Observational Studies: Identifying Causal Effects

Sebastien Haneuse, PhD

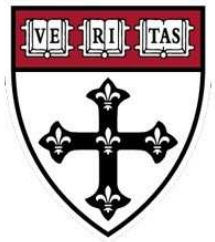

Department of Biostatistics
Harvard T.H. Chan School of Public Health

Non-small cell lung cancer

- Surgical resection is the current treatment of choice
- There are a number of surgical options:
  - ⋆ wedge resection
  - ⋆ segementectomy
  - ⋆ lobectomy
- Differ in terms of how much lung tissue is removed
- Decision between procedures is based, in part at least, on an attempt to balance between post-operative lung function and risk of recurrence
- First step might be to establish if there is a difference in the risk of recurrence
  - ⋆ let $Y{=}0/1$ denote recurrence within some frame (e.g. 1 year)

- For simplicity, we can consider a binary surgical 'treatment', denoted A
  - ⋆ $A = 0$ for wedge resection
  - ⋆ $A = 1$ for anatomic resection
    - ∗ segementectomy/lobectomy

Q What is the causal association between A and Y?

- To answer this question we need to

(i) explicitly define what we mean by 'causal association'

(ii) come up with a method for estimating it given data from either

- ∗ a randomized study
- ∗ an observational study

Counterfactual outcomes

Q: What is the difference in risk of recurrence be if a patient undergoes anatomic resection versus wedge resection?

- Question is framed in terms of the application of two treatments to the same person
  - ⋆ natural framing that aligns with how we think and make decisions
- Towards a formal treatment of this question, we introduce the notion of a *counterfactual outcome*
  - ⋆ denoted by $Y_{(a)}$
  - ⋆ value of Y had, possibly contrary to fact, the individual been exposed to $A = a$
  - ⋆ also called a *potential outcome*

- For the lung surgery setting, each individual has two counterfactuals:
  - $Y_{(0)}$: the value of response had they undergone wedge resection
  - $Y_{(1)}$: the value of response had they undergone anatomic resection

- In more general settings, an individual has a counterfactual outcome for *each* possible treatment scenario
  - continuous treatments
    - A takes on a infinite number of possible values
    - e.g. dose for some drug therapy
    - e.g. amount of time exercising
  - a course of treatment
    - A is a vector describing treatment at a series of time points
    - e.g. surgery followed by chemotherapy for 3 or 6 months
    - e.g. 200mg every other day or 100mg every day
  - define $Y_{(a)}$ for all possible treatment regimens

Individual causal effects

- Focusing on binary treatments, A has a causal effect if

$$Y_{(0)} \neq Y_{(1)}$$

- A has no causal effect if

$$Y_{(0)} = Y_{(1)}$$

  - ⋆ say that the *sharp causal null hypothesis* holds if this is true for all subjects in the population

- While these align with how we make decisions, in the *factual* world, at most one counterfactual outcome can be observed for any given individual
  - ⋆ i.e. the counterfactual corresponding to the treatment actually experienced

- Hence we are faced with a missing data problem
  - ⋆ one that can never be resolved with data
- Individual causal effects, therefore, cannot be determined
  - ⋆ on the basis of either a randomized study or an observational study
  - ⋆ not without extrapolating beyond the observable world
- Need another definition of 'causal effects'
  - ⋆ one that can, in principle, be determined with data
  - ⋆ one that requires weaker assumptions

Average causal effects

- One way forward is to define 'causation' in terms of a comparison between the *distributions* of the counterfactuals across the population of interest or some *summary feature* of the distribution

- For a binary outcome, it would be natural to examine

$$P(Y_{(a)} = 1)$$

  - ⋆ proportion of individuals who experience the outcome for a scenario where everyone receives treatment $A = a$

- A comparison between $P(Y_{(1)} = 1)$ and $P(Y_{(0)} = 1)$ therefore describes a comparison between two (possibly hypothetical) treatment scenarios for an entire population:

  A=0: everyone underwent wedge resection

  A=1: everyone underwent anatomic resection

- Causal risk difference:

$$\text{RD} = P(Y_{(1)} = 1) - P(Y_{(0)} = 1)$$

- Causal risk ratio:

$$\text{RR} = \frac{P(Y_{(1)} = 1)}{P(Y_{(0)} = 1)}$$

- Causal odds ratio:

$$\text{OR} = \frac{P(Y_{(1)} = 1)/(1 - P(Y_{(1)} = 1))}{P(Y_{(0)} = 1)/(1 - P(Y_{(0)} = 1))}$$

- When $P(Y_{(1)} = 1) = P(Y_{(0)} = 1)$ then there is no causal effect
  - ⋆ regardless of the choice of effect measure

- If we could estimate $P(Y_{(a)} = 1)$ then we could report any (or all) of the above effect measures

Estimating causal effects: a randomized study

- Suppose we recruit 100 subjects in a study of surgical treatment for lung cancer
  - ⋆ randomize half to $A = 0$ and half to $A = 1$
  - ⋆ record their outcomes:

$A = 0$	$A = 1$
$Y_{0,1}$	$Y_{1,1}$
$\vdots$	$\vdots$
$Y_{0,50}$	$Y_{1,=50}$
$\overline{Y}_0$	$\overline{Y}_1$

- Estimate $P(Y_{(a)} = 1)$ by the empirical mean, $\overline{Y}_a$
  - ⋆ the mean of Y among the 50 individuals who *actually* received $A = a$

- Plug-in the values of $\overline{Y}_a$ to estimate the causal risk difference:

$$\widehat{\mathsf{RD}} \;=\; \overline{Y}_1 - \overline{Y}_0,$$

or the causal risk ratio:

$$\widehat{\mathsf{RR}} \;=\; \frac{\overline{Y}_1}{\overline{Y}_0},$$

or the causal odds ratio:

$$\widehat{\mathsf{OR}} \;=\; \frac{\overline{Y}_1/(1-\overline{Y}_1)}{\overline{Y}_0/(1-\overline{Y}_0)}.$$

Q: Why is it that can we interpret these estimates in terms of causation?

- Intuitively, we appeal to randomization but what is the formal justification?

- We first need to outline two assumptions:

(i) **Consistency**

$$\text{if } A = a \text{ then } Y_{(a)} = Y$$

⋆ in principle, the (two) counterfactual outcomes exist regardless of which treatment is actually given

⋆ guarantees that what we observe when we do give a specific treatment is the corresponding counterfactual

(ii) **Exchangeability**

$$A \perp\!\!\!\perp Y_{(a)}, \ \forall \ a$$

⋆ treatment allocation is independent of the counterfactual outcomes

⋆ that is, treatment decisions are not based on what the outcomes might be under various scenarios for treatment

- Given consistency and exchangeability we have that

$$\begin{aligned} P(Y_{(a)} = 1) &= P(Y_{(a)} = 1|\ A) \\ &= P(Y_{(a)} = 1|\ A = a) \\ &= P(Y = 1|\ A = a) \end{aligned}$$

- That is, the probability of experiencing the outcome had, possibly contrary to fact, everyone been exposed to a is the same as the probability of experiencing the outcome among individuals who did get exposed to a
  - ⋆ the first is generally hypothetical, unobservable
  - ⋆ the second is observable

- The core reason why randomized studies are viewed as superior to observational studies is that exchangeability is (in theory) guaranteed

  Q: Why?

Estimating causal effects: an observational study

- In practice, it isn't always possible to conduct a randomized study
  - ⋆ most often because of ethical considerations
  - ⋆ sometimes for logistical reasons

- In the absence of a randomized study, we must appeal to observational data to answer questions of interest
  - ⋆ data collected as part of a specific research study
  - ⋆ data collected for some other (primary) purpose
    - ∗ e.g. an administrative claims data or EHR data

- For example, to learn about differences in surgical treatment options for NSCLC we could use data on patients who underwent surgery at BWH
  - ⋆ data collected during the course of usual clinical care at BWH

Q: Are there any issues associated with using observational data to establish causation?

- Key issue: surgical procedure was not randomized at BWH

Q: How were the decisions among possible procedures made?

- Choice of procedure was likely driven by a number of factors
  - ⋆ age
  - ⋆ lung function e.g. FEV_1, COPD
  - ⋆ smoking history
  - ⋆ other co-morbidities e.g. BMI
  - ⋆ expertise of the surgeon
- Each of these factors may also be related to the risk of recurrence

Q: What impact does this 'structure' have on our ability to establish causality?

- For simplicity, lets consider smoking history

$$L = \begin{cases} 0 & \text{never smoker} \\ 1 & \text{former/current smoker} \end{cases}$$

- Reasonable to suppose that smokers are more likely to
  - ⋆ have diminished lung function
  - ⋆ experience a recurrence

- Summarize these relationships with a directed acyclic graph (DAG):

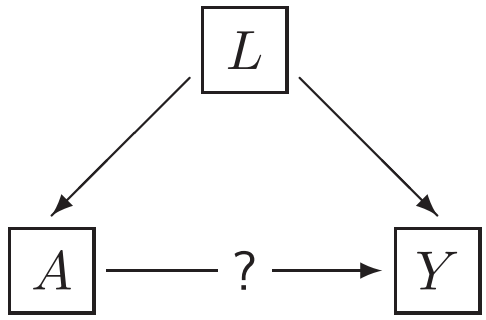

- Consider those who actually underwent wedge resection
  - ⋆ loss of lung function is likely a concern
  - ⋆ more likely to be former/current smokers

- Consider those who actually underwent anatomic resection
  - ⋆ loss of lung function is of less concern
  - ⋆ more likely to be never smokers

Q: What would have happened to the wedge resection patients had they undergone anatomic resection?

Actual treatment	Smoking distribution	Counterfactuals
$A = 0$	more former/current smokers	$Y_{(0)},\ Y_{(1)}$
$A = 1$	more never smokers	$Y_{(0)},\ Y_{(1)}$

- The distribution of their counterfactual outcomes would likely not be the same as the distribution of the counterfactual outcomes among the patients who did undergo anatomic resection
  - ⋆ $[A = 0]$ is 'enriched' with smokers who tend to have worse outcomes
- So, the distributions of the counterfactual outcomes are not independent of treatment allocation

$$Y_{(a)} \not\perp\!\!\!\perp A$$

  - ⋆ the exchangeability assumption does not hold

Confounding

Q: What if we ignore this absence of exchangeability?

- ⋆ we can, of course, continue to estimate $P(Y = 1|\ A = a)$ by $\overline{Y}_a$
- ⋆ and compute contrasts such as

$$\widehat{\mathsf{RD}} \ = \ \overline{Y}_1 - \overline{Y}_0$$

- The problem is that $P(Y = 1|\ A = a) \ \neq \ P(Y_{(a)} = 1)$
- Consequently, in the absence of exchangeability, effect measures based on $P(Y = 1|\ A = a)$ (i.e. quantities that are estimable) will not have a causal interpretation

Q: So what are we estimating?

- If we don't control for L, we are estimating some mixture of two components:

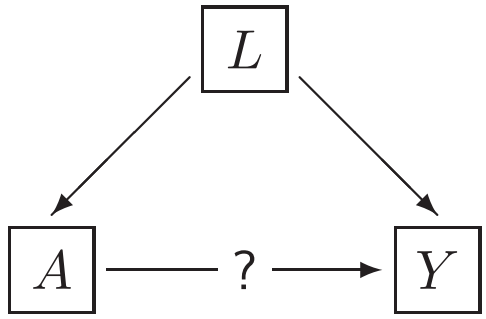

(i) the 'direct' causal association from $A \rightarrow Y$

(ii) an 'indirect' association, through L

- ⋆ sometimes referred to as a *backdoor* path

- The result is a *spurious* or *distorted* association

- Formally, the difference between what we are estimating and the causal association is referred to as *confounding bias*

- Covariates that contribute to confounding bias are *confounders*

- The extent of confounding bias will depend on many aspects of the problem
  - ⋆ number and type of confounders that have not been included in the analysis
    - ∗ i.e. unmeasured confounders, U
  - ⋆ strength of the associations between U and both A and Y
  - ⋆ strength of association between A and Y

- In practice, the interrelationships between A, Y and all potential confounders will likely be complex
  - ⋆ use the DAG framework to identify sufficient sets of covariates that permit the control of confounding
  - ⋆ domain knowledge will be critical as statistical criteria are generally inadequate
  - ⋆ Hernan and Robins (2017; Chapter 7)

Estimation of causal contrasts

- Within the counterfactual framework, to estimate causal associations from observational data we need to extend our assumptions:

(i) **Consistency**

$$\text{if } A = a \text{ then } Y_{(a)} = Y$$

(ii) **Conditional exchangeability**

$$A \perp\!\!\!\perp Y_{(a)} | \, \boldsymbol{L} = \boldsymbol{l}, \; \forall \; a, \boldsymbol{l}$$

(iii) **Positivity**

$$\text{if } f_{\boldsymbol{L}}(\boldsymbol{l}) > 0 \text{ then } f_{A|\boldsymbol{L}}(a|\boldsymbol{l}) > 0, \; \forall \; a$$

- Conditional exchangeability states that exchangeability holds *within* strata defined by $\boldsymbol{L}$

Smoking status, L	Actual treatment	Counterfactuals
never smokers	$A = 0$	$Y_{(0)},\ Y_{(1)}$
never smokers	$A = 1$	$Y_{(0)},\ Y_{(1)}$
former/current smokers	$A = 0$	$Y_{(0)},\ Y_{(1)}$
former/current smokers	$A = 1$	$Y_{(0)},\ Y_{(1)}$

Q: Is treatment allocation independent of $Y_{(a)}$ among never smokers?

Q: Is treatment allocation independent of $Y_{(a)}$ among former/current smokers?

- Referred to as the 'no unmeasured confounding' assumption

- Positivity states that all treatment choices are available to all sub-populations defined by $\boldsymbol{L}$

- Would not hold, for example, if never smokers were never treated with wedge resection
  - ⋆ i.e. don't worry as much about lung function and always go for the anatomic resection which minimizes future risk

- The practical implications are that if one is comparing $A = 0$ vs. $A = 1$, one shouldn't include sub-populations that would never receive one of the treatment choices
  - ⋆ if one does, estimation will rely on extrapolation
  - ⋆ interpretation of the estimated causal effects is unclear

Interpreting causal effects

- Although in some settings we have to use observational data/methods, thinking about the randomized study we would have conducted is a useful thought experiment

Q: If we could perform a randomized study, what would it look like?

  - ⋆ use this to drive decisions in our analysis of the observational data

- A key aspect of a randomized study is a well-defined intervention
  - ⋆ patients are randomized to a finite number of treatment regimens
  - ⋆ study protocols specify regimens in great detail

- For the lung surgery example, we would randomize patients to receive one of two treatment/intervention options:
  - ⋆ wedge resection
  - ⋆ anatomic resection

- In some instances, it may be difficulty to conceptualize a hypothetical randomized study

- Consider a study published by Rosenman et al (*JAMA* 233:872-877, 1975)
  - ⋆ Western Collaborative Group Study
  - ⋆ prospective study of coronary heart disease
  - ⋆ participants were categorized into one of two behavior pattern groups:

    Type A: characterized by *enhanced aggressiveness, ambitiousness, competitive drive, and chronic sense of urgency*

    Type B: characterized by *more relaxed and non-competitive*

Q: For the WCGS data, what would this intervention look like?

Q: Can we conceive of an intervention that directly corresponds to this 'exposure'?

- We could conceive of an intervention that seeks to make Type A individuals more relaxed
  - ⋆ herbal tea, yoga, a holiday?!?!
- But this is not what the paper was reporting on
- If we cannot conceive of an intervention that directly corresponds to behavior type

  Q: What does 'causation' mean in this context?

  Q: What are we estimating when we 'adjust for confounding'?

  Q: Can we call it a 'causal effect'?
- These issues arise in a broad range of public health settings:
  - ⋆ age, gender, race, weight, blood pressure, ...
- The key is to be precise in our interpretation of the results

The control of confounding bias

- Study design:
  - ⋆ *randomization*
    - ∗ known and unknown confounders
  - ⋆ *matching and/or restriction*
    - ∗ known confounders only
- Analysis:
  - ⋆ *inverse-probability weighting/standardization*
    - ∗ known confounders only
  - ⋆ *stratification-based methods (stratified analyses, restriction, regression)*
    - ∗ known confounders only

Reading

- Rothman K, Greenland S, and Lash T. **Modern Epidemiology** (2008)
  - ⋆ Chapter 4, *Measures of effect and measures of association*
- Hernan M and Robins J. **Causal Inference** (2017)
  - ⋆ Chapter 7, *Counfounding*
  - ⋆ available in pdf form through Miguel Hernan's website
- Krieger N and Davey Smith G. *The tale wagged by the DAG: broadening the scope of causal inference and explanation for epidemiology.* **International Journal of Epidemiology** (2016)