STAT6061/STAT5008 – Causal Inference

Part 8. Causal Inference in Longitudinal and Survival Data

An-Shun Tai

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

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

Longitudinal data

- Longitudinal data consist of repeated observations of the same subjects over time. This structure allows researchers to examine within-subject changes, assess temporal dynamics, and distinguish individual variability from population-level trends.

Note: In contrast to longitudinal data, cross-sectional data provide only a snapshot of observations at a single point in time, but they are often easier to collect, less costly, and useful for estimating prevalence and identifying associations in a population.

- Examples

Repeated measurements of blood pressure over several clinic visits; Annual cognitive test scores in an aging cohort study; Monthly HbA1c levels in diabetic patients; Student academic performance tracked across semesters; Weekly symptom reports in a clinical trial

- Common statistical methods

1. Linear mixed-effects models / Generalized linear mixed models
2. Generalized estimating equations
3. Growth curve models / Latent trajectory models
4. Marginal structural models / G-formula (causal inference)

Causal inference for longitudinal data

- In Parts 1 to 4, treatments, confounders (or covariates), and outcomes were all measured at a single time point.

(Note: Although each variable was observed only once, the actual measurement times may differ across variables.)

- In the longitudinal setting, data typically consist of:
  - Time-varying confounders
  - Time-varying treatments
  - Time-varying (repeated) outcomes

- Notations for two time points (baseline $t = 0$ and follow-up $t = 1$)
  - Y_0, Y_1 : Outcome at time 0 and time 1, respectively
  - A_0, A_1 : Treatment or exposure at time 0 and time 1
  - L_0, L_1 : Time-varying confounders at time 0 and time 1
  - C: Baseline (time-invariant) confounder

(L_0 vs. C and L_1 vs. Y_0)

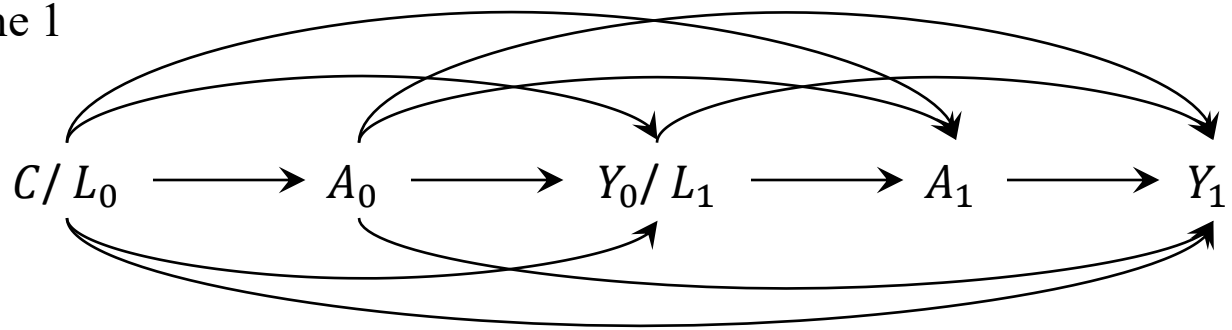

Post-treatment covariates

- A post-treatment covariates (L_1) is a variable that occurs after treatment and is associated with the outcome. While it can act as a time-varying confounder, not all post-treatment covariates change over time (time-invariant).

- Examples

(Truncation by Death)
In trials with severely ill patients, some may die before the outcome (like quality of life) is measured. The post-treatment covariates L_1 is a binary indicator of survival.

(Unemployment)
In job training studies, people are assigned to treatment or control and then report their employment status L_1 and wage Y. Here, L_1 indicates whether they are employed.

(Surrogate Endpoint)
In clinical trials, long-term outcomes (e.g., 30-year survival) are costly and time-consuming to measure. Researchers often use early, easier-to-measure variables called **surrogate endpoints**. For example, in HIV trials, **CD4 cell count** is used as a surrogate for long-term survival.

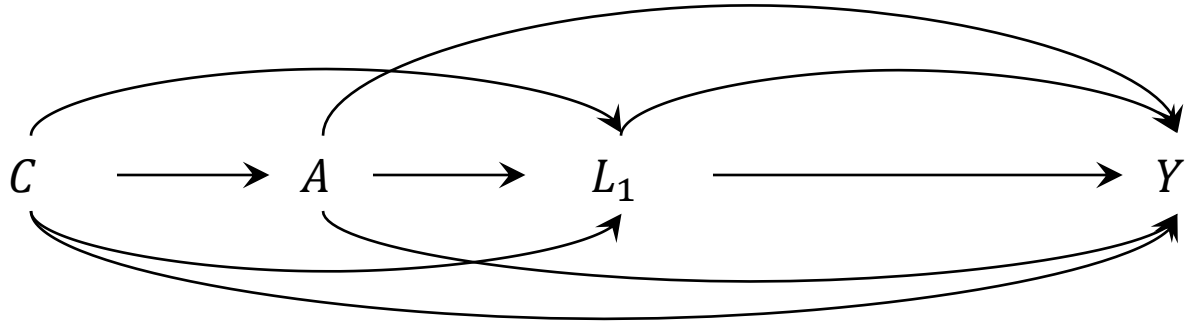

Conditioning on the post-treatment covariate

- A common but naive approach is to condition on the post-treatment covariates L_1 as if it were a baseline covariate.
- However, L_1 differs fundamentally from baseline covariates C because L_1 is generally affected by the treatment, while C is not.
- A general rule of thumb (Cochran, 1957; Rosenbaum, 1984):

  Avoid conditioning on post-treatment covariates when estimating the ATE.
- Frangakis and Rubin (2002) offered a more formal explanation using the potential outcomes framework.

1. Suppose we ignore the baseline confounder C; in that case, the exchangeability assumption is modified as follows:

$$\{Y(1), Y(0), L_1(1), L_1(0)\} \perp A$$

2. Conditioning on $L_1 = l$, we compute

$$\tau_L = \mathbb{E}(Y|A = 1, L_1 = l) - \mathbb{E}(Y|A = 0, L_1 = l)$$

- If L_1 is a baseline covariate, this comparison exactly reflects the conditional ATE.
- However, L_1 is a post-treatment covariates now, and the interpretation becomes problematic due to potential bias.

Conditioning on the post-treatment covariate (cont.)

3. Under the modified exchangeability assumption, we have

$$\mathbb{E}(Y|A=1, L_1=l) = \mathbb{E}(Y(1)|A=1, L_1(1)=l) = \mathbb{E}(Y(1)|L_1(1)=l)$$

and

$$\mathbb{E}(Y|A=0, L_1=l) = \mathbb{E}(Y(0)|A=0, L_1(0)=l) = \mathbb{E}(Y(0)|L_1(0)=l)$$

Therefore, τ_L correspond to the contrast

$$\mathbb{E}(Y(1)|L_1(1)=l) - \mathbb{E}(Y(0)|L_1(0)=l)$$

which compares the distributions of $Y(1)$ and $Y(0)$ cross different subsets of units—those defined by the post-treatment values $L_1(1)=l$ and $L_1(0)=l$, respectively.

- In general, comparisons conditional on $L_1=l$ do not have a valid causal interpretation, unless $L_1(1)=L_1(0)$, that is, the post-treatment covariate is unaffected by treatment.

- More on truncation by death
  - When the treatment increases survival, it may allow more weak individuals to survive compared to the control group.
  - Consequently, those with $L_1(1)=1$ (survivors under treatment) may be systematically weaker than those with $L_1(0)=1$ (survivors under control), highlighting that survival subsets under different treatments are not directly comparable.

Conditioning on the potential values of post-treatment covariates

- As mentioned earlier, why don't we simply ignore post-treatment covariates?

- Reasons

1. **When post-treatment covariates introduces confounding effects:**

- If a post-treatment covariates affects the outcome and is affected by treatment, it becomes a time-varying confounder.
- Ignoring it violates the exchangeability assumption, biasing your total effect estimate.

2. When selection bias occurs (e.g., Truncation by Death):

- If only a subset of individuals (e.g., survivors) have outcome data,
- And survival is influenced by treatment,
- Then the observed outcome is conditionally dependent on a post-treatment covariate (L_1).

→ You're comparing non-comparable groups unless you account for this.

3. When post-treatment covariates are used in outcome definitions:

- E.g., wage observed only if employed, or quality of life only if alive.
- Ignoring L_1 in such cases leads to estimands that are not well-defined.

Conditioning on the potential values of post-treatment covariates (cont.)

- Frangakis and Rubin (2002) suggests principal stratification (Part 4-3, monotonicity for IV)

$$\begin{cases} \tau(1,1) = \mathbb{E}(Y(1) - Y(0) | L_1(1) = 1, L_1(0) = 1) \\ \tau(1,0) = \mathbb{E}(Y(1) - Y(0) | L_1(1) = 1, L_1(0) = 0) \\ \tau(0,1) = \mathbb{E}(Y(1) - Y(0) | L_1(1) = 0, L_1(0) = 1) \\ \tau(0,0) = \mathbb{E}(Y(1) - Y(0) | L_1(1) = 0, L_1(0) = 0) \end{cases}$$

- Since $\{L_1(1), L_1(0)\}$ is unaffected by treatment, it can be treated as a baseline covariate.

→ In this case, $\tau(l_1, l_0)$ represents a subgroup (conditional) ATE.

- For subgroups where $L_1(1) = L_1(0)$:

→ The treatment does not alter the intermediate variable.

→ $\tau(1,1)$ and $\tau(0,0)$ capture **dissociative effects** (effects independent of the intermediate variable).

- For subgroups where $L_1(1) \neq L_1(0)$:

→ The treatment does affect the intermediate variable.

→ $\tau(l_1, l_0)$ reflects **associative effects**, showing how changes in L_1 are associated with changes in the outcome.

Conditioning on the potential values of post-treatment covariates (cont. 2)

➢ **Truncation by Death**

- The outcome (e.g., quality of life) is only defined for survivors.
- Therefore, the only meaningful subgroup effect is:

$$\tau(1,1) = \mathbb{E}(Y(1) - Y(1)|L_1(1) = 1, L_1(0) = 1)$$

- This is known as the **Survivor Average Causal Effect (SACE)**—the treatment effect among those who would survive under both treatment and control.

➢ **Unemployment**

- Similar to truncation by death, wages are only defined for those who are employed.
- Thus, the relevant effect is again $\tau(1,1)$, interpreted as the **Employed Average Causal Effect**.
- Earlier methods like the Heckman Selection Model treated wages for the unemployed as missing.
- But under the potential outcomes framework, $\tau(1,1)$ offers a more meaningful causal interpretation.

➢ **Surrogate Endpoint**

- We want to evaluate treatment effects **via a surrogate** (e.g., CD4 count for long-term survival).
- A good surrogate must satisfy:
  - **Causal necessity**: If the treatment doesn't change the surrogate, it doesn't change the outcome.
    → Requires $\tau(1,1) = 0$ and $\tau(0,0) = 0$.
  - **Causal sufficiency**: If the treatment changes the surrogate, it changes the outcome.
    → Requires $\tau(1,0) \neq 0$ and $\tau(0,1) \neq 0$.

Back in the longitudinal setting

➢ **Example: Time-varying Treatment in an HIV Study**

To illustrate time-varying causal inference, consider a setting in which treatment may change at each time point. Let A_k denote a **binary treatment indicator** at month k, where $k = 0,1,2,\ldots,59$, representing a 5-year follow-up period.

For instance, in a cohort of individuals infected with HIV:
- $A_k = 1$ if the individual receives antiretroviral therapy (ART) in month k,
- $A_k = 0$ if the individual does not receive treatment in that month.

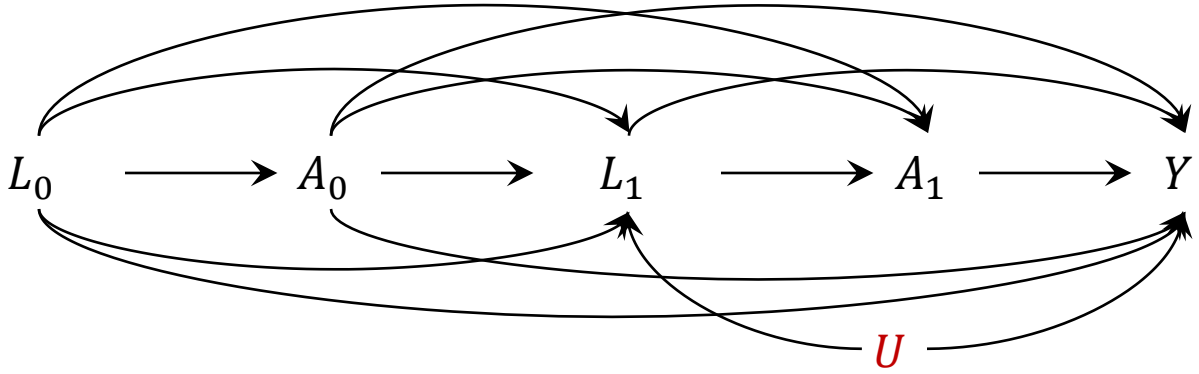

- Time-varying treatment (A_0 and A_1): ART can be initiated, discontinued, or re-initiated across time.
- Time-varying confounders (L_0 and L_1): CD4 count, viral load, or adherence status may evolve and also affect future treatment decisions and outcomes.
- Outcome (Y): Could be a long-term clinical endpoint like progression to AIDS or death.
- Unmeasured confounder (U): e.g., adherence, health behavior

Treatment–confounder feedback

- Standard regression or stratification **fails** — because:

1. Adjusting for a confounder that lies on the causal path blocks part of the treatment effect.
2. If the confounder is also influenced by unmeasured variables, adjusting for it can introduce collider bias.

- L_1 is affected by prior treatment and affects future treatment → feedback loop.

- Even in randomized trials, if treatment depends on prior L_k, feedback arises. Standard models yield biased estimates of causal effects.

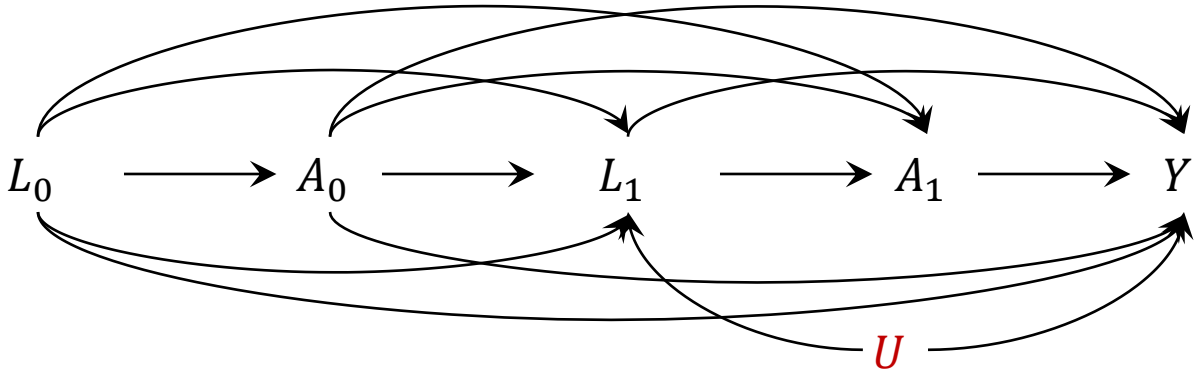

Hazard Function

- The hazard function describes the instantaneous risk of an event occurring at a specific time, given that the individual has survived up to that time.

It captures the rate at which events happen over time and is central to models like the Cox proportional hazards model.

- For a nonnegative random variable T denoting time to event, the hazard function $\lambda(t)$ is defined as:

$$\lambda(t) = \lim_{\Delta t \to 0} \frac{P(t \leq T < t + \Delta t | T \geq t)}{\Delta t} = \frac{f(t)}{S(t)}$$

where $f(t)$ is the probability density function of T and $S(t) = P(T > t)$ is the survival function.

For a survival outcome T

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

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

The Hazards of Hazard Ratios

➢ **Example: Women's Health Initiative (WHI)** (Manson et al., 2003)

- Study Design: Randomized controlled trial comparing combined estrogen plus progestin therapy vs. placebo in over 16,000 postmenopausal women.

- Follow-up Duration: Average of 5.2 years before the trial was halted due to safety concerns.

- Primary Outcome: Coronary heart disease (CHD) events.

➢ Main results from Manson et al. (2003)

1. Single HR of 1.24, interpreted as a 24% increased risk of CHD with hormone therapy. (weighted average over follow-up)

2. Year-specific HRs:

Year 1: 1.81; Year 2: 1.34; Year 3: 1.27; Year 4: 1.25; Year 5: 1.45; Year 6+: 0.70

→ Does this look off to you?

The Hazards of Hazard Ratios (cont.)

➢ **Issue 1: Interpretational Problem**

- If the study had ended after 1 year → HR = **1.8** → large harmful effect.
- After 5 years → HR = **1.2** → moderate effect.
- If it continued longer → possibly HR = **1.0** or less → no effect.

This demonstrates how the average HR depends on follow-up duration, not just on biology or treatment effect.

➢ **Issue 2: Selection Bias (Depletion of Susceptibles)**

- Susceptible women (those at higher CHD risk due to hormone therapy) were more likely to experience the event early.
- Over time, the remaining women in the treatment arm were less susceptible, biasing later HRs downward.
- The HR of 0.70 after year 5 might falsely suggest a protective effect, when in reality it's a result of selective survival.

While HRs can be useful in some contexts, they should not be the default causal estimand, especially in observational studies where time-varying confounding and selection effects are common.

References

Frangakis, C. E., & Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, *58*(1), 21-29.

Hernán, M. A. (2010). The hazards of hazard ratios. *Epidemiology*, *21*(1), 13-15.

Manson, J. E., Hsia, J., Johnson, K. C., Rossouw, J. E., Assaf, A. R., Lasser, N. L., ... & Cushman, M. (2003). Estrogen plus progestin and the risk of coronary heart disease. *New England Journal of Medicine*, *349*(6), 523-534.