.

We will discuss estimation in more detail later in this course.

Conditional randomisation

- Let $L \in \{0, 1\}$
  In the heart transplant example, let $L = 1$ if the individual is critically ill, 0 otherwise.
- Suppose A is *conditionally* randomised as a function of L such that $P(A = 1 \mid L = 0) = p_0$ and $P(A = 1 \mid L = 1) = p_1$,
  where $p_0 \neq p_1$ and $p_0, p_1 \in (0, 1)$.

How do we identify $\mathbb{E}(Y^a)$?

Illustrative *conditional* experiment (trial) on heart transplant

	L	A	Y
Rheia	0	0	0
Kronos	0	0	1
Demeter	0	0	0
Hades	0	0	0
Hestia	0	1	0
Poseidon	0	1	0
Hera	0	1	0
Zeus	0	1	1
Artemis	1	0	1
Apollo	1	0	1
Leto	1	0	0
Ares	1	1	1
Athena	1	1	1
Hephaestus	1	1	1
Aphrodite	1	1	1
Cyclope	1	1	1
Persephone	1	1	1
Hermes	1	1	0
Hebe	1	1	0
Dionysus	1	1	0

In this conditional randomised trial $p_0 = 0.5, p_1 = 0.75$
Compute an estimator based on the numbers above, and you will find that $\hat{\mathbb{E}}(Y^{a=1}) - \hat{\mathbb{E}}(Y^{a=0}) = 0$.

Identification in a conditional randomised experiment

A is *conditionally* randomised such that $P(A = 1 \mid L = 0) = p_0$ and $P(A = 1 \mid L = 1) = p_1$, where $p_0 \neq p_1$ and $p_0, p_1 \in (0, 1)$.

$Y^a \not\perp\!\!\!\perp A, \forall a \in \{0, 1\}$ (Exchangeability from Slide 50 may fail), but

1. $Y^a \perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$ (Exchangeability).
2. $P(A = a \mid L = l) > 0 \ \forall a \in \{0, 1\}, \forall l$ s.t. $P(L = l) > 0$. (positivity).
3. $Y^a = Y$ for every unit with $A = a$ (consistency).

When 1-3 hold, then

$$\mathbb{E}(Y^a) = \sum_l \mathbb{E}(Y \mid L = l, A = a) P(L = l).$$

These conditions hold by design in a conditional randomised experiment.

Identification in a conditional randomised experiment

Proof.

$$\begin{aligned}
\mathbb{E}(Y^a) &= \sum_l \mathbb{E}(Y^a \mid L = l) P(L = l) \\
&= \sum_l \mathbb{E}(Y^a \mid L = l, A = a) P(L = l) \quad \text{(positivity and exchangeability)} \\
&= \sum_l \mathbb{E}(Y \mid L = l, A = a) P(L = l). \quad \text{(consistency)}
\end{aligned}$$

□

We say that the 3rd line is an identification formula for $\mathbb{E}(Y^a)$.
This is a special case of a so-called G-formula (or truncation formula)[12].

[12] James M Robins. "A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect". In: *Mathematical modelling* 7.9-12 (1986), pp. 1393–1512.

Alternative weighted identification formula

$$\begin{aligned}\mathbb{E}(Y^a) &= \sum_l \mathbb{E}(Y \mid L = l, A = a) \Pr(L = l) \\ &= \mathbb{E}\left[\frac{I(A = a)}{\pi(A \mid L)} Y\right].\end{aligned}$$

where $\pi(a \mid l) = P(A = a \mid L = l)$.
Why bother with equivalent expressions?
Because they motivate different *estimators*.

Proof of IPW

Proof.

$$
\begin{aligned}
&\mathbb{E}\left[\frac{I(A=a)}{\pi(A \mid L)} Y\right] \\
=&\mathbb{E}\left[\frac{I(A=a)}{P(A=a \mid L)} Y^a\right] \quad \text{(consistency and positivity)} \\
=&\mathbb{E}\left[\mathbb{E}\left\{\frac{I(A=a)}{P(A=a \mid L)} Y^a \mid L\right\}\right] \\
=&\mathbb{E}\left\{\mathbb{E}\left[\frac{I(A=a)}{P(A=a \mid L)} \mid L\right] \mathbb{E}\left[Y^a \mid L\right]\right\} \quad \text{(exchangeability)} \\
=&\mathbb{E}\left\{\mathbb{E}\left[Y^a \mid L\right]\right\} = \mathbb{E}\left[Y^a\right].
\end{aligned}
$$

What if questions can be assessed in experiments

Later in this course we will discuss experiments and design,[13]

... but experiments are often not available because they are

- impractical,
- expensive,
- time consuming,
- unethical,

... and experiments *are often not be perfectly executed*.

So, what do we do? Decisions in real life have to be made...

[13] David Roxbee Cox and Nancy Reid. *The theory of the design of experiments*. CRC Press, 2000.

Section 7

Effect modification and conditional effects

Effect modification

Definition (Effect modification)

We say that V is a modifier of the effect of A on Y when the average causal effect of A on Y varies across levels of V.

Since the average causal effect can be be defined on different scales, effect modification depends on the scale.

Definition (Qualitative effect modification)

We say there is qualitative effect modification if the average causal effects if there exist v, v' such that the effect given $V = v$ are in the opposite direction of effects given $V = v'$.

Note that:

- V may or may not be equal to L.
- *"Effect heterogeneity across strata of V"* is often used interchangeably with *"effect modification by V"*.

Why bother with effect modification?

- So far we have focused on average causal effects.
- However, effects will often be different in different subpopulations of individuals (between men and women, Greek and Romans etc.).
- It is often of practical interest to target future intervention to subsets of the full population (If the treatment has a positive effect in men and negative effect in women, we would like to give men and women different treatments).
- Some individuals will have different benefit of treatment than others (towards precision medicine and personalised medicine...).
- Later in the course, we will also see that this is important when we are going to *generalize* (or *transport*) effects from a study to other populations (for example, we have done an experiment in a selected population, and now we want to make decisions in another population. Therefore our question is how the intervention will work in this other population).

Illustrative experiment (trial) on heart transplant.

We may be interested in effects conditional on a baseline variable V.

	V	Y^0	Y^1
Rheia	1	0	1
Demeter	1	0	0
Hestia	1	0	0
Hera	1	0	0
Artemis	1	1	1
Leto	1	0	1
Athena	1	1	1
Aphrodite	1	0	1
Persephone	1	1	1
Hebe	1	1	0
Kronos	0	1	0
Hades	0	0	0
Poseidon	0	1	0
Zeus	0	0	1
Apollo	0	1	0
Ares	0	1	1
Hephaestus	0	0	1
Cyclope	0	0	1
Hermes	0	1	0
Dionysus	0	1	0

Here, $V = 1$ if woman, $V = 0$ if man.

Concrete example

Suppose that:

- $\mathbb{E}(Y^{a=1} \mid V = 1) = 0.6 > \mathbb{E}(Y^{a=0} \mid V = 1) = 0.4.$
- $\mathbb{E}(Y^{a=1} \mid V = 0) = 0.4 < \mathbb{E}(Y^{a=0} \mid V = 0) = 0.6.$

We conclude that there is qualitative effect modification by gender.
Treatment $A = 1$

- increases mortality in women, but
- reduces mortality in men.

Let $P(V = 0) = 0.5$. Then, the average causal effect
$\mathbb{E}(Y^{a=1}) - \mathbb{E}(Y^{a=0}) = 0.$

Identification of effects modified by V.

For simplicity suppose that V and L are disjoint.

1. $Y^a \perp\!\!\!\perp A \mid L, V, \forall a \in \{0, 1\}$ (Exchangeability).
2. $P(A = a \mid L = l, V = v) > 0 \; \forall a \in \{0, 1\}, \forall l \in \mathcal{L}, \forall v \in \mathcal{V}$ (Positivity).
3. $Y^a = Y$ for every unit with $A = a$ (Consistency).

How to identify effect modification

- Strategy for identification:
  1. Stratify by V.
  2. Identify the effect within each level $V = v$.
- For example, in a conditional randomised trial, an identification formula for the average causal effect of $A = a$ in the stratum defined by $V = v$ is

$$\mathbb{E}(Y^a \mid V = v) = \sum_l \mathbb{E}(Y \mid L = l, V = v, A = a) P(L = l, V = v).$$

Romans vs Greeks.

Consider a conditional randomised study on Heart transplant, and let V indicate whether the individual is Roman ($V = 0$) or Greek ($V = 1$)[14]

Stratum $V = 0$

	L	A	Y
Cybele	0	0	0
Saturn	0	0	1
Ceres	0	0	0
Pluto	0	0	0
Vesta	0	1	0
Neptune	0	1	0
Juno	0	1	1
Jupiter	0	1	1
Diana	1	0	0
Phoebus	1	0	1
Latona	1	0	0
Mars	1	1	1
Minerva	1	1	1
Vulcan	1	1	1
Venus	1	1	1
Seneca	1	1	1
Proserpina	1	1	1
Mercury	1	1	0
Juventas	1	1	0
Bacchus	1	1	0

[14] Hernan and Robins, *Causal inference: What if?*

Concrete example from Slide 63

Suppose that:

- $\mathbb{E}(Y^{a=1}) = 0.55$ and $\mathbb{E}(Y^{a=0}) = 0.40$.
- $\mathbb{E}(Y^{a=1} \mid V = 1) = 0.5 = \mathbb{E}(Y^{a=0} \mid V = 1) = 0.5$ (in Greeks).
- $\mathbb{E}(Y^{a=1} \mid V = 0) = 0.6 > \mathbb{E}(Y^{a=0} \mid V = 0) = 0.3$. (in Romans)

We conclude that there is qualitative effect modification by nationality.

Section 8

Interaction is different from effect modification

Terminology

Remember the difference between the following terms:

- **Estimand** (a parameter of interest, e.g. $\mathbb{E}(Y^a)$).
- **Estimator** (an algorithm / function / rule that can be applied to data).
- **Estimate** (an output from applying the estimator to data).

We talk about bias of an estimator with respect to an estimand.
That is, the term *bias* (biased / unbiased) is always defined with respect to an estimand.

Interaction requires multiple interventions

- Consider two binary treatments $A \in \{0, 1\}$ and $E \in \{0, 1\}$.
  For example, chemotherapy and surgery.
- For each individual we can imagine 4 potential outcomes, that is, $Y^{a=0,e=0}$, $Y^{a=1,e=0}$, $Y^{a=0,e=1}$ and $Y^{a=1,e=1}$.

Definition (Additive interaction)

There is additive interaction if

$$\mathbb{E}(Y^{a=0,e=0}) - \mathbb{E}(Y^{a=1,e=0}) \neq \mathbb{E}(Y^{a=0,e=1}) - \mathbb{E}(Y^{a=1,e=1}).$$

Additive interaction is symmetric wrt. A and E,

$$\begin{aligned} & \mathbb{E}(Y^{a=0,e=0}) - \mathbb{E}(Y^{a=1,e=0}) \neq \mathbb{E}(Y^{a=0,e=1}) - \mathbb{E}(Y^{a=1,e=1}) \\ \implies & \mathbb{E}(Y^{a=0,e=0}) - \mathbb{E}(Y^{a=0,e=1}) \neq \mathbb{E}(Y^{a=1,e=0}) - \mathbb{E}(Y^{a=1,e=1}). \end{aligned}$$

`Remember that, unlike interactions, effect heterogeneity did only involve interventions on` A, `not the modifier` V.

Multiplicative interaction

Definition (Multiplicative interaction)

There is multiplicative interaction if

$$\frac{\mathbb{E}(Y^{a=0,e=0})}{\mathbb{E}(Y^{a=1,e=0})} \neq \frac{\mathbb{E}(Y^{a=0,e=1})}{\mathbb{E}(Y^{a=1,e=1})}.$$

Example: Interaction

- A chemotherapy, E radiation therapy, Y being cured of cancer.
- Interaction question: Is there interaction between the effect of receiving both A chemotherapy and E radiation therapy?

	E = 0	E = 1
A = 0	0.02	0.05
A = 1	0.04	0.10

Table 1: Experiment where A and E are randomised[15]

[15] Tyler J VanderWeele and Mirjam J Knol. "A tutorial on interaction". In: *Epidemiologic Methods* 3.1 (2014), pp. 33–72.

Conceptual example

- Let Y indicate being cured. There is additive interaction because

$$\mathbb{E}(Y^{a=0,e=0}) - \mathbb{E}(Y^{a=1,e=0}) \neq \mathbb{E}(Y^{a=0,e=1}) - \mathbb{E}(Y^{a=1,e=1})$$
$$0.2 - 0.4 \neq 0.05 - 0.10,$$

  but no multiplicative interaction because $\frac{0.2}{0.4} = \frac{0.5}{0.10}$.
- Suppose we had 100 versions of drug E after A was randomly assigned. Then, we would expect to cure 3 additional persons if we used all of the drug supply among those with $A = 0$. However, we would expect to cure 6 additional people if we used all the supply among those with $A = 1$.

Interaction and its relation to factorial experiments[16]

- How would you conduct an experiment to evaluate interactions between varialbes?
- We need a *factorial* design.
  - Each treatment (A and E in our example) has different levels ($A, E \in \{0, 1\}$ in our example). A factorial design consists of an equal number of replicates of all possible combinations of the levels of the factors.
  - In our Example from Slide 73, there are $2^2 = 4$ different combination of treatment levels.

[16] Cox and Reid, *The theory of the design of experiments*.

Interaction summary

- Just to say that there is an interaction on some scale is relatively uninteresting; all it means is that both exposures have some effect on the outcome.
- Additive interaction is more relevant to public health.

Section 9

Causal inference from observational data

Observational data

Definition (Observational data)

A sample from a population where the treatment (exposure) is not under the control of the researcher.

That is, the treatment (exposure) of interest is not randomly assigned.

Following Robins[18], let's be slightly more abstract

- A dataset is a string of numbers.
- These data represent empirical measurements (for example, for each study subject, a series of treatments and outcomes).
- In an analysis, calculations are performed on these numbers.
- Based on the calculations, causal inference is drawn.
- "*Since the numerical strings and the computer algorithm applied to them are well-defined mathematical objects, it would be important to provide formal mathematical definitions for the English sentences expressing the investigator's causal inferences that agree well with our informal intuitive understanding*"[17].

[17] James M Robins. "Addendum to "a new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect"". In: *Computers & Mathematics with Applications* 14.9-12 (1987), pp. 923–945.

[18] Robins, "A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect".

Observational studies

- In an observational study, treatment is not assigned according to randomisation, but according to someone's choice, for example the patient, the costumer or the medical doctor.
- People who choose to take treatment may be different from those who choose not to take treatment, in the sense that they have different risk of the outcome even before the decision is made. $Y^a \not\perp\!\!\!\perp A, \forall a \in \{0, 1\}$.
- The question is, can we find the characteristics L, which are associated with treatment and the outcome such that $Y^a \perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$?
  In other words, exchangeability does no longer hold by design, but can we *assume* that it holds? What do we need to include in L for this to hold?
- Yet, humans have learned a lot from *observations*, and many scientific studies are not experiments. We have learned about *effects of* smoking, global warming, evolution, astrophysics etc.

Same data, different story

Suppose the data (identical numbers to the slide 54) were from an *observational* study (now A is not randomly assigned), where the doctors tended to provide transplants ($A = 1$) to those with most severe disease ($L = 1$)

	L	A	Y
Rheia	0	0	0
Kronos	0	0	1
Demeter	0	0	0
Hades	0	0	0
Hestia	0	1	0
Poseidon	0	1	0
Hera	0	1	0
Zeus	0	1	1
Artemis	1	0	1
Apollo	1	0	1
Leto	1	0	0
Ares	1	1	1
Athena	1	1	1
Hephaestus	1	1	1
Aphrodite	1	1	1
Cyclope	1	1	1
Persephone	1	1	1
Hermes	1	1	0
Hebe	1	1	0
Dionysus	1	1	0

Example continues

- Suppose first that L is the *only* outcome predictor unequally distributed between those with $A = 1$ and $A = 0$. Then $Y^a \perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$.
- Now, suppose that the doctors not only used L to make treatment decisions, but also used smoking status, $S \in \{0, 1\}$, where smoking status is an outcome predictor. Then, $Y^a \not\perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$.
- Thus, $Y^a \perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$ may not hold in observational studies.
- Suppose the investigators did not measure S. Can they use the observed data to evaluate whether $Y^a \not\perp\!\!\!\perp A \mid L, \forall a \in \{0, 1\}$ holds? The answer is *no*.

Finally, we need consistency

- Well-defined interventions.
- How do we reason about exchangeability for a treatment A that is ill-defined?
- Suppose now that our exposure (treatment) is obesity A.
- How can we identify common causes of obesity L and the outcome mortality Y?
- And does positivity hold? There can be some Ls (say, related to exercise) for which nobody is obese.
- The target trial where obesity is the exposure seems to involve unreasonable interventions. How can we instantly make people non-obese? By forcing them to exercise? By doing surgery? By diet? All of these interventions may have different effects.

The target trial

- We have argued that contrast between average counterfactual outcomes under different treatments are often of substantial interest.
- We have also clarified that conducting an experiment guarantees identification of a causal effect. However, conducting an experiment is not always feasible.
- For each causal effect of interest, we can conceptualize a (hypothetical) randomised experiment to quantify it. This hypothetical randomised experiment is called the **target experiment** or **target trial**.
- Being explicit about specifying the target trial forces us to be explicit about the causal question of interest. We ask the question: "What randomised experiment are you trying to emulate?"

Specification of the target trial

To make a causal question practically interesting and useful, it is important to clarify the following, which is part of the specification of the target trial:

- Target population (eligibility criteria).
- Interventions (the treatment strategies).
- Outcome (what is the outcome and when will the outcome be measured)
- Statistical analysis (application of estimators and their statistical properties).

Also clarifies how the claims made can be falsified in the future (in principle), by conducting the target trial. This fits with a positivist (Popperian) view of science.