Introduction to Causal Inference

Lecture 1

Introduction & The Neyman-Rubin Framework

邓婉璐

WANLUDENG@TSINGHUA.EDU.CN



清华大学统计与数据科学系

Outline

- Motivation
- Introduction to the course
  - Big names in history
  - Content of the course
  - Goal of the course
- Neyman-Rubin framework
  - Key concepts
  - Challenges
- Lord’s paradox - Example

Introduction

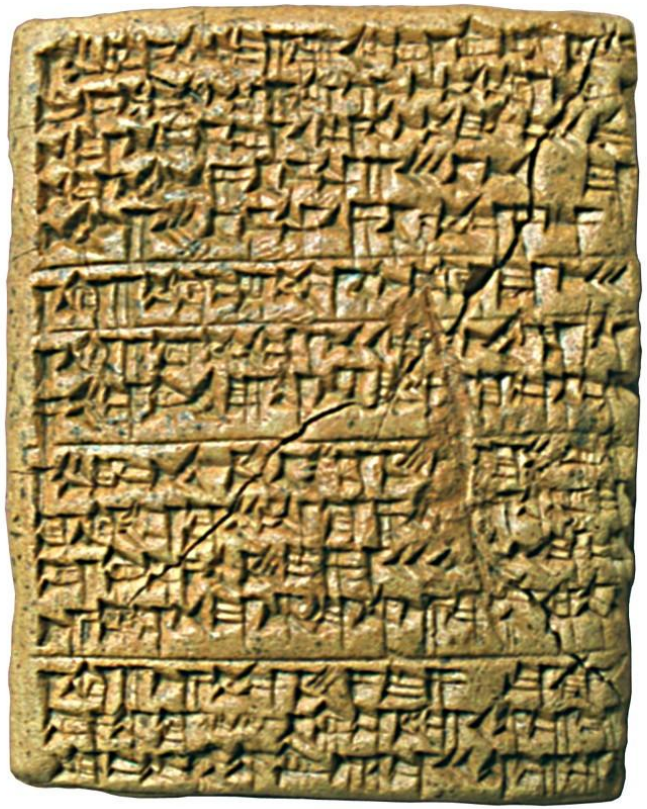


Concern of Causality

- Causality has been of concern since the dawn of civilization.

- 一种新药在特定人群中是否有效?
- 空气质量的改变是否是由汽车限号政策引起的?
- 从传统课堂换到翻转课堂能否提高学习成绩?

- Goal: Guide our behavior or decisions!



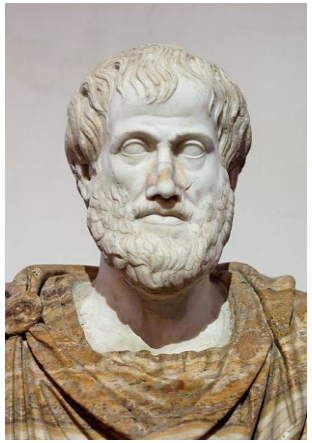
MS 4575
Medical text; the scribe's exemplar is broken.
Uruk, ca. 300 BC

http://www.academia.dk/MedHist/Bibliotheket/mesopotamien_farmakope.php

Causality in Philosophy

- The word ‘cause’ is used to mean “explanation” or “answer to a why question”

—— Aristotle (384-322 BC)
Ancient Greek philosopher



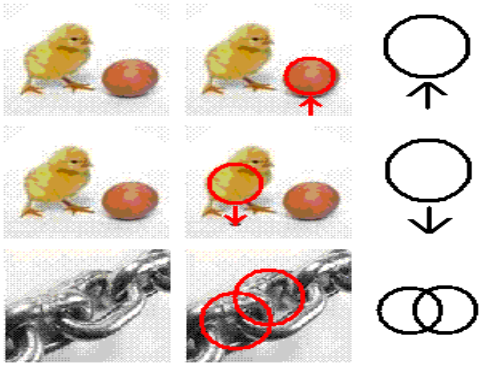
- “I would rather discover one causal law than be the king of Persia”

—— Democritus (460-370 BC)
Ancient Greek philosopher



What is Causality?

- Causality is an abstraction that indicates how the world progresses.
- Causality is the agency or efficacy that connects one process (the cause) with another process or state (the effect), where the first is understood to be partly responsible for the second, and the second is dependent on the first.
  - 前因后果



- What is the cause of a given effect?
- What is the effect of a given cause?

How to express causality?

Some statements

- My headache went away because I took an aspirin
- She got a good job last year because she went to college



How to express causality? Does correlation work?

- Cocks crow v.s. sunrise
- Same time on different watches.

correlation => causality?



- Healthy worker effect
- Taiji v.s. life time

causality => correlation?



Causality vs Correlation in Statistics

- Older textbooks avoid any mention of the term other than in settings of randomized experiments (Neyman 1932)
- Some even caution their readers to avoid using causal language in statistics
- Some mention it mainly to stress that correlation or association is not the same as causation

Causality better than Correlation?

- Correlation: Prediction
- Causality: Decision



Causal Inference

- “你理解的因果推断指代什么？”



Thinking: Base for causal inference?

全国II卷

阅读下面的材料，根据要求写作。

“二战”期间，为了加强对战机的防护，英美军方调查了作战后幸存飞机上弹痕的分布，决定哪里弹痕多就加强哪里。然而统计学家沃德力排众议，指出更应该注意弹痕少的部位，因为这些部位受到重创的战机，很难有机会返航，而这部分数据被忽略了。事实证明，沃德是正确的。

要求；综合材料内容及含义，选好角度，确定立意，明确文体，自拟题目。不要套作，不得抄袭，不少于800字。



Assumption about data collection matters!

Base for Causal Inference

3 parts:

- Causal model
- Assumption about data collection
- Collecting data, to get the conclusion about parameters in causal models.



Big Names

- Rigorous statistical frameworks for causality that can be applied to many types of data have been established for only a century.



R.A. Fisher
(1890-1962)



Jerzy Neyman
(1894-1981)



Donald B. Rubin
(1943-)



James M. Robins
(1946-)

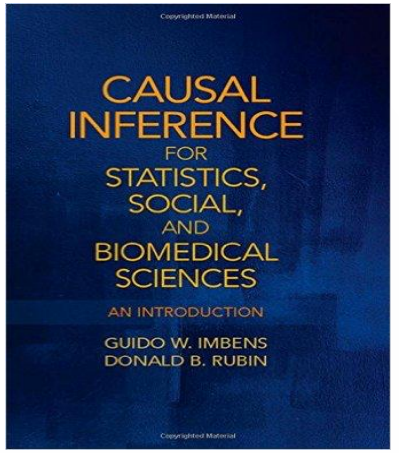


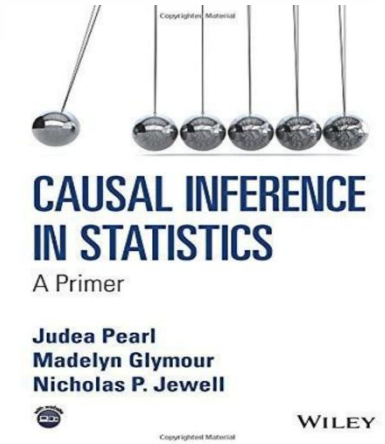
Judea Pearl
(1936-)

Course information

Main topics:

Two causal frameworks, classical randomized experiments, observational studies, directed acyclic graphs, instrumental variables.





Textbooks:

- Guido W. Imbens & Donald B. Rubin (2015) *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press.
- Judea Pearl (2016) *Causal Inference in Statistics - A Primer*, Wiley.

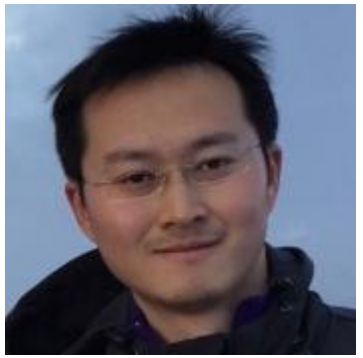
References

Other books:

- Judea Pearl (2011) *Causality: Models, Reasoning, and Inference* (2nd edition), Cambridge University Press.
- Hernán MA, Robins JM (2020). *Causal Inference: What If.* Boca Raton: Chapman & Hall/CRC.
- Peng Ding (2024) *A First Course in Causal Inference*. Chapman & Hall/CRC.



James M. Robins (1946-)



Peng Ding

Goal for The Course

- Causal inference vs other Statistical courses?

After the course:

- Capacity of basic frameworks / concepts for causal inference
- Basic ability to judge a causal statement
- Basic ability to estimate causal effects
- Basic ability to find causal relationships



Course information

- Evaluation:
  - Homework (≈every 2 weeks) 30%,
  - Midterm Exam (11th Nov) 30%,
  - Final Exam (30th Dec) 40%
- TA:
  - 陆畅 lu-c26@mails.tsinghua.edu.cn
- Office Hour:
  - ✓ Lecturer: Thursday 11am-12pm, 吕大龙楼630
  - ✓ TA: Thursday 2-3pm. 吕大龙楼8层讨论区

群聊：2026秋因果推断课程群

该二维码7天内(9月18日前)有效，重新进入将更新

The Neyman-Rubin Framework

FIND THE EFFECT OF A GIVEN CAUSE



Causal Effects in Common Usage







你后悔过么？

Causality in Everyday Life

Causal statements in an informal way

- My headache went away because I took an aspirin
- She got a good job last year because she went to college
- She has long hair because she is a girl

Thinking & Discussion

- How to express the causality?
  - Corr(X,Y)?
  - Y|X?

Applying Actions to a Unit

- A unit:
  - a physical object, a firm, an individual person
  - or collection of objects or persons, such as a classroom or a market
- An action:
  - a manipulation, treatment, or intervention that can be potentially applied to a unit
  - *Example:* in the aspirin-headache example, we can have different actions:
    - take an aspirin
    - not to take the aspirin
    - take an alternative medicine
- The key idea: causality is tied to an *action* applied to a *unit*



Treatment & Control

- Two actions:
  - We primarily considers settings with two actions
  - extensions to multi-valued treatments are conceptually straightforward.
- Active treatment (treatment):
  - the more active one of the two treatment
  - e.g., taking an aspirin
- Control treatment (control):
  - the more passive one of the two treatment
  - e.g., not taking an aspirin
- "treatment" and "control" are merely labels and formally are viewed symmetrically



Potential Outcomes & Causal Effect

- Each action-unit pair is associated with a potential outcome
  - given a unit and a set of actions
- The causal effect of one action relative to another involves the comparison of these potential outcomes:

$$Y_t - Y_c \quad \text{or} \quad \frac{Y_t}{Y_c}$$

Remark: any treatment must occur temporally before the observation of any associated potential outcome is possible



Unit-Level Causal Effect

Table 1.1. ***Example of Potential Outcomes and Causal Effect with One Unit***

Unit	Potential Outcomes		Causal Effect
	Y(Aspirin)	Y(No Aspirin)	
You	No Headache	Headache	Improvement due to Aspirin

Four possibilities of mappings for the potential outcomes in the aspirin-headache example:

1. Headache gone only with aspirin:
   Y(Aspirin) = No Headache, Y(No Aspirin) = Headache
2. No effect of aspirin, with a headache in both cases:
   Y(Aspirin) = Headache, Y(No Aspirin) = Headache
3. No effect of aspirin, with the headache gone in both cases:
   Y(Aspirin) = No Headache, Y(No Aspirin) = No Headache
4. Headache gone only without aspirin:
   Y(Aspirin) = Headache, Y(No Aspirin) = No Headache



Causal Statements under N-R Model (I)

- The informal causal statement I:

  "My headache went away because I took an aspirin"

- In this example, it is clear that:
  - Unit: I
  - Treatment: I took an aspirin ⟹ Potential outcome Y_t (Had no headache)
  - Control: at the time that I took the aspirin, I could have followed the alternate course of not taking an aspirin ⟹ Potential outcome Y_c (Had headache or not)

- The causal effect of *taking aspirin* with respect to *not taking aspirin* for *me* is defined as:

$$Y_t - Y_c$$



Causal Statements under N-R Model (II)

- The informal causal statement II:

  "She got a good job last year because she went to college"

- In this example, the alternative action is less clear:
  - Unit: she
  - Treatment: she went to college → Potential outcome Y_t (Got the current job)
  - Control: at the point in time when she decided to go to college, she could have decided not to go to college
  - But, she might have many different choices at that time point: Y_c

    she could have enrolled in the military → Potential outcome Y_c^1

    she could have joined an artist's colony → Potential outcome Y_c^2

- The causal effect of *going to college* with respect to *not going to college* for "*her*" is not clearly defined: $Y_t - Y_c^1$ or $Y_t - Y_c^2$ or others?



Causal Statements under N-R Model (III)

- The informal causal statement III:

  "She has long hair because she is a girl"

- In this example, the alternative action is not at all clear
  - Unit: she
  - Treatment: she is a girl ⟹ Potential outcome Y_t (Has long hairs)
  - Control: she is boy.
  - But, there is no action articulated that would have made her a boy and allowed us to observe the alternate potential outcome of hair length for this person as a boy.
  - Are we talking about: surgical procedures or hormone treatments (at a certain age)? ⇢ Different actions lead to different potential outcomes
- Hence, in our approach, this "because" statement is ill-defined as a causal statement

A Summary of Neyman-Rubin Model

- Units: the subjects on which the causal effects to be defined
- Actions
  - A group of "operations"
  - typically two: treatment t & control c
- Potential outcomes: $Y_i(t), Y_i(c)$ for ith unit
  - each associated with an action-unit pair
  - indicating the overall consequences of the action on the unit
- Unit-level causal effect:
  - comparison of the potential outcomes $Y_i(t) - Y_i(c)$ or $Y_i(t)/Y_i(c)$
  - plays a central role in causal inference as it defines the "causal effect"



Key Goal

- Unit Level Causal Estimand:

  For some given unit i

  $\{0,1\}$ ------▸ control & treatment

$$\tau_i = Y_i(t) - Y_i(c)$$
$$= Y_i(1) - Y_i(0)$$

- Population Level Causal Estimand
  - How to define?
  - Will be shown later



Challenges

ASSUMPTIONS



Causal Effects in Common Usage







- In the movie we can observe both, but in reality we would never be able to see both worlds.

The Basic Challenge of Causal Inference

Table 1.2. ***Example of Potential Outcomes, Causal Effect, Actual Treatment, and Observed Outcome with One Unit***

Unit	Not Observable			Known	
	Potential Outcomes		Causal Effect	Actual Treatment	Observed Outcome
	Y(Aspirin)	Y(No Aspirin)			
You	No Headache	Headache	Improvement due to Aspirin	Aspirin	No Headache

At most one of the two potential outcomes is realized and observable!

Learning Causal Effects via Multiple Units

- The basic challenge of causal inference
  - only one potential outcome is realized and observable
  - the other not realized and therefore not observable
- Possible solution: observe multiple units, some exposed to the active treatment, some exposed to the alternative (control) treatment
  - Option I: observe the same physical object under different treatment levels at different points in time
  - Option II: observe different physical objects at approximately the same time
- Be aware that:
  - different physical objects at the same time point are different units
  - the same physical object at different times are different units



Same Object at Different Times

A causal statement

"I might feel confident that an aspirin is going to relieve my headache within an hour, based on previous experiences, including episodes when my headache went away when I took an aspirin, and episodes when my headache did not go away when I did not take aspirin."

- The causal statement are shaped by comparisons of multiple units:
  - mysel
- There is different at
  - We vi
- Time ma
  - I may
  - eveni
  - the initial intensity of my headache may affect the result.

➢ it is often reasonable to assume that time makes little difference for inanimate objects

➢ but this assumption is typically less reasonable with human subjects

➢ It is never correct to confuse assumptions with definitions

Different Objects at the Same Time

A causal statement

"Both you and I have headaches, but only one of us takes an aspirin, we may attempt to infer the efficacy of taking aspirin by comparing our subsequent headaches."

- The causal statement are shaped by comparisons of multiple units:
  - "you" and "I" at the same point
- It is more obvious here that "you" and "I" at the same point in time are different units
  - Your headache status after taking an aspirin can obviously differ from what my headache status would have been had I taken an aspirin.
  - I may be more or less sensitive to aspirin
  - I may have started with a more or less severe headache



Presence of Multiple Units does not Solve the Problem

- Problem setting (an example):
  - consider the aspirin example with two units: You and I
  - two possible treatments for each unit: aspirin or no aspirin
  - assume that the two available aspirin tablets are equally effective
- There are four treatment levels (instead of two): $\vec{t} = (t_1, t_2)$
  - you take an aspirin and I do not - we both take an aspirin
  - I take an aspirin and you do not - neither of us does
- Leads to
  - four potential outcomes for "each of us" : $Y_i(\vec{t})$
  - six different comparisons defining causal effects for each of us
    - we can compare the status of my headache if we both take aspirin with the status of my headache if neither of us takes an aspirin, or
    - we can compare the status of my headache if only you take an aspirin to the status of my headache if we both do.



Simplify the Problem with Assumptions

- A typical assumption:

  "whether you take an aspirin does not affect my headache status"

  it is important to understand the force of such an assumption

- We should never forget that it is an assumption, not a fact
  - often a strong and controversial assumption
  - and may be false
- A counter example with "interactions":

  "we are in the same room, and unless you take an aspirin to ease your own headache, your incessant complaining will maintain my headache"



Stable Unit Treatment Value Assumption

Assumption 1.1 (SUTVA)

- **No Interference:** the potential outcomes for any unit do not vary with the treatments assigned to other units $\Longrightarrow Y_i(\vec{t}) = Y_i(t_i)$
- **No Hidden Variations of Treatments:** for each unit, there are no different forms or versions of each treatment level, which lead to different potential outcomes $\Longrightarrow Y_i(\vec{t}) = Y_i(t_i) \in \{Y_i(t), Y_i(c)\}$

- *exclusion restrictions*: assumptions that rely on external, substantive, information to rule out the existence of a causal effect of a particular treatment relative to an alternative.
- Causal inference is generally impossible without such assumptions
- It is critical to be explicit about their content and their justifications.



The Power of SUTVA (example)

Table 1.3. ***Example of Potential Outcomes and Causal Effects under SUTVA with Two Units***

Unit	Unknown		
	Potential Outcomes		Causal Effect
	Y(Aspirin)	Y(No Aspirin)	
You	No Headache	Headache	Improvement due to Aspirin
I	No Headache	No Headache	None

- SUTVA simplified the model (potential outcomes, def. of causal effects)
- But SUTVA is not enough for causal inference. E.g. in this example, observed outcomes showed no effect, yet it may due to different initial headache level.



Population Level Causal Estimand

- a population of units, indexed by $i = 1, \ldots, N$
- $\mathbb{T}_i = \mathbb{T} = \{0, 1\}$, ----→ control & treatment

Unit-level causal effect

average causal effect over the finite sample population ←---- $\tau_{\text{fs}} = \frac{1}{N} \sum_{i=1}^{N} \big(Y_i(1) - Y_i(0)\big),$

Average over the population

Assignment Mechanism (Introduction)

Treatment indicator representing the treatment assigned to unit i

Potential outcomes of unit i

The realized outcome

$$Y_i^{\text{obs}} = Y_i(W_i) = \begin{cases} Y_i(0) & \text{if } W_i = 0, \\ Y_i(1) & \text{if } W_i = 1. \end{cases}$$

The missed outcome

$$Y_i^{\text{mis}} = Y_i(1 - W_i) = \begin{cases} Y_i(1) & \text{if } W_i = 0, \\ Y_i(0) & \text{if } W_i = 1. \end{cases}$$

The alternative treatment that is not assigned to unit i

An Example of Ideal Assignment

Table 1.4. ***Medical Example with Two Treatments, Four Units, and SUTVA: Surgery (S) and Drug Treatment (D)***

Unit	Potential Outcomes		Causal Effect
	$Y_i(0)$	$Y_i(1)$	$Y_i(1) - Y_i(0)$
Patient #1	1	7	6
Patient #2	6	5	−1
Patient #3	1	5	4
Patient #4	8	7	−1
Average	4	6	2

Table 1.5. ***Ideal Medical Practice: Patients Assigned to the Individually Optimal Treatment; Example from Table 1.4***

Unit i	Treatment W_i	Observed Outcome Y_i^{obs}
Patient #1	1	7
Patient #2	0	6
Patient #3	1	5
Patient #4	0	8

Ignoring Assignment mechanism may lead to wrong judgement.

Pre-Treatment Variables or Covariates

Consider the case with many units & SUTVA

- At least half of all potential outcomes will be unobserved or missing
- We need to predict/impute the missing potential outcome to estimate the causal effect for any particular unit
- Often the presence of unit-specific background attributes (i.e., pre-treatment variables or covariates) X_i for unit i can assist in making these predictions.

Covariates

- *Key feature*: a priori known.
- *Main function*:
  - Subgroup
  - Estimate more precisely
  - Effect on assignment mechanism

⟹ Subgroup

Average Causal Effects on Sub-populations

average causal effect over the finite sample sub-population of females

$$\tau_{\mathrm{fs}}(f) = \frac{1}{N(f)} \sum_{i:X_i=f} \big(Y_i(1) - Y_i(0)\big).$$

Unit-level causal effect

Average over the sub-population of females

average causal effect over the finite sample sub-population of treated units

$$\tau_{\mathrm{fs},t} = \frac{1}{N_t} \sum_{i:W_i=1} \big(Y_i(1) - Y_i(0)\big)$$

Unit-level causal effect

Average over the sub-population of treated units

average causal effect over the finite sample sub-population of selected units

$$\tau_{\mathrm{fs,pos}} = \frac{1}{N_{\mathrm{pos}}} \sum_{i:Y_i(0)>0,Y_i(1)>0} \big(Y_i(1) - Y_i(0)\big),$$

Unit-level causal effect

Average over the sub-population of selected units

Population Level Causal Estimand

- a population of units, indexed by $i = 1, \ldots, N$,
- $\mathbb{T}_i = \mathbb{T} = \{0, 1\}$, ----→ control & treatment

Unit-level causal effect

average causal effect over the finite sample population ←---- $\tau_{\text{fs}} = \frac{1}{N} \sum_{i=1}^{N} \big(Y_i(1) - Y_i(0)\big)$,

Average over the population

More general form: $\tau = \tau(\underset{N\times 1}{\mathbf{Y}(0)}, \underset{N\times 1}{\mathbf{Y}(1)}, \underset{N\times k}{\mathbf{X}}, \underset{N\times 1}{\mathbf{W}})$

Example: Lord’s Paradox

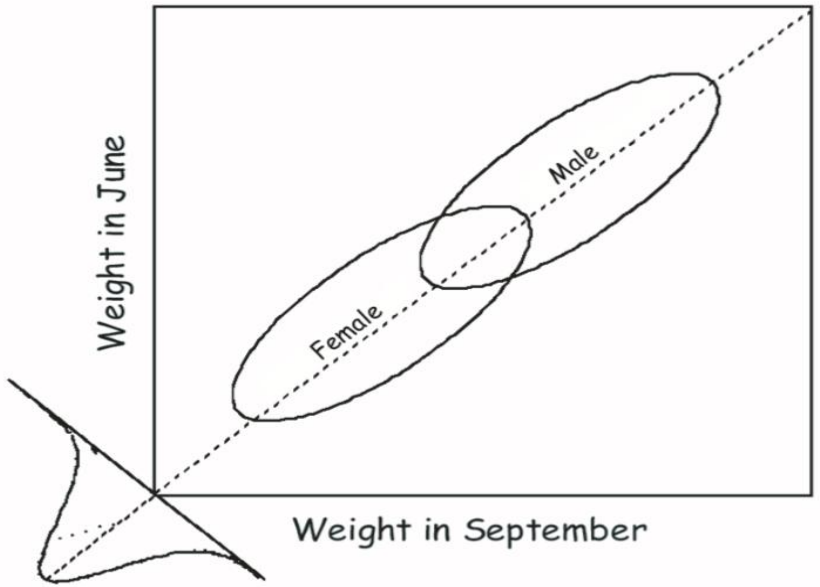


Lord’s Paradox

- The original problem:

  A large university is interested in investigating the effects on the students of the diet provided in the university dining halls and any sex differences in these effects. Various types of data are gathered. In particular, the weight of each student at the time of his arrival in September and his weight the following June are recorded. (Lord, 1967, p. 304)

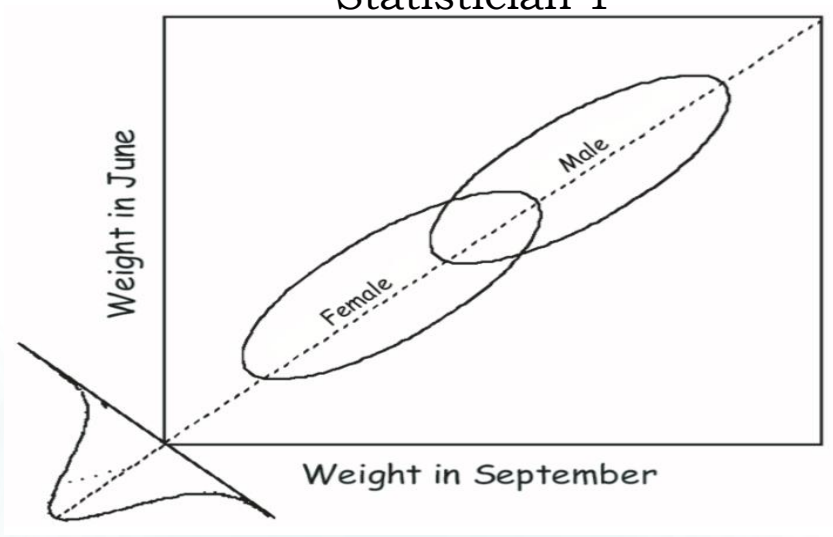
- Key features of the data:
  - For the males, the whole distribution of weights is unchanged
  - The same thing is true for the females. The only difference is that the females started and ended the year lighter on average than the males.
  - On average, there is no weight gain or weight loss for either males or females.



Lord’s Paradox

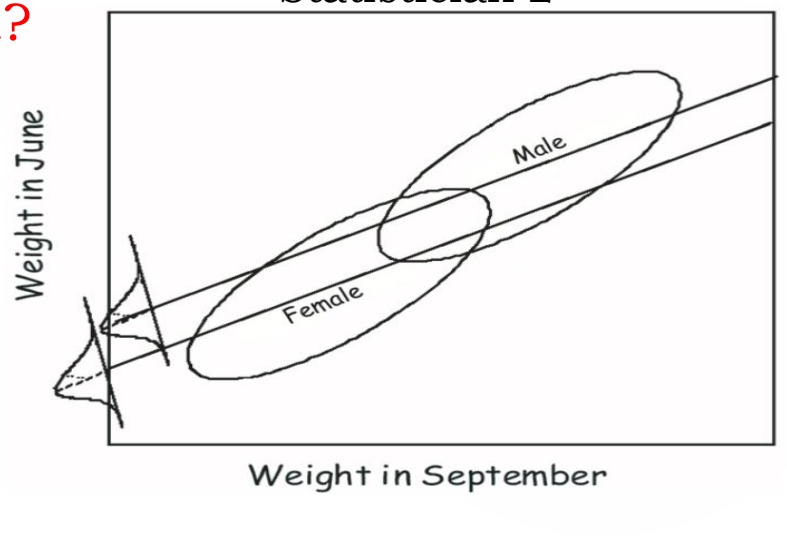
Two hypothetical statisticians got opposite conclusions for the same data

Statistician 1



Who is correct?

Statistician 2



- Observes that there are no differences between the September and June weight distributions for either males or females.
- Concludes that there is no evidence of any interesting effect of diet (or of anything else) on student weight.

- Uses “parallel regression adjustment” to fit the data
- Conclude that after “controlling for” initial weight, the diet has a differential positive effect on males relative to females

Lord's Paradox

Explain of the paradox:

- If both statisticians confined their comments to describing the data, both would be correct
- But, for causal inference, both are wrong because these data cannot support any conclusions about the causal effect of the diet without making some very strong, and arguably implausible, assumptions.

The Neyman-Rubin framework

Unit: the students

Treatment: the university diet (Sep. - June) ⟹ Potential outcome Y_t (The observed weights in June)

Control: not defined at all ⟹ Potential outcome Y_c ⇢ ?

Covariate: the weights in Sep.

Assignment mechanism: assigned everyone to the treatment

Nobody received the "control treatment" at all!

Summary



Summary

- What we’ve learned?
  - Concepts
    - potential outcomes, causal effect, assignment mechanism
  - Goal: causal estimand for causal effect
  - Tools:
    - assumption: SUTVA, about assignment mechanism
    - covariates
  - Example – Lord’s paradox: the importance of clarifying potential outcomes.
- What to gain?
  - identify formal / informal (correct / incorrect) causal statement,
  - build up causal framework facing new problem.



Causality in Statistics: a Brief History

- The two most important early developments
  - the introduction of potential outcomes in randomized experiments by Neyman (Neyman, 1923)
  - the introduction of randomization as the "reasoned basis" for inference by Fisher (Fisher 1935, p. 14)
- Rubin's contributions for observational studies
  - puts the potential outcomes center stage in the analysis of causal effects for both randomized experiments and observational studies (Rubin, 1974)
  - discusses the benefits of randomization in terms of eliminating systematic differences between treated and control units, and formulates the assignment mechanism in general mathematical terms (Rubin 1974, 1975, 1978)



Reference

- Reference: [T1] Chapter 1 & 2
- Further reading:
  - Rubin, D. B. (1974), “Estimating Causal Effects of Treatments in Randomized and Non-randomized Studies,” Journal of Educational Psychology, Vol. 66: 688 – 701.
  - Lord, F. M. (1967), “A Paradox in the Interpretation of Group Comparisons”, Psychological Bulletin, Vol. 68: 304-305.

