STAT6061/STAT5008 – Causal Inference

Part 1-2. Counterfactual Framework and Causal Estimands

An-Shun Tai

[1]*Department of Statistics*
National Cheng Kung University

[2]*Institute of Statistics and Data Science*
National Tsing Hua University

How to define and estimate causal effects?

- Suppose the three causal conditions are met, and all sources of systematic bias have been properly adjusted for.

- Can we then use a conventional regression model to assess causality? Consider the following two scenarios.

Scenario 1: assessing the causal effect of smoking on lung function (continuous outcome)

$$Y_{\text{lung function}} = \alpha_0 + \alpha_A A_{\text{smoking}} + \alpha_C C_{\text{confounder}} + \varepsilon$$

Scenario 2: assessing the causal effect of smoking on lung cancer development (binary outcome)

$$\text{logit}\{P(Y_{\text{lung cancer development}} = 1)\} = \beta_0 + \beta_A A_{\text{smoking}} + \beta_C C_{\text{confounder}}$$

- Do these parameters (i.e., α_A and β_A) represent the causal effects of interest? See Part 1.3 for the answer.

Counterfactual framework

- Remember that causality is fundamentally tied to **a manipulation or intervention** (such as treatment, exposure, action, or strategy) applied to **a unit**.

- Accordingly, the causal effect of a treatment or exposure on the outcome for unit *i* can be established if there is a difference between *"the outcome of unit i when receiving the treatment/exposure"* and *"the outcome of unit i when not receiving it"*—**measured at the same time**.

- Only one outcome can be observed, while the other is the **counterfactual** or **potential outcome**.

- Building on this concept, let's introduce the *counterfactual framework*, also known as the potential outcome framework or the Rubin Causal Model (Rubin, 1974).

Counterfactual framework

- **Setup**

Y: Outcome of interest (e.g., disease)

A: Intervention (treatment or exposure)

(NOTE1: In this course, we will primarily focus on **binary treatment** variables for A: **1 for the treatment and 0 for the control**)

C: Pretreatment measured confounders/covariates

U: Pretreatment unmeasured confounders/covariates

(NOTE2: Post-treatment confounders add complexity to causal inference. We will explore this further in Part 8.)

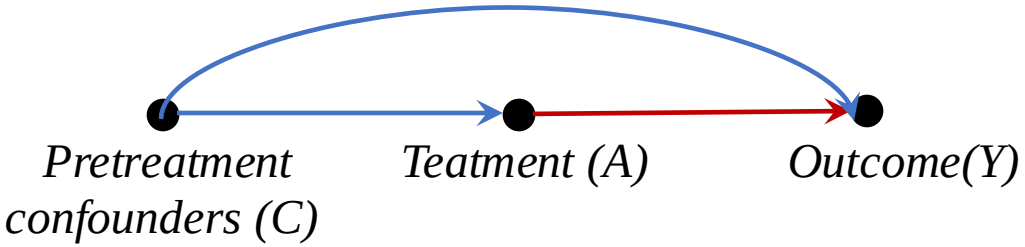

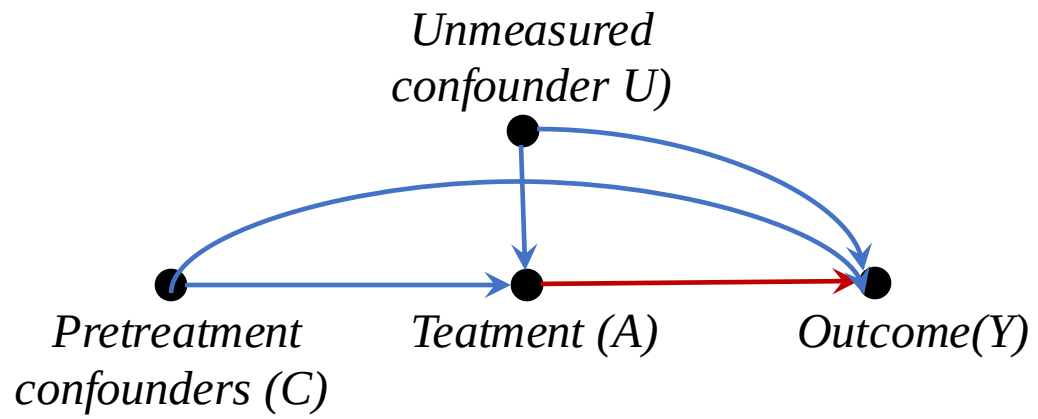

Counterfactual framework

- **Setup**

Consider a study with n experimental units, indexed by $i = 1, \ldots, n$, drawn from a well-defined target population, whose outcomes will be compared.

For each unit i, the observed data are (C_i, A_i, Y_i).

Definition of counterfactual outcomes (potential outcomes)

For each unit i, there are two potential outcomes :
- $Y_i(1)$: the outcome if unit i received the treatment,
- $Y_i(0)$: the outcome if unit i did not receive the treatment (control).

- For any given unit, only one of the two outcomes can be observed. The unobserved outcome is called as the "*counterfactual outcome*" or "*potential outcome*".

Science Table (God's Table)

- **Definition of the causal effect**

The causal effect for unit j is intuitively defined as the contrast between $Y_i(1)$ and $Y_i(0)$, such as the difference

$$Y_i(1) - Y_i(0).$$

Observations

Unit	Aspirin (A=1)	No Aspirin (A=0)
Y_1	No Headache (NH)	?
Y_2	?	Headache (H)
Y_3	?	No Headache (H)
Y_4	Headache (H)	?

Science Table

Unit	Aspirin (A=1)	No Aspirin (A=0)	Causal effect
Y_1	$Y_1(1)$ = NH	$Y_1(0)$ = NH	None
Y_2	$Y_2(1)$ = NH	$Y_2(0)$ = H	Improvement due to Aspirin
Y_3	$Y_3(1)$ = H	$Y_3(0)$ = NH	Aspirin-induced side effect
Y_4	$Y_4(1)$ = H	$Y_4(0)$ = H	None

Causal effects (Causal estimands, Causal parameters)

- Causal effects (also referred to as causal estimands, causal parameters) are the functions of potential outcomes $\{Y_i(1), Y_i(0)\}_{i=1}^{n}$.

Individual Treatment Effect, ITE (or Individual Causal Effect, ICE)

$$\tau_i \equiv Y_i(1) - Y_i(0)$$

Average Treatment Effect, ATE (or Average Causal Effect, ACE)*

$$\mathbb{E}[Y_i(1)] - \mathbb{E}[Y_i(0)] = \mathbb{E}_c[\tau_i(c)]$$

Conditional Average Treatment Effect, CATE

$$\mathbb{E}[Y_i(1)|X_i = x] - \mathbb{E}[Y_i(0)|X_i = x]$$

*This ATE is sometimes referred to as the causal risk differences (RD).

Some questions

1. Statistical parameters vs. causal parameter

Do you define statistical parameters first or establish the statistical model first?
Do you start with a scientific problem or formulate the statistical model first?

⇒ Causal effects are defined by potential outcomes, not by model parameters.

2. ITE vs. ATE

Can individual-level causal effects be directly estimated?
Can population-level causal effects be estimated?

⇒ Causal inference as a fundamental missing data challenge

3. Why do we need the conditional causal effect (i.e., CATE)?

Estimand, Estimator, Estimate

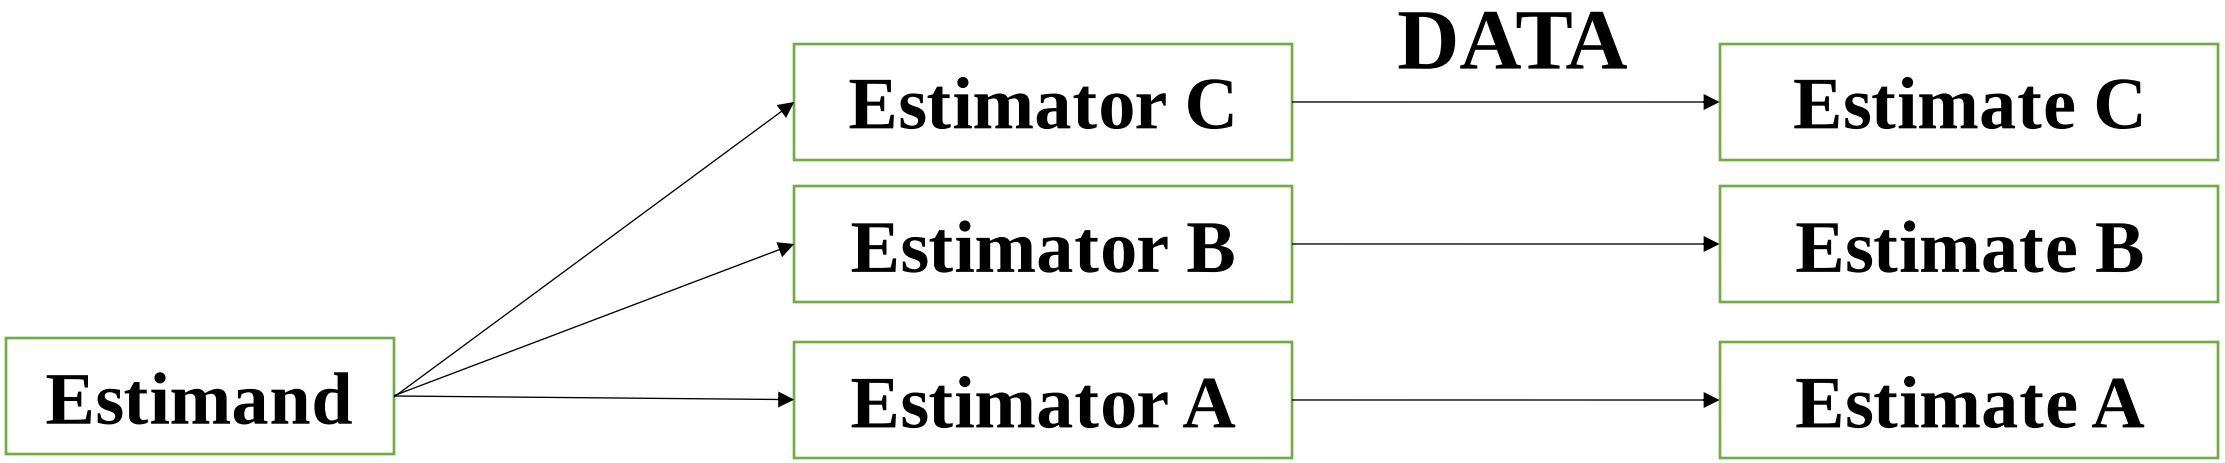

- Estimand is the specific quantity or parameter that a study aims to estimate to address a scientific problem.
- It defines the target of estimation before any data analysis begins, ensuring clarity in what the study seeks to measure.

- Estimator is the statistical method used to estimate the estimand/parameter.

- Estimate is the numerical result obtained from applying the estimator to data.

Conditional causal effect

$$\mathbb{E}[Y_i(1)|X_i = x] - \mathbb{E}[Y_i(0)|X_i = x]$$

- **Conditional causal effect help account for treatment effect heterogeneity**

- Personalized decision-making and targeted interventions
- Effect modification (Part 6)

- **Average treatment effect on the treated (ATT)**

$$\mathbb{E}[Y_i(1)|A_i = 1] - \mathbb{E}[Y_i(0)|A_i = 1]$$

- Helps determine if a treatment or program benefits those who actually participate.
- More practically relevant than the ATE in real-world scenarios where treatment assignment is not random.
- Requires fewer assumptions than ATE.

Causal effects on other scales

- The previously defined ITE, ATE, and CATE are all expressed on the (risk) difference scale.

$$\mathbb{E}[Y_i(1)] - \mathbb{E}[Y_i(0)]$$

- Other scales:

On the (risk) ratio scale: $\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)]$

For a binary outcome Y_i

On the odds ratio scale: $\mathrm{OR} = P(Y_i(1) = 1)[1 - P(Y_i(0) = 1)]/P(Y_i(0) = 1)[1 - P(Y_i(1) = 1)]$

For a survival outcome T_i

On the log hazard ratio scale: $\log\big(\lambda_T(1;t)\big) - \log\big(\lambda_T(0;t)\big)$, where

$$\lambda_T(a;t) = \lim_{h \to 0} P(t \le T(a) < t + h | T(a) > t)/h$$

- Causal effects measured on different scales offer distinct interpretations and, from a practical perspective, correspond to different statistical models.

Causal effects on the (risk) ratio scale

- For continuous outcomes

Individual causal ratio:

$$Y_i(1)/Y_i(0)$$

Population causal ratio:

$$\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)]$$

- For binary outcomes

Population causal (risk) ratio:

$$P(Y_i(1) = 1)/P(Y_i(0) = 1)$$

- Importantly, the population causal risk ratio is NOT the average of individual treatment effects on the ratio scale.

$$\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)] \neq \mathbb{E}[Y_i(1)/Y_i(0)]$$

- For example, $\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)] = 2$ indicates that the average potential outcome under treatment is twice that under control, but it does not imply that each individual's outcome doubles with treatment.

Causal effects on the odds ratio scale

➢ Odds

For a binary outcome Y with probability $\mu = P(Y = 1)$. The **odds** of the event occurring are given by $\mu/(1-\mu)$, representing the likelihood of occurrence relative to non-occurrence.

➢ Odds Ratio (OR)

OR compares the odds of an event occurring between the treatment ($A = 1$) and control ($A = 0$) groups, defined as

$$\frac{P(Y=1|A=1)}{1-P(Y=1|A=1)} \Big/ \frac{P(Y=1|A=0)}{1-P(Y=1|A=0)}$$

- OR remains unchanged regardless of the sampling method used in *case-control studies*, making it a robust measure of association.

➢ Causal Odds Ratio

$$\frac{P(Y_i(1)=1)}{1-P(Y_i(1)=1)} \Big/ \frac{P(Y_i(0)=1)}{1-P(Y_i(0)=1)}$$

Further insights on OR: Collapsibility

(Martinussen and Vansteelandt, 2013; Greenland, Pearl, and Robins, 1999)

➢ Collapsibility

- A measure is *collapsible* if the weighted average of the conditional measures (given a third variable X, which is not a confounder) equals the marginal measure.

(In regression contexts, a measure is *strictly collapsible* in a generalized linear model if its estimate remains unchanged whether or not Z is included in the model.)

- While (causal) risk differences and risk ratios are collapsible, the (causal) odds ratio is NOT.

- The concept of noncollapsibility is not a causal one, but rather a much more basic and general arithmetical one.

	$X = 1$		$X = 0$		Marginal	
	$A = 1$	$A = 0$	$A = 1$	$A = 0$	$A = 1$	$A = 0$
$Y = 1$	200	150	100	50	300	200
$Y = 0$	50	100	150	200	200	300
Risk	0.8	0.6	0.4	0.2	0.6	0.4
Risk difference	0.8 - 0.6 = 0.2		0.5 – 0.2 = 0.2		0.6 – 0.4 = 0.2	
Risk ratio	0.8/0.6 = 1.33		0.4/0.2 = 2		0.6/0.2 = 1.5	
Odds ratio	2.67		2.67		2.25	

NOTE1: Non-existence of Simpson's paradox

- **Simpson's paradox cannot happen to ATE**

- Let S_i denote the group corresponding to different stone sizes.

($S_i = 1$ for small kidney stones; $S_i = 0$ for large kidney stones)

- CATE

$$\tau_s = \frac{\sum_{i=1}^{n} I(S_i = s)\{Y_i(1) - Y_i(0)\}}{\sum_{i=1}^{n} I(S_i = s)}, \qquad s = 0, 1$$

- ATE

$$\tau = \frac{\sum_{i=1}^{n} I(S_i = 1)}{n}\tau_1 + \frac{\sum_{i=1}^{n} I(S_i = 0)}{n}\tau_0 = \pi_1\tau_1 + \pi_0\tau_0$$

- **If $\tau_1 > 0$ and $\tau_0 > 0$, we must have $\tau > 0$.**

NOTE2: Definition of the causal effect

➢ Question

If I previously did not take aspirin and my headache persisted, but now I take aspirin and my headache disappears, can I conclude that taking aspirin had a causal effect on my headache?

➢ Potential outcomes

$$Y_{i,\text{before}}(0) = 0, Y_{i,before}(1) = ?, Y_{i,\text{after}}(0) = ?, Y_{i,after}(1) = 1$$

- The individual causal effects, $Y_{i,\text{before}}(1) - Y_{i,\text{before}}(0)$ and $Y_{i,\text{after}}(1) - Y_{i,\text{after}}(0)$ are unknown, unless a strong assumption is made, such as $Y_{i,\text{before}}(0) = Y_{i,\text{after}}(0) = 0$

➢ Same unit at the same time

The causal effect compares potential outcomes **for the same unit at the same time post-treatment,** rather than comparing outcomes at different time points, such as a before-and-after assessment of a headache before and after taking aspirin (Imbens and Rubin, 2015).

NOTE3: A very important issue "Effects of Cause" not "Cause of Effects"

(Li and Mealli, 2014)

Fan Li asks:

In the Rubin Causal Model (RCM), cause/intervention should always be defined before you start the analysis. In other words, the **RCM is a framework to investigate the "effects of a cause," but not the "causes of an effect."** Some criticize this as a major limitation. Do you regard this as a limitation?

Do you think it is ever possible to draw inference on the causes of effects from data, or is it, per se, an interesting question worth further investigation?

Donald B. Rubin said:

I regard "the cause" of an event topic as more of a cocktail conversation topic than a scientific inquiry, because it leads to an essentially infinite regress.

Someone says, "He died of lung cancer because he smoked three packs a day"; then someone else counters, "Oh no, he died of lung cancer because both of his parents smoked three packs a day and, therefore, there was no hope of his doing anything other than smoking three packs a day"; then another one says, "No, no, his parents smoked because his grandparents smoked—they lived in North Carolina where, back then, everyone smoked three packs a day, so the cause is where the grandparents lived," and so on.

How far back should you go? **You can't talk sensibly about *the cause* of an event**; you can talk about "but for that cause (and there can be many 'but for's), what would have happened?" All these questions can be addressed hypothetically. But *the cause*? The notion is meaningless to me.

NOTE3: Cause of effects

(Dawid and Musio, 2022)

➢ Effects of cause

Forecasting (forward-looking causal inference)

Peter has started smoking daily. What is his risk of developing lung cancer?

Decision

Peter is considering smoking but is concerned about the risk of lung cancer. How should he decide?

➢ Cause of effects

Backcasting (backward-looking causal inference)

Peter has been diagnosed with lung cancer. Did he have a history of smoking? If so, what was the intensity and duration of his smoking?

Attribution

Peter has smoked daily for 10 years and developed lung cancer. Is his lung cancer caused by smoking?

➢ Probability of necessary causation (Pearl, 2000)

$$P(Y(0) = 0 | A = 1, Y = 1)$$

References

Dawid, A. P., & Musio, M. (2022). Effects of causes and causes of effects. *Annual Review of Statistics and Its Application*, *9*(1), 261-287.

Greenland, S., Pearl, J., & Robins, J. M. (1999). Confounding and collapsibility in causal inference. *Statistical science*, *14*(1), 29-46.

Imbens, G. W., & Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge university press.

Li, F. & Mealli, F. (2014). A conversation with Donald B. Rubin. *Statistical science*, *29*(3), 439-457.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*.. Cambridge university press.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, *66*(5), 688.