Estimation of causal effects

Arvid Sjölander

Department of Medical Epidemiology and Biostatistics
Karolinska Institutet

Causal Inference for Epidemiological Research

Ideal data

- Let Y_x be the outcome that we would observe, for a given subject, if the subject potentially received exposure level x
  - Y_0 is the outcome when not exposed
  - Y_1 is the outcome when exposed
- Y_0 and Y_1 are referred to as **potential outcomes**
- Ideally - **and very unrealistically** - we could observe both potential outcomes for any given subject

Subject	Y_0	Y_1
August	0	1
Selma	0	0
Fjodor	1	1

Subject-specific causal effects

subject	Y_0	Y_1
August	0	1
Selma	0	0
Fjodor	1	1

- X has a causal effect on Y, for a given subject, if the potential outcomes Y_0 and Y_1 differ for this subject
  - for August, the exposure has an effect: $Y_0 \neq Y_1$
  - for Selma and Fjodor, the exposure has no effect; $Y_0 = Y_1$

Observed data

- August is exposed ($X = 1$). Thus, for August
  - Y_1 is observed and equal to the factual outcome Y
  - Y_0 is unobserved, or **counterfactual**
- Selma and Fjodor are unexposed ($X = 0$). Thus, for Selma and Fjodor
  - Y_0 is observed and equal to the factual outcome Y
  - Y_1 is unobserved, or **counterfactual**

subject	X	Y	Y_0	Y_1
August	1	1	?	1
Selma	0	0	0	?
Fjodor	0	1	1	?

A fundamental problem of causation

- It is very difficult to say whether the exposure causes the outcome for a specific subject
  - because we cannot observe the same subject under two exposure levels simultaneously
- Fortunately, it is much easier to justify causal claims on population levels
  - e.g. 'if nobody would smoke, then the incidence of liver cancer would be 15% less then if everybody would smoke'

Population causal effects

- $p(Y_x = 1)$ is the proportion of subjects that would develop the outcome, if **everybody** would receive exposure level x
  - the probability of the outcome if everybody would receive x
- X has a population causal effect on Y if

$$p(Y_0 = 1) \neq p(Y_1 = 1)$$

- X has no population causal effect on Y if

$$p(Y_0 = 1) = p(Y_1 = 1)$$

How to estimate population causal effects

Everybody unexposed

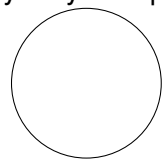

Everybody exposed

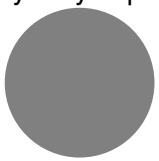

$$p(Y_0 = 1) \text{ vs } p(Y_1 = 1)$$

- Direct computation of population causal effects requires comparing
  - the whole population being unexposed, with
  - the whole population being exposed
- But just like for any given subject, we cannot in general observe the whole population under two exposure levels
- *How can we estimate population causal effects?*

A naive solution

Factually unexposed

Factually exposed

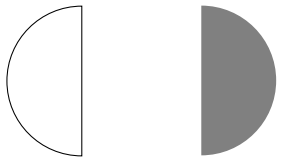

$$p(Y = 1|X = 0) \text{ vs } p(Y = 1|X = 1)$$

- A naive solution would be to use the statistical association as an estimate of the population causal effect
- E.g, to use

$$\frac{p(Y = 1|X = 1)}{p(Y = 1|X = 0)}$$

as an estimate of

$$\frac{p(Y_1 = 1)}{p(Y_0 = 1)}$$

- When does this naive solution work?

Exchangeability

- Intuitively, the statistical association

$$\frac{p(Y=1|X=1)}{p(Y=1|X=0)}$$

  is equal to the causal effect

$$\frac{p(Y_1=1)}{p(Y_0=1)}$$

  if there are no systematic differences between exposed and unexposed
- If so, we say that have **exchangeability** between exposed and unexposed
- The notion of exchangeability can be formalized with potential outcomes

Outline

Outline

Example

ID	X	Y
1	0	0
2	0	0
3	0	1
4	0	1
5	0	0
6	1	0
7	1	1
8	1	1
9	1	1
10	1	0

- *Compute the risk ratio*

Solution

From association to causation

- What we ultimately want is the causal risk ratio

$$\frac{p(Y_1 = 1)}{p(Y_0 = 1)},$$

not the 'associational' risk ratio

$$\frac{p(Y = 1|X = 1)}{p(Y = 1|X = 0)}$$

- But perhaps the 'associational' risk ratio is equal to the causal risk ratio?
- Let's see if we can prove this!
- The proof will require two steps

Step 1

ID	X	Y	Y_0	Y_1
1	0	0	0	?
2	0	0	0	?
3	0	1	1	?
4	0	1	1	?
5	0	0	0	?
6	1	0	?	0
7	1	1	?	1
8	1	1	?	1
9	1	1	?	1
10	1	0	?	0

- From the definition of potential outcomes, we have that
  - $Y_0 = Y$ for subjects with $X = 0$
  - $Y_1 = Y$ for subjects with $X = 1$
- It follows that

$$\frac{p(Y = 1|X = 1)}{p(Y = 1|X = 0)} = \frac{p(Y_1 = 1|X = 1)}{p(Y_0 = 1|X = 0)}$$

Step 2

ID	X	Y	Y_0	Y_1
1	0	0	0	?
2	0	0	0	?
3	0	1	1	?
4	0	1	1	?
5	0	0	0	?
6	1	0	?	0
7	1	1	?	1
8	1	1	?	1
9	1	1	?	1
10	1	0	?	0

- Next, we want to have that

$$\frac{p(Y_1 = 1|X = 1)}{p(Y_0 = 1|X = 0)} = \frac{p(Y_1 = 1)}{p(Y_0 = 1)}$$

- In order for this equality to hold, we must require that

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$

and

$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- *Fill in values of Y_0 and Y_1 such that these equalities hold*
- *Verify that the 'associational' risk ratio is equal to the causal risk ratio*

Solution

The two steps

- We saw that, in order for the 'associational' risk ratio to equal to causal risk ratio, we must require two conditions

1.
$$p(Y = 1|X = 0) = p(Y_1 = 1|X = 0)$$
$$p(Y = 1|X = 1) = p(Y_0 = 1|X = 1)$$

and

2.
$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- The first condition follows from the definition of potential outcomes and is always valid
- How about the second condition?

Example

ID	X	Y	Y_0	Y_1
1	0	0	0	?
2	0	0	0	?
3	0	1	1	?
4	0	1	1	?
5	0	0	0	?
6	1	0	?	0
7	1	1	?	1
8	1	1	?	1
9	1	1	?	1
10	1	0	?	0

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- *Can you fill in values of Y_0 and Y_1 such that these equalities **do not** hold?*

Solution

Step 2, revisited

2.

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- Clearly, it is possible that these equalities do not hold, in which case the second step is not valid
- As a consequence, association $\neq$ causation
- We will discuss when and why these equalitites may hold, and when they may be violated
- Before doing so we will put a name on these equalities

Exchangeability

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- If these equalitites hold, then we say that we have **exchangeability**
- If we have exchangeability, then association = causation

Alternative formulation of exchangeability

ID	X	Y	Y_0	Y_1
1	0	0	0	0
2	0	0	0	1
3	0	1	1	1
4	0	1	1	1
5	0	0	0	0
6	1	0	0	0
7	1	1	0	1
8	1	1	1	1
9	1	1	1	1
10	1	0	0	0

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

- Exchangeability means that
  - the potential outcome Y_0 has the same distribution among those with $X = 0$, as among everybody
  - the potential outcome Y_1 has the same distribution among those with $X = 1$, as among everybody
- In other words, that the potential outcomes (Y_0, Y_1) are independent of the exposure

$$(Y_0, Y_1) \amalg X$$

Testability

ID	X	Y	Y_0	Y_1
1	0	0	0	?
2	0	0	0	?
3	0	1	1	?
4	0	1	1	?
5	0	0	0	?
6	1	0	?	0
7	1	1	?	1
8	1	1	?	1
9	1	1	?	1
10	1	0	?	0

$$p(Y_0 = 1|X = 0) = p(Y_0 = 1)$$
$$p(Y_1 = 1|X = 1) = p(Y_1 = 1)$$

$$(Y_0, Y_1) \amalg X$$

- *Can we use the observed data to check if you have exchangeability or not?*

Solution

Outline

Exchangeability in randomized trials

- In a randomized trial we assign exposure levels randomly to each study participant
  - e.g. by the flip of a coin
- **This study design guarantees that we have exchangeability**
- ... so that association = causation

Motivation

- Under randomization, all pre-exposure variables are independent of the exposure
  - e.g. same distribution of age, sex and genetics in exposed and unexposed
- **The potential outcomes (Y_0, Y_1) are pre-exposure variables**
  - they describe how the subject will react to $X = 0$ and $X = 1$
  - this depends on factors that are determined before the factual exposure level is received. e.g. age, sex and genetics
- Thus, under randomization (Y_0, Y_1) are independent of X

$$(Y_0, Y_1) \amalg X$$

- *This is amazing! Why then not always randomize?*

Example

- Research question: does obesity (X) increase the risk of cardiovascular disease (Y)?
- To answer this question with a randomized trial we must:
  - get funding and ethical approval
  - collect study participants
  - randomize to different levels of BMI
  - follow for, say, 5 years
  - compare the incidence of cardiovascular disease across BMI levels
- *Is this feasible?*

Ethical problems

- We know today that obesity is dangerous for various reasons
- Thus, we would most likely not get ethical approval for a study where people are randomized to obesity

Difficult interventions

- To make sure that people reach their assigned BMI level, we must be able to control BMI by intervention
- In principle, we can think about possible interventions such as exercise, diet etc
- In practice, there is no guarantee that these interventions would really 'work' (e.g. give the desired BMI)
- And if they work, they may require unreasonable efforts from the study participants

Non-compliance

- Some subjects who are assigned to 'high BMI' may not comply with the required intervention
- Traditional analyses:
  - Intention To Treat (ITT)
  - As Treated (AT)
- Both these analyses are likely to be biased
  - alternative 'instrumental variables methods' exist (later)

Non-ignorable drop out

- Some people may drop out of the study ($D = 1$) before end of follow up
  - can compute $p(Y = 1|X, D = 0)$, but not $p(Y = 1|X)$
- Possibly systematic differences (e.g. in disease severity) between those who drop out and those who remain
- If so, then we may not have exchangeability among those who remain

$$(Y_0, Y_1) \not\perp\!\!\!\perp X \mid D = 0$$

Conclusion

- Real randomized trials often suffer from several important problems
- Observational studies are needed
  - in fact, most human knowledge comes from observations, e.g. evolution theory, smoking causes lung cancer etc
- And so are methods for causal inference from observational studies

Outline

Exchangeability in observational studies

- In observational studies, the exposure is not controlled by design
  - e.g. each study participants chooses his/her own exposure level
- **This study design does not guarantee that we have exchangeability**

Motivation

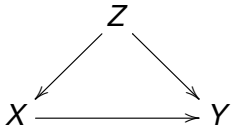

- Suppose that there is a covariate, Z, which affects both X and Y
  - e.g. Z = age; old people have higher BMI (X) than young people, and are more likely to get cancer (Y)
- If so, then X and Y will be associated, even in the complete absence of a causal effect
  - e.g. those with high BMI are older than those with low BMI, and for that reason more likely to get cancer
- Thus, association $\neq$ causation, i.e. we don't have exchangeability

Confounding

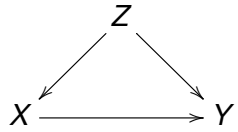

- When X and Y have common causes, we say that there is **confounding**
  - more on confounding later
- **Confounding leads to non-exchangeability**

Randomization and confounding

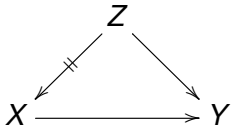

- When X is randomized the influence of all confounders is broken by design
- Thus, there is no confounding in an ideal randomized trial

How can we estimate causal effects if we don't have exchangeability?

- There are several ways to control/adjust for potential confounders in the analysis
  - stratification
  - matching
  - standardization
  - propensity scores
  - regression modeling
  - inverse probability weighting
  - etc
- Often combined

Conditional exchangeability

- Controlling for a potential confounder Z gives a causal effect if we have conditional exchangeability, given Z:

$$(Y_0, Y_1) \amalg X \mid Z$$

  - we say that Z is sufficient for confounding control
- Conditional exchangeability is untestable, and must be judged by subject matter knowledge
- Exchangeability can be achieved by controlling for covariates, but can also be 'destroyed'
  - much more later

Example

	$Z=0$		$Z=1$	
	$Y=0$	$Y=1$	$Y=0$	$Y=1$
$X=0$	270	30	120	80
$X=1$	180	20	60	240

- *Assume conditional exchangeability, given Z, and compute the conditional causal risk ratio, given Z, for $Z=1$ and $Z=0$.*
- *Where in the computation do you use the assumption of conditional exchangeability?*

Solution

Solution

Solution

Summary

- Exchangeability is defined as

$$(Y_0, Y_1) \amalg X$$

- If we have exchangeability, then the crude (without controlling for covariates) association is a causal effect
- In (ideal) randomized trials, exchangeability follows by design
- In observational studies, we typically don't have exchangeability because of confounding
- Conditional exchangeability, given Z, is defined as

$$(Y_0, Y_1) \amalg X \mid Z$$

- If we have conditional exchangeability, given Z, then controlling for Z gives a causal effect
  - we say that Z is sufficient for confounding control