The Salk vaccine field trials: The first large-scale randomized experiment

Unit 5 Lecture 1

Jonathan Auerbach
STAT 489 Pre-Cap Prof Development
jauerba@gmu.edu

November 4, 2021

Learning Objectives

After this lecture, you will be able to:

1. Describe how the National Foundation for Infantile Paralysis conducted field trials to determine the efficacy of Salk's vaccine.
2. Describe the Fundamental Problem of Causal Inference and explain how it is overcome by randomization.
3. Estimate the Average Treatment Effect and calculate its standard error.
4. Display data in a table using `kable`.

These slides use the following `R` packages

Setup:

```
library("knitr")
library("kableExtra")
library("HistData")
library("tidyverse")
theme_set(theme_bw())
```

The Salk vaccine field trials were the first large-scale randomized experiment

- Jonas Salk developed his vaccine for Poliomyelitis (Polio) in 1953.
  - Polio paralyzed and killed children in inexplicable, epidemic waves.
  - The disease was relatively rare but affluent children more likely to exhibit severe symptoms—most famously Franklin D. Roosevelt.
- The National Foundation for Infantile Paralysis (NFIP) tested Salk's vaccine in highly publicized field trials to inspire public confidence.
  - The large-scale randomized controlled trial demonstrated Salk's vaccine halved the rate of Polio from 2/2500 to 1/2500.
- Randomization was convincing because it eliminated competing theories that might explain the apparent effectiveness of the vaccine
  - e.g. Could vaccinated patients have been healthier at baseline? No. Randomization means the average vaccinated patient was as healthy as the average control patient at baseline.
  - Though already known among statisticians, the trials helped establish randomization as the gold standard of scientific evidence.

Vaccination is an old medical practice

- Inoculation dates back to at least 10th century China.
  - Patients are exposed to a disease in order to stimulate the production of antibodies and induce immunity
  - Inoculation is risky—a patient can contract and spread the disease.
- In 1796, Edward Jenner introduced the first vaccine.
  - Vaccines inoculate with comparably safe forms of the disease, typically killed or weakened.
  - Jenner inoculated patients from smallpox with cowpox—a related but relatively harmless disease. Vaccinus Latin for "derived from cow"
- Salk created his vaccine from a killed virus, which was easier to develop than a live virus vaccine but could have been ineffective.
  - Nevertheless, Salk was so confident in his vaccine that he announced success on the radio in 1953 before independent review—much to the annoyance of his peers.

Jonas Salk's Polio vaccine (1957)

Source: https://www.history.com/news/8-things-you-may-not-know-about-jonas-salk-and-the-polio-vaccine

Salk vaccine field trials were highly publicized

- After Salk's announcement, the National Foundation for Infantile Paralysis (NFIP) organized field trials to test the vaccine's efficacy.
- The trials were designed to inspire public confidence in the vaccine.
  - Previous vaccines had been withdrawn because they caused Polio.
  - The NFIP field trials were the largest experiment in human history at the time, including 1 million children and costing $5 million dollars (roughly $50 million today).
- The field trials were famous in part because NFIP had a famous fundraising campaign—the March of Dimes.
  - Starting during the Depression, the campaign convinced everyday Americans to send more than seven billion dimes to the White House—a portion of which was devoted to research.
  - Ads were shown in movie theaters, which were extremely popular at the time. They included stars like Judy Garland and Elvis Presley.

NFIP considered three different experiments:

1. Vital Statistics Approach
   - design: assign everyone vaccine and observe whether Polio cases fall
   - potential problems: Polio fluctuates season-season/year-year
2. Observed Control Approach (Observed Control)
   - design: assign second graders the vaccine and use first and third graders as a control
   - potential problems:
     a. Volunteer effect. Some second graders may refuse to participate.
     b. Diagnostic bias. Knowing whether the student was vaccinated may change the diagnosis.
3. Placebo Control Approach (Randomized Control)
   - design: randomly assign participants vaccine or placebo (control). Avoids above problems:
     a. Second graders may refuse to participate—but such students are equally likely to be assigned vaccine or placebo (focus of this lesson).
     b. Assignment double blind. Neither student nor doctor knows who is vaccinated.

Factors that influenced the decision to participate according to Francis et al. (1955)

1. The frequency of vaccination against small pox, diphtheria, and whooping cough strongly correlated with participation.
2. Participants more frequently stated that "shots always work" than non-participants.
3. Mothers of participants were more likely to spend two or more evenings a week in outside activities than were mothers of non-participants.
4. Mothers of participants were more likely to have completed high school than mothers of non-participants.
5. A much smaller percentage of participants had family incomes under $4,500. Participation rate increased steadily with increasing income
6. The interviewer's rating of the quality of the respondent's neighborhood and condition of his house was highly correlated with participation status. Participants lived in better neighborhoods, and their homes were better kept.

Results of the Salk field trials (detailed)

```
PolioTrials %>%
  kable(format.args = list(scientific = F, big.mark = ","),
                              booktabs = T) %>%
  kable_styling(font_size = 6, position = "center")
```

Experiment	Group	Population	Paralytic	NonParalytic	FalseReports
RandomizedControl	Vaccinated	200,745	33	24	25
RandomizedControl	Placebo	201,229	115	27	20
RandomizedControl	NotInoculated	338,778	121	36	25
RandomizedControl	IncompleteVaccinations	8,484	1	1	0
ObservedControl	Vaccinated	221,998	38	18	20
ObservedControl	Controls	725,173	330	61	48
ObservedControl	Grade2NotInoculated	123,605	43	11	12
ObservedControl	IncompleteVaccinations	9,904	4	0	0

Results of the Salk field trials (summary, long)

```
salk <- PolioTrials %>%
  mutate(Cases = Paralytic + NonParalytic + FalseReports,
         Pop   = Population,
         Prop  = Cases / Pop,
         Group = if_else(Group == "Placebo", "Controls",
                         as.character(Group))) %>%
  select(Experiment, Group, Pop, Prop) %>%
  filter(Group %in% c("Vaccinated", "Controls"))

salk %>%
  kable(format.args = list(scientific = F, big.mark = ","),
        digits = 4, booktabs = T) %>%
  kable_styling(position = "center")
```

Experiment	Group	Pop	Prop
RandomizedControl	Vaccinated	200,745	0.0004
RandomizedControl	Controls	201,229	0.0008
ObservedControl	Vaccinated	221,998	0.0003
ObservedControl	Controls	725,173	0.0006

Results of the Salk field trials (summary, wide)

```
salk <- salk %>% select(Experiment, Group, Prop, Pop) %>%
   nest(data = c(Prop, Pop)) %>%
   spread(key = Group, value = data) %>%
   unnest(Vaccinated, Controls, names_sep = ' ')

salk %>%
  kable(format.args = list(scientific = F, big.mark = ","),
        digits = 4, booktabs = T) %>%
  kable_styling(font_size = 7, position = "center")
```

Experiment	Controls Prop	Controls Pop	Vaccinated Prop	Vaccinated Pop
ObservedControl	0.0006	725,173	0.0003	221,998
RandomizedControl	0.0008	201,229	0.0004	200,745

- Polio twice as likely in vaccinated as control for both experiments. Thus both trials suggests the risk of Polio is halved by the vaccine.
- But Polio is fifty-percent more likely among randomized control trial participants than observational control trial participants. Thus the estimated number of cases prevented is larger.

Randomized trial suggests more cases prevented than obervational—but not statistically significant

```
salk %>%
  mutate(effect = `Vaccinated Prop` - `Controls Prop`,
         se_vac = sqrt(`Vaccinated Prop` *
            (1 - `Vaccinated Prop`) / `Vaccinated Pop`),
         se_cnt = sqrt(`Controls Prop` *
            (1 - `Controls Prop`) / `Controls Pop`),
         se_effect = sqrt(se_vac^2 + se_cnt^2),
         `lower limit` = effect - 2 * se_effect,
         `upper limit` = effect + 2 * se_effect) %>%
  select(Experiment, effect,
         `lower limit`, `upper limit`)%>%
  kable(digits = 5, booktabs = T) %>%
  kable_styling(position = "center")
```

Experiment	effect	lower limit	upper limit
ObservedControl	-0.00026	-0.00036	-0.00017
RandomizedControl	-0.00040	-0.00055	-0.00024

Did NFIP prove Salk's vaccine was effective?

- Although the randomized control and observational control trials agreed that vaccinating reduced the proportion of Polio cases by half, the randomized control trial is particularly convincing.
  - Randomization eliminates competing theories that might explain the apparent effectiveness of the vaccine in the observational control.
  - e.g. Could vaccinated patients have been healthier at baseline? No. Randomization means the average vaccinated patient was as healthy as the average control patient at baseline.
- Thus, NFIP established effectiveness in the following sense:
  - If all participants were vaccinated, there would be half as many cases than if all participants were given the control.
  - Statisticians often use potential outcomes notation to make this point mathematically precise.

Causal inferences compare potential outcomes

Subject ID i	Vaccine $Y_i(1)$	Placebo $Y_i(0)$	Effect $Y_i(1) - Y_i(0)$
1	1	0	1
2	1	1	0
3	1	0	1
4	0	1	-1
⋮	⋮	⋮	⋮
401,974	0	0	0

- Each patient has two potential treatments: take vaccine or placebo.
- Thus each patient has two potential outcomes: the outcome if the patient takes the vaccine, and the outcome if the patient takes the placebo. (n.b. for each table entry, 1 indicates patient develops Polio, 0 patient does not—data are not real and are just for illustration.)
- A causal effect is a comparison of potential outcomes.

The fundamental problem of causal inference

Subject ID i	Vaccine $Y_i(1)$	Placebo $Y_i(0)$	Effect $Y_i(1) - Y_i(0)$
1	1	?	?
2	1	?	?
3	?	0	?
4	?	1	?
⋮	⋮	⋮	⋮
401,974	0	?	?

- In practice scientists usually only observe one potential outcome. A patient cannot both take and not take the vaccine.
- This is called the "fundamental problem of causal inference."
- This complicates inference. If only patients who never get Polio take the vaccine, the vaccine will look effective—even if it's not.

Randomization to estimate the average causal effect

Subject ID i	Vaccine $Y_i(1)$	Placebo $Y_i(0)$	Effect $Y_i(1) - Y_i(0)$
1	1	?	?
2	1	?	?
3	?	0	?
4	?	1	?
⋮	⋮	⋮	⋮
401,974	0	?	?
Average:	$^1/_{2500}$	$^2/_{2500}$	$-^1/_{2500}$

- If you randomly assign some patients vaccine, their outcomes will represent the population if vaccinated.
- If you randomly assign other patients placebo, their outcome will represent the population if not vaccinated.
- We can measure the effect of vaccine ON AVERAGE OVER POPULATION by randomizing and comparing groups.

NFIP proved effectiveness in a specific sense

- Randomization made the vaccinated and placebo groups "representative" of a hypothetical population in which everyone vaccinated AND everyone not vaccinated.
  - NFIP established that the vaccine reduced cases on average over this hypothetical population by ${}^{1}/_{2500} - {}^{2}/_{2500} = -{}^{1}/_{2500}$.
- The average outcome if everyone got the vaccine minus the average outcome if everyone did not get the vaccine is called the Average Treatment Effect (ATE).
  - Another common statistic is the ratio of these averages.
  - e.g. Salk's vaccine ${}^{1}/_{2500} \div {}^{2}/_{2500} = 50\%$ effective against Polio.
- Potential outcomes were first used to study randomized experiments by Jerzy Neyman (1923).
  - They were subsequently developed by Donald Rubin (1974).
  - Many researchers refer to the use of potential outcomes as the Neyman–Rubin Causal Model.

Jerzy Neyman (1955) and Donald Rubin (2022)

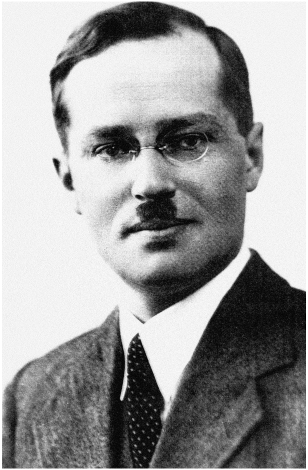

Source: http://magazine.amstat.org/wp-content/uploads/2017/04/Neyman.jpg
https://statistics.fas.harvard.edu/people/donald-b-rubin

Randomized controlled trials are the gold standard—but not perfect

- Randomized controlled trials are limited by their design.
  - Results may only hold for individuals like those who participated in the trials.
  - Experimental conditions may deviate from real life. Patients may behave differently in the context of an experiment. Vaccines produced after the trial may also be different than those used in the trial. (e.g. the Cutter Incident)
  - Bias can be reintroduced post-randomization. Double blind helps eliminate some of these sources.

- Conclusive evidence comes from a combination of experimental and observational studies.
  - Vaccination has largely eradicated once common diseases such as Smallpox, Polio, and Measles.
  - Links between vaccines and autism, for example, have been debunked.

References

1. Brownlee, K. A. Statistics of the 1954 Polio Vaccine Trials, Journal of the American Statistical Association, 50, 1005-1013. 1955.

2. Francis T, Jr, Korns R, Voight R, Boisen M, Hemphill F, Napier J, et al. An Evaluation of the 1954 Poliomyelitis Vaccine Trials: Summary Report. Am J Public Health, 45, 1–50. 1955.

3. Friendly, Michael. HistData: Data sets from the history of statistics and data visualization. R package version 0.7-5. 2014.

4. Meier, Paul. The Biggest Public Health Experiment Ever: The 1954 Field Trial of the Salk Poliomyelitis Vaccine. Statistics: A Guide to the Unknown. Second Edition. pg. 3-15. 1974.

5. Thompson, Dennis. CBS News. The Salk polio vaccine: Greatest public health experiment in history. December 2, 2014. https://www.cbsnews.com/news/the-salk-polio-vaccine-greatest-public-health-experiment-in-history/