Estimands and Causal Inference

Daniel Scharfstein
Johns Hopkins University
dscharf@jhu.edu

November 2, 2017

Disclosures

- Regularly consult with pharmaceutical and device companies
  - DSMB's
  - Mock advisory panels
  - Data analysis
  - Study design
  - Training
  - Litigation
- Served on the National Academy of Sciences (NAS) Panel that issued the report "The Prevention and Treatment of Missing Data in Clinical Trials"
- Funding from NIH, PCORI and FDA for the development of statistical methods and associated software.

Acknowledgements

- Miguel Hernan, Harvard University
- Michael Hudgens, University of North Carolina

Estimand

- A population parameter that quantifies the *effect* of treatment relative to control.
- NAS Report
  - Target of inference in a randomized clinical trial
  - *Causally* interpretable
  - Motivated the ICH Addendum
- ICH Addendum
  - Avoids the word "causal"
  - But uses the language of causal inference
  - *"What would have happened to the same subjects under different treatment conditions?"*

Causal Inference

- Mathematical language for
  - defining causal estimands
  - formulating identification assumptions
- Extension of standard mathematical language of statistics

Outline

- Casual Inference
  - Intuitive Definition of Cause
  - Mathematical Formalism
  - Individual-Level Causal Effects
  - Population-Level Causal Effects
  - Association vs. Causation
- Randomized Experiments
  - Ideal Experiments
  - Intercurrent Events
- Examples of Estimands
- Discussion

Intuitive Definition of Cause

- Ian took the red pill on September 1, 2015
  - Five days later, he died
- Had Ian not taken the red pill (all other things being equal)
  - Five days later, he would have been alive
- Did the red pill have a causal effect on Ian's death?

Intuitive Definition of Cause

- Jim did not take the red pill on September 1, 2015
  - Five days later, he was alive
- Had Jim taken the red pill (all other things being equal)
  - Five days later, he would have been alive
- Did the red pill have a causal effect on Jim's death?

Human Reasoning

- We compare (often mentally)
  - the outcome when action A is present with
  - the outcome when action A is absent
  - assuming other things being equal
- If the two outcome differs, we say that action A has a causal effect
- A is commonly referred to as **exposure** or **treatment**

Mathematical Notation for Observed Data

Binary treatment and outcome

- $A = 1$: individual is treated with red pill, $A = 0$ otherwise
  - Ian: $A_i = 1$; Jim $A_j = 0$
- $Y = 1$: individual died, $Y = 0$ otherwise
  - Ian: $Y_i = 1$; Jim $Y_j = 0$

ID	A	Y
Ian	1	1
Jim	0	0

Notation for Complete Data

- $Y(0) = 1$: individual dies if untreated, $Y(0) = 0$ lives if untreated.
    - Ian: $Y_i(0) = 0$; Jim: $Y_j(0) = 0$
- $Y(1) = 1$: individual dies if treated, $Y(1) = 0$ lives if treated.
    - Ian: $Y_i(1) = 1$; Jim: $Y_j(1) = 0$

ID	A	Y	$Y(0)$	$Y(1)$
Ian	1	1	0	1
Jim	0	0	0	0

Individual-level Causal Effects

- Ian
  - The red pill has a causal effect because $Y_i(1) \neq Y_i(0)$
- Jim
  - The red pill has a no causal effect because $Y_j(1) = Y_j(0)$

Individual-level (sharp) causal null hypothesis:

$$Y_k(1) = Y_k(0) \text{ for all } k$$

Potential vs. Counterfactual Outcomes

- Potential outcomes: $Y(0)$, $Y(1)$
- $Y(0)$ and $Y(1)$ should be viewed as fixed attributes
- Counterfactual outcomes: outcome under treatment not actually received.
    - If $A = 1$, $Y(0)$ is counterfactual
    - If $A = 0$, $Y(1)$ is counterfactual

Consistency Assumption

$$Y = AY(1) + (1 - A)Y(0)$$

- If $A = 1$ then $Y = Y(1)$.
- If $A = 0$ then $Y = Y(0)$.
- "There are a number of ways that treatment could have been assigned but all those ways would have resulted in the same outcome."

Observed Data

ID	A	Y	$Y(0)$	$Y(1)$
Ian	1	1	?	1
Jim	0	0	0	?
Ken	1	0	?	0
Leo	0	1	1	?
Mike	1	1	?	1
Nick	0	0	0	?
...				

Fundamental Problem of Causal Inference

- Individual causal effects cannot generally be determined (except in cross-over experiments with strong assumptions)
- Causal inference is a missing data problem

Causal Inference

- $P[Y(a) = 1]$: proportion of individuals in the population who would have developed the outcome had everyone received treatment a
- Language used in the ICH Addendum
- Marginal probability, **not conditional**

Population-Level Causal Effect (Estimand)

Casual effect at the population-level:

$$P[Y(1) = 1] \neq P[Y(0) = 1]$$

Population-level (average) causal null hypothesis:

$$P[Y(1) = 1] = P[Y(0) = 1]$$

Causal effects can be measured on many scales

- Casual risk difference: $P[Y(1) = 1] - P[Y(0) = 1]$
- Causal risk ratio: $P[Y(1) = 1]/P[Y(0) = 1]$
- Causal odds ratio: $\frac{P[Y(1)=1]}{P[Y(1)=0]} / \frac{P[Y(0)=1]}{P[Y(0)=0]}$

Identification of Causal Effect

- Identification means that the quantity of interest can be determined mathematically from the distribution of the observed data.
- Population-level causal effect can be identified under no assumptions in ideal randomized studies (more later).

Association

- $P[Y = 1|A = a]$: proportion of individuals who developed the outcome among those who received treatment a
- Conditional probability
- Treatment A and outcome Y are said to be associated if

$$P[Y = 1|A = 1] \neq P[Y = 1|A = 0]$$

- No association (independence)

$$P[Y = 1|A = 1] = P[Y = 1|A = 0]$$

- Independence: $A \perp Y$
- Associaton can be measured on many scales
  - Risk difference: $P[Y = 1|A = 1] - P[Y = 1|A = 0]$
  - Risk ratio: $P[Y = 1|A = 1]/P[Y = 1|A = 0]$
  - Odds ratio: $\frac{P[Y=1|A=1]}{P[Y=0|A=1]} / \frac{P[Y=1|A=0]}{P[Y=0|A=0]}$

“Association is NOT Causation”

- Association: measures difference in risk between disjoint subsets of the population determined by individual’s actual treatment value
- Causation: measures difference in risk in the *entire population* under two treatment values

Association vs. Causation"

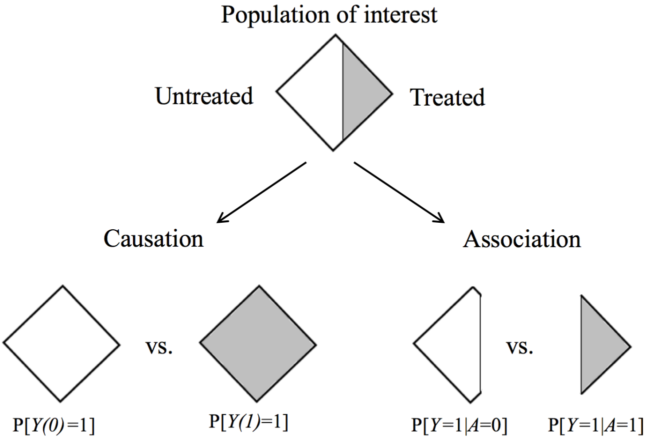

Statistics and Causation

- Need potential outcomes/counterfactuals to talk about causation
- Otherwise, statistics is a language for association not causation
- Causal concepts cannot be represented using purely statistical language
- For example, confounding occurs when

$$P[Y(a) = 1] \neq P[Y = 1 | A = a]$$

Experiments

- **Experiment:** A scientific study in which the investigators intervene in the assignment of treatment to the individuals participating in the study. Also known as "clinical trial" when the goal is to study the effects of medications or devices in humans.
- **Randomized Experiment:** An experiment in which the investigators use a random procedure to allocate treatment, e.g., flip of a coin, computer-generated random number

Ideal Randomized Experiment

- No loss to follow-up
- Full compliance with (adherence to) assigned treatment
- One version of treatment
- Double blind assignment: neither the study participants nor the investigators know who is receiving which treatment

Ideal Randomized Experiment

$$A \perp Y(a) \text{ for } a = 0, 1$$

- $P[Y(1) = 1] = P[Y = 1|A = 1]$
- The marginal distribution of $Y(1)$ is equal to the conditional distribution of Y given $A = 1$
- $P[Y(0) = 1] = P[Y = 1|A = 0]$
- The marginal distribution of $Y(0)$ is equal to the conditional distribution of Y given $A = 0$
- Individuals in the treatment groups are exchangeable.
- This does not mean $A \perp Y$.

Randomized Experiments in the Real World

- Intercurrent Events
  - Initiation of rescue medication
  - Treatment switching
  - Stopping treatment
  - ...

Example: Rescue Medication

- Some individuals initiate rescue medication prior to time of outcome measurement
- Everything else ideal

Notation

- A is treatment assignment
- $R(a)$ is the use of rescue medication under treatment a
- $R = R(A)$ is observed use of rescue medication
- $Y(a, r)$ is outcome under treatment a and rescue medication use r
- $Y(a) = Y(a, R(a))$ is the outcome under treatment a
- $Y = Y(A)$ is observed outcome

Example

ID	A	R	Y	$R(0)$	$R(1)$	Y(0,0)	Y(0,1)	Y(1,0)	Y(1,1)	Y(0)	Y(1)
Fay	0	0	1	0	?	1	?	?	?	1	?
George	0	0	0	0	?	0	?	?	?	0	?
Tom	0	0	0	0	?	0	?	?	?	0	?
Mary	0	1	1	0	?	?	1	?	?	1	?
Chris	0	1	0	0	?	?	0	?	?	0	?
Anna	0	1	1	0	?	?	1	?	?	1	?
Adam	0	1	1	1	?	?	1	?	?	1	?
John	0	1	0	1	?	?	0	?	?	0	?
Ian	0	1	0	1	?	?	0	?	?	0	?
Rose	1	0	1	?	0	?	?	1	?	?	1
Jack	1	0	0	?	0	?	?	0	?	?	0
Lee	1	0	0	?	0	?	?	0	?	?	0
Betsy	1	0	0	?	1	?	?	0	?	?	0
Claus	1	0	0	?	1	?	?	0	?	?	0
Sara	1	0	1	?	1	?	?	1	?	?	1
Lisa	1	1	1	?	1	?	?	?	1	?	1
Peter	1	1	0	?	1	?	?	?	0	?	0
Sue	1	1	0	?	1	?	?	?	0	?	0

Randomization

$$A \perp R(a) \text{ for } a = 0, 1$$

$$A \perp Y(a, r) \text{ for } a, r = 0, 1$$

Estimands

- Intention to Treat (Treatment Policy)
- Hypothetical (Prescriptive)
- Hypothetical (Natural)
- Principal Stratum (PS)
- Composite
- Per-protocol Effect

Intention to Treat (Treatment Policy)

$$P[Y(1) = 1] \text{ vs. } P[Y(0) = 1]$$

- Measures the causal effect of being assigned to treatment vs. control
- No additional assumptions required
- Can be impacted by intercurrent events.
- For example, suppose rescue medication is effective and more patients on control initiate rescue medication

Hypothetical (Prescriptive)

Intervene on intercurrent event

$$P[Y(1,1)=1] \text{ vs. } P[Y(0,1)=1]$$

$$P[Y(1,0)=1] \text{ vs. } P[Y(0,0)=1]$$

- Is it realistic/ethical to intervene?
- Requires additional untestable assumptions, e.g.,

$$R \perp Y(a,r) \mid A, X$$

where X is baseline covariates.

Estimand - Hypothetical (Natural)

Intervene on intercurrent event in treatment group by setting rescue medication use to its natural level under the control condition

$$P[Y(1, R(0)) = 1] \text{ vs. } P[Y(0, R(0)) = 1]$$

- Is it realistic/ethical to intervene?
- Requires additional untestable assumptions

Estimand - Principal Stratum

Stratum based on potential intercurrent outcomes, e.g.,

$$P[Y(1) = 1 | R(1) = 0, R(0) = 0] \text{ vs. } P[Y(0) = 1 | R(1) = 0, R(0) = 0]$$

$$P[Y(1) = 1 | R(0) = 0] \text{ vs. } P[Y(0) = 1 | R(0) = 0]$$

$$P[Y(1) = 1 | R(0) = 1] \text{ vs. } P[Y(0) = 1 | R(0) = 1]$$

- Unknown stratum membership, clinical relevance (?)
- Requires additional untestable assumptions

Estimand - Composite

Composite outcome $U(a)$ which takes on the value 1 if $R(a) = 0$ and $Y(a) = 0$

$$P[U(1) = 1] \text{ vs. } P[U(0) = 1]$$

- Mixes the effect of treatment on (1) rescue medication use and (2) the outcome; clinical relevance (?)
- ITT effect on composite outcome, no assumptions

Estimand - Per Protocol Effect

- Suppose the protocol prescribes that if an intercurrent V occurs then rescue medication is allowed, otherwise it not.
- Further, suppose some patients use rescue medication when V does not occur (i.e., non-compliance with protocol)

$$Z(a) = V(a)Y(a,1) + \{1 - V(a)\}Y(a,0)$$

$$P[Z(1) = 1] \text{ vs. } P[Z(0) = 1]$$

- Not the same as the hypothetical where one intervenes on rescue medication for all patients
- Should not be confused with per-protocol analysis
- Requires additional untestable assumptions

Take Home Messages

- Estimands should not be *conditional* on post-randomization events - apples vs. oranges
- Marginal quantities
  - apply to all patients in the population
- Conditional on pre-randomization covariates
  - apply to subgroup of the population
- Choice of estimand depends on context and perspective
- Multiple estimands may be important for understanding treatment effects
- Deciding on proper estimands requires close interaction between statisticians, clinicians and decision makers.
- Sample size/design will depend on choice of estimand(s).

Take Home Messages

- Potential/counterfactual outcomes is a mathematical language for formalizing estimands and ultimately understanding identification assumptions.
- Untestable assumptions beyond randomization may be required
  - may be so strong that estimand might not be ideal
  - can inform design/data collection
- When assumptions are required, the robustness of inference should be checked via sensitivity analysis.

Modern History of Causal Inference



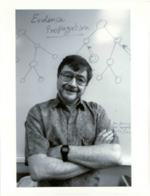

Resources

- Causality - Pearl
- Causal Inference - Hernan, Robins
- Counterfactuals and Causal Inference - Morgan and Winship
- Causal Inference for Statistics, Social Sciences and Biomedical Sciences - Imbens, Rubin
- Applied Bayesian Modeling and Causal Inference from an Incomplete Data Perpsective - Gelman, Meng
- Observational Studies - Rosenbaum
- Blogs - Pearl (UCLA), Gelman (Columbia)

Discussion Points

- Relative merits of different estimands from a regulatory perspective.
- Inference about some estimands require untestable assumptions.