Toward Causal Inference in Networks

Ashley L. Buchanan

Spring 2019

An Introduction for DATA 2080: Data and Society

Thanks to my advisor, Dr. Michael Hudgens, at the University of North Carolina Chapel Hill for materials used to develop this lecture!

Outline

Potential Outcomes

- Dichotomous (binary) treatment A taking values 0 (untreated), 1 (treated)
- Dichotomous outcome Y taking values 0, 1
- Potential outcome $Y^{a=1}$ outcome that would have been observed if, possibly counter to fact, treatment $a = 1$ was received
- Potential outcome $Y^{a=0}$ defined similarly
  aka counterfactuals
- E.g., A is receive heart transplant or not, Y is die or not 5 days later
- Assumptions: Consistency, exchangeability, and positivity

Causation Versus Association

- Q: Under what circumstances does association imply causation?
- Figure 1.1 (Hernan and Robins)

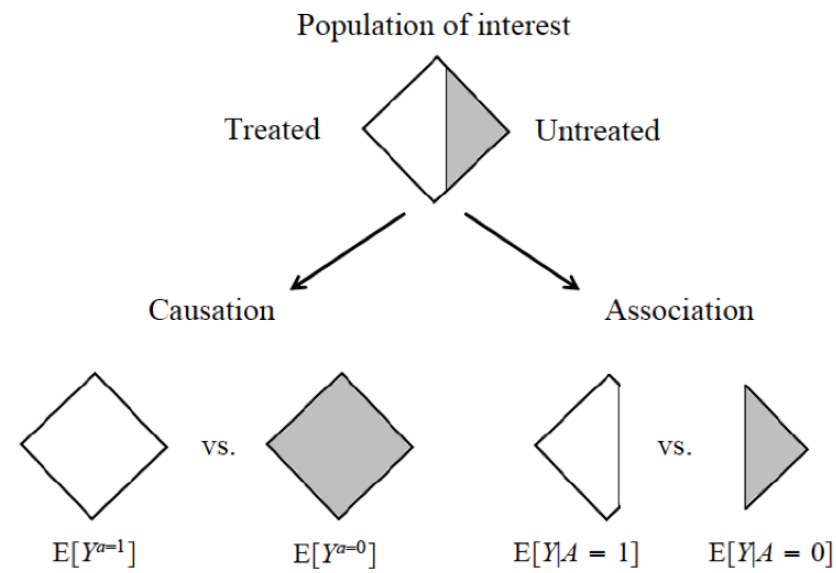

Randomized Experiments and Exchangeability

- Consider a randomized experiment where individuals are assigned treatment $a = 1$ or $a = 0$ randomly, i.e., independent of their potential outcomes. Thus, marginal exchangeability holds:

$$Y^a \perp\!\!\!\perp A \text{ for } a = 0, 1$$

- Conditionally randomized experiments can be viewed as two separate marginal experiments. Thus, conditional exchangeability holds

$$Y^a \perp\!\!\!\perp A | L = l \text{ for } a, l = 0, 1$$

- Causal effects can be identified under conditional exchangeability, e.g., using standardization or inverse probability weighting

Randomized Experiment Paradigm

An observational study can be conceptualized as a conditionally randomized experiment under the following three conditions:

1. Values of treatment under comparison correspond to well-defined interventions (Sections 3.4-3.5)
2. Conditional probability of receiving every value of treatment, though not decided by investigators, depends only on the measured covariates (Section 3.2)
3. Conditional probability of receiving every value of treatment is greater than zero, i.e., positive (Section 3.3)

Causation and Prediction

- What if we are not willing to assume positivity, conditional exchangeability, and well-defined interventions?
- There is always prediction (association)
  Obesity (A) may predict (or be associated with) mortality risk (Y)
- This does not imply causation, just as carrying lighter in pocket (A) being predictive of lung cancer (Y) does not imply carrying a lighter causes cancer
- However, associations may be helpful in generating hypothesis.
  Why is obesity associated with mortality risk? Is it diet? Exercise? Other well-defined interventions?
- From a public health or scientific viewpoint we may want to go beyond associations to attempt to understand why such associations exist

Introduction to Interference

- No Interference
  - Outcome of one individual assumed to be unaffected by the treatment assignment of others
  - Typical assumption of causal inference
  - Part of SUTVA
- Clearly not true in some settings
  - Infection diseases, education interventions. social sciences
- Phenomenon of interest vs. nuisance
- The following slide shows different possible effects of pre-exposure prophylaxis for HIV prevention adapted from Halloran and Struchiner (1991, 1995)

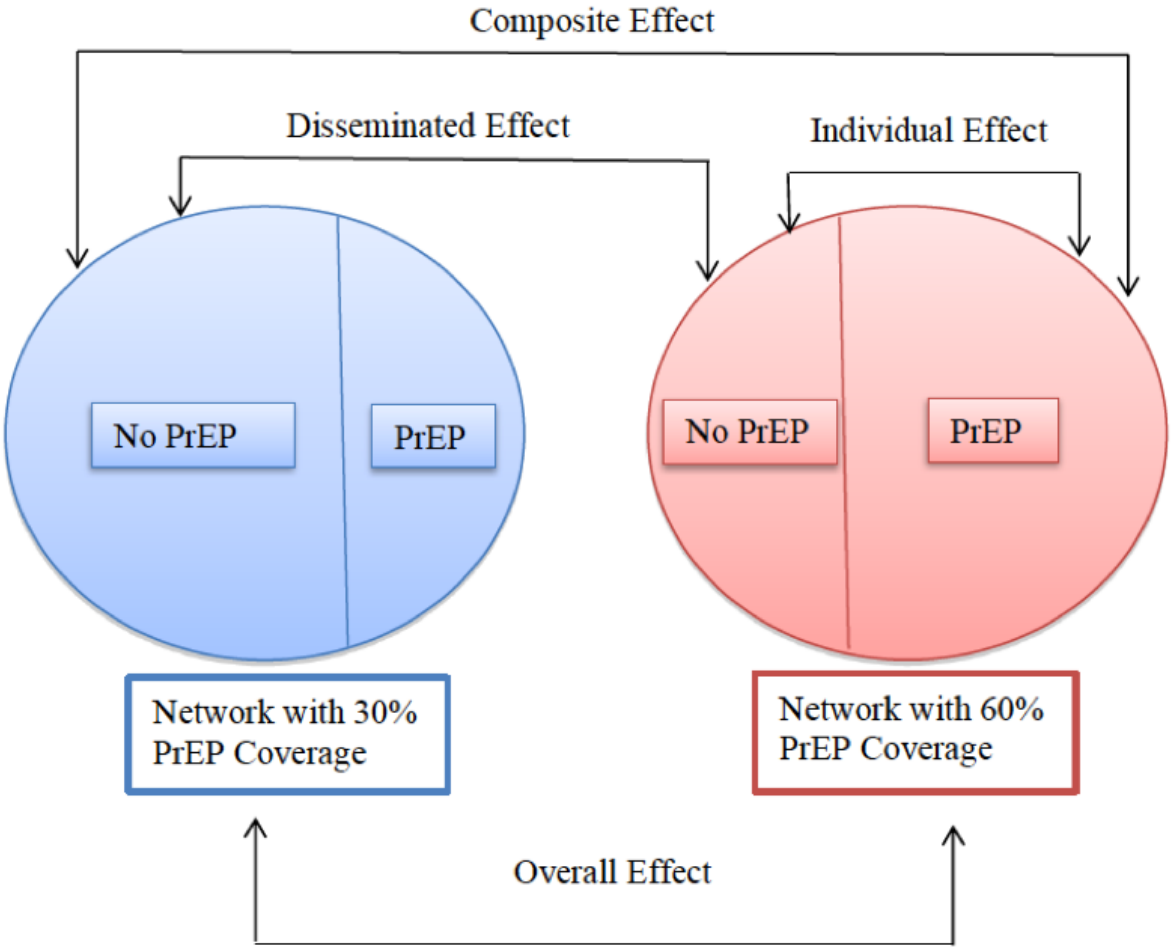
Composite Effect
Disseminated Effect
Individual Effect
No PrEP
PrEP
No PrEP
PrEP
Network with 30% PrEP Coverage
Network with 60% PrEP Coverage
Overall Effect

Interference Definition

- That two potential outcomes sufficiently represent all potential outcomes for an individual assumes no interference between individuals
  i.e., the treatment of one individual does not affect the outcome of other individuals (Cox 1958)
- The no interference assumption may not hold in some settings
- Examples: Vaccine studies, educational intervention studies, HIV prevention studies
- Settings: Epidemiology, medical research, econometrics, social network analysis

General Approach

- Population of groups of individuals (blocks of units)
- Assume partial interference: Possibly inteference between individuals in a group but not between groups.
- Define direct, disseminated (indirect), composite (total) causal effects
- Two-stage randomization
  1. Groups to allocation strategies α_1, α_0
  2. Given 1, individuals randomized to treatment/controls $A \in 0, 1$
- Unbiased estimators, variance using randomization-based inference or M-estimation

Illustrative Example

Two-stage randomized placebo-controlled cholera vaccine trial based on data from Ali et al. (2005)

	α	Vaccine ($X_{ij}=1$)		Placebo ($X_{ij}=0$)	
		Total $\sum_j X_{ij}$	Cases $\sum_j X_{ij} Y_{ij}$	Total $\sum_j (1-X_{ij})$	Cases $\sum_j (1-X_{ij}) Y_{ij}$
1	α_1	12541	16	12541	18
2	α_1	11513	26	11513	54
3	α_1	10772	17	25134	119
4	α_0	8883	22	20727	122
5	α_0	5627	15	13130	92

α_0 is the allocation strategy for the group that randomized 50% to the treatment.
α_1 is the allocation strategy for the group that randomized 30% to the treatment.

Estimates of population average effects per 1000 individuals per year

Effect	Parameter	Estimate	Estimated varaince
Direct	$\overline{CE}^{D}(\alpha_1)$	1.30	0.856
Direct	$\overline{CE}^{D}(\alpha_0)$	3.64	0.178
Indirect	$\overline{CE}^{I}(\alpha_1, \alpha_0)$	2.81	3.079
Total	$\overline{CE}^{T}(\alpha_1, \alpha_0)$	4.11	0.672
Overall	$\overline{CE}^{O}(\alpha_1, \alpha_0)$	2.37	1.430

- Indirect: 50% vaccine coverage results in 2.8 fewer cholera cases per 1000 unvaccinated individuals per year compared to 30% vaccine coverage
- Overall: 50% vaccine coverage results in 2.4 fewer cholera cases per 1000 individuals per year compared to 30% vaccine coverage

Introduction to Networks

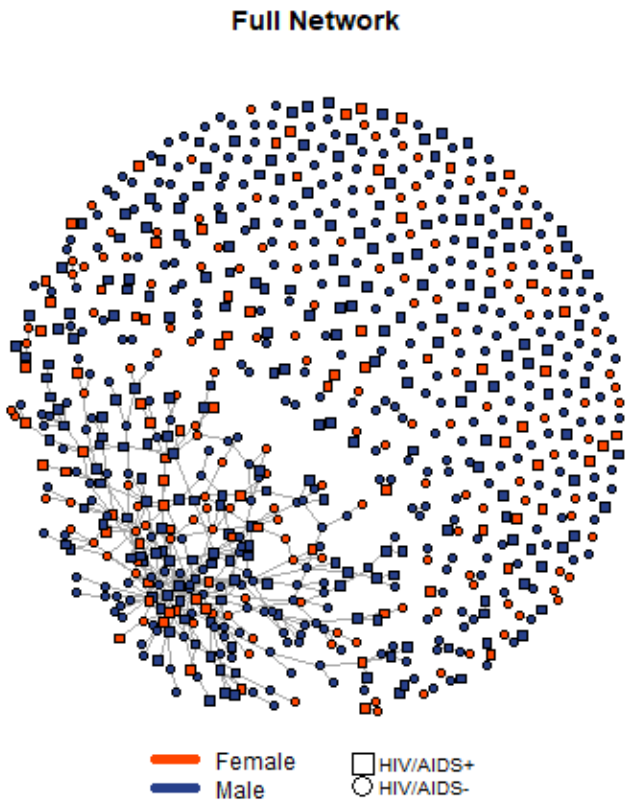

- Each node has an outcome, treatment and covariates
- Nodes connected through edges (e.g., sexual, drug use, social connections)
- Estimands: peer effects, treatment effects, spillover/interference effects, effects of network interventions

Approaches in the Literature

- Christakis and Fowler (2007, 2008, 2009, 2010, 2011, 2012) estimated peer effects in social network data
  - Model: $Y_{ego}^{t} \sim Y_{alter}^{t-1}, Y_{alter}^{t-2}, Y_{ego}^{t-2}, \mathbf{C}_{ego}$
  - Results included significant peer effects for obesity, smoking, alcohol consumption, etc.
  - Peer effects evaluated in other settings (Ali and Dwyer, 2009, Cacioppo et al, 2009; 2008; Lazer et al., 2010; Rosenquist et al, 2010, Wasserman, 2012)

Randomization-based Inference

- Randomization-based inference for networks (e.g., Toulis and Kao, 2013; Bowers et al., 2013; Aronow and Samii, 2013; Eckles et al., 2014, Choi 2016).
  - Assumes on finite population of N individuals and for each individual there is a set of individuals that may interfere with that individual (i.e., interference sets, neighborhoods, friends)
  - Interference sets can be represented by an adjacency matrix and often assumed to be known and fixed
- For an observational study, Tchetgen Tchetgen and VanderWeele (2012) suggest IPW estimator where all observations from group i are weighted by the inverse of probability of the treatment assignment vector $\mathbf{A}_i$ given $\mathbf{X}_i$
- Methods in the presence of interference often rely on randomization and the assumption of partial interference, but provides a solution to the problem of network dependence in cluster randomized trials (e.g. Sobel, 2006; Hong and Raudenbush, 2006; Rosenbaum, 2007; Hudgens and Halloran, 2008; Tchetgen Tchetgen and VanderWeele, 2012; Liu and Hudgens, 2014; Buchanan et al, 2018).

Non-randomized Interventions

- In many studies, the intervention or treatment is not randomized
- There may be confounding at either the individual, network-level or both
- Methods employ
  - A generalized propensity score (Forastiere, 2018) or a Bayesian generalized propensity score (Forastiere, 2018) that account for individual and neighborhood covariates
  - Targeted maximum likelihood estimation (TMLE) (Sofrygin, 2015)

Statistical Dependence in Networks

- Latent variables (i.e., homophily) lead to similar outcomes among close contacts
- Contagion: Indirect effect that treating one individual may have on another by preventing the treated individual from getting the disease and thereby from passing it on
- Networks often observed at a single time point, so difficult to disentangle homophily from an effect
- Why is this a problem?
  - We cannot assume independence (i.e., cannot assume independent and identically distributed (iid))
  - Central limit theorem may not hold
  - Standard error estimates and confidence intervals will be anti-conservative!

Possible Solutions for Dependence in Networks

- Create conditionally independent units; analyze with standard models, but conditional on information barriers (Ogburn and Vanderweele, 2017)
- Extension of influence function from iid setting with interference set (van der Laan, 2014) and social network setting with contagion and homophily (Ogburn, et al, 2017)
- Nearest neighbor approach: Potential outcomes of any individual only depends on their nearest neighbors (or two-step neighbors)
- Subsampling: Implementation and conditions may not be applicable to networks (e.g., bootstrap individual nodes)
- K-dependence: $Cov(W_i, W_j) = \sigma_k$, where $k = ||i, j||$ and estimate using a plug-in estimator

Nearest Neighbor Approach

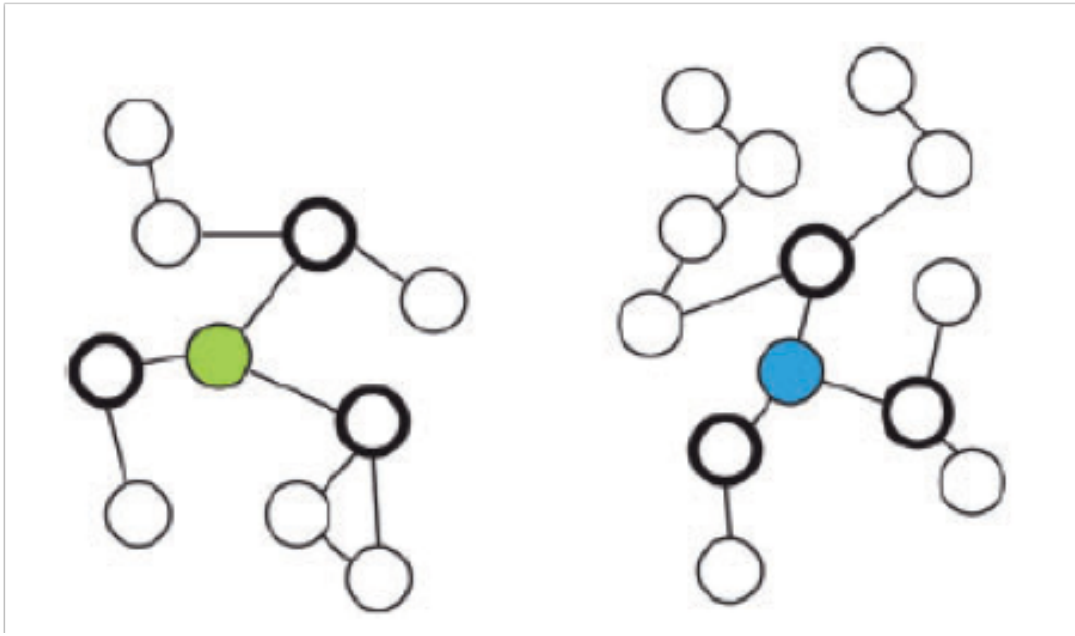

Diagram of the network disseminated effect with intervention subnetwork (left) and control subnetwork (right) (Index shaded green or blue) (Benjamin-Chung et al, 2017)

Motivating Example with Community Detection

The Social Factors and HIV Risk Study (SFHR)

- Sociometric network study conducted between 1991 and 1993 in Bushwick, Brooklyn, New York among *street-recruited injection drug users*
- Investigated how HIV/AIDS infection spread through shared sexual and injection risk behaviors.
- 767 participants along with 3,162 dyadic relationships (i.e. a connection b/w two people).
- Connections were shared **risk behaviors** (i.e. inject drug together and/or having sexual intercourse) within 30 days before the interview.
- Assess attitudes toward HIV/AIDS risk among PWIDs and their effect on health-seeking behaviors.

Causal Inference with Observational Study

- Tchetgen Tchetgen and VanderWeele (2012) developed an **inverse probability weighting (IPW)** method to replicate two-stage randomized design with observational study.
- We apply this method to evaluate how PWIDs locus of control/ blame attitudes affect on both their own health-seeking behavior and that of other members in their subnetworks/ communities.

APPROACH

1. Split the SFHR network into smaller subnetworks/ communities of PWIDs.
2. Calculate group-level propensity score (i.e. probability of having specific attitude toward HIV/AIDS risk) for each subnetwork based on individual-level covariates of sex, race, education, age and their pairwise interactions.
3. Use the inverse of propensity score as weights to create IPW estimator of potential outcome.

Causal Inference: IPW estimation

Table 1: Estimated risk differences (RDs) with 95% CIs of **locus of control** (external vs. internal) on likelihood of **receiving HIV test results** in SFHR (coverage of internal)

Effect	Coverage (α, α')	Unadjusted		Adjusted with interactions	
		RD	95% CI	RD	95% CI
Direct	(50%, 50%)	**-0.148**	(-0.230, -0.065)	**-0.160**	(-0.265, -0.055)
Direct	(70%, 70%)	**-0.142**	(-0.246, -0.038)	**-0.162**	(-0.268, -0.055)
Direct	(99%, 99%)	-0.101	(-0.258, 0.056)	-0.130	(-0.268, 0.008)
Indirect	(50%, 70%)	**-0.041**	(-0.071, -0.012)	**-0.031**	(-0.054, -0.008)
Indirect	(50%, 99%)	-0.070	(-0.156, 0.019)	-0.062	(-0.123, 0.000)
Indirect	(70%, 99%)	-0.029	(-0.098, 0.040)	-0.030	(-0.072, 0.011)
Total	(50%, 70%)	**-0.183**	(-0.271, -0.096)	**-0.193**	(-0.286, -0.100)
Total	(50%, 99%)	**-0.172**	(-0.278, -0.066)	**-0.192**	(-0.291, -0.093)
Total	(70%, 99%)	**-0.130**	(-0.254, -0.006)	**-0.161**	(-0.272, -0.049)
Overall	(50%, 70%)	**-0.067**	(-0.096, -0.038)	**-0.065**	(-0.089, -0.041)
Overall	(50%, 99%)	**-0.097**	(-0.183, -0.010)	**-0.111**	(-0.190, -0.032)
Overall	(70%, 99%)	-0.030	(-0.095, 0.035)	-0.046	(-0.105, 0.013)

Why Does Spillover Matter for Society?

- Interventions can be tailored to maximize impact by capitalizing on network structure
- Informs community-level treatments and policy
  - How many and which community members to target for Treatment as Prevention and Pre-exposure Prophylaxis?
  - How to best distribute naloxone for opioid overdose prevention in the community?
  - What are the social influences on injection drug use among youth with NMU-PO ?
  - Who to train as a peer education within a community of people who use drugs?
- Design interventions to leverage beneficial disseminated effects beyond those directly treated
- Implications for cost-effectiveness evaluations

Discussion Questions

- Which settings would you expect there to be interference? Which settings would you find the assumption of no interference plausible?
- Can you think of other settings where this type of analysis could be useful?
- Why is it important to understand these types of effects for health outcomes?

Example

Table 1.1 of Potential Outcomes and Table 1.2 of Observed Data

Patient	$Y^{a=0}$	$Y^{a=1}$
Rheia	0	1
Kronos	1	0
Demeter	0	0
Hades	0	0
Hestia	0	0
Poseidon	1	0
Hera	0	0
Zeus	0	1
Artemis	1	1
Apollo	1	0
Leto	0	1
Ares	1	1
Athena	1	1
Hephaestus	0	1
Aphrodite	0	1
Cyclope	0	1
Persephone	1	1
Hermes	1	0
Hebe	1	0
Dionysus	1	0

Patient	A	Y
Rheia	0	0
Kronos	0	1
Demeter	0	0
Hades	0	0
Hestia	1	0
Poseidon	1	0
Hera	1	0
Zeus	1	1
Artemis	0	1
Apollo	0	1
Leto	0	0
Ares	1	1
Athena	1	1
Hephaestus	1	1
Aphrodite	1	1
Cyclope	1	1
Persephone	1	1
Hermes	1	0
Hebe	1	0
Dionysus	1	0

Average Causal Effects

- There is an *average causal effect* in the (super-)population if

$$\Pr[Y^{a=1} = 1] \neq \Pr[Y^{a=0} = 1]$$

or more generally (i.e., for a non-dichotomous Y)

$$E[Y^{a=1}] \neq E[Y^{a=0}]$$

- There is no average causal effect in the population if

$$\Pr[Y^{a=1} = 1] = \Pr[Y^{a=0} = 1]$$

This is implied by, but does not imply, the sharp null

- E.g., treatment may have an effect on some individuals, but no average effect

Measures of Association

- Associational risk difference (RD)

$$\Pr[Y = 1|A = 1] - \Pr[Y = 1|A = 0] = E(Y = 1|A = 1) - E(Y = 1|A = 0)$$

- Associational risk ratio (RR)

$$\frac{\Pr[Y = 1|A = 1]}{\Pr[Y = 1|A = 0]}$$

- Associational odds ratio (OR)

$$\frac{\Pr[Y = 1|A = 1] / \Pr[Y = 0|A = 1]}{\Pr[Y = 1|A = 0] / \Pr[Y = 0|A = 0]}$$

- If $A \perp\!\!\!\perp Y$ (independent) then assoc. RD equals 0, assoc. RR and OR equal 1

Conditional Randomization

- In a marginally randomized experiment, full exchangeability holds, which implies exchangeability

$$Y^a \perp\!\!\!\perp A \text{ for } a = 0, 1$$

- This is not necessarily true in a conditionally randomized experiment
- Conditionally randomized experiments can be viewed as two separate marginal experiments
- Thus, conditional exchangeability holds

$$Y^a \perp\!\!\!\perp A|L = l \text{ for } a, l = 0, 1$$

where is written as $Y^a \perp\!\!\!\perp A|L = l$ for $a = 0, 1$ for notational ease

Standardization

- Under conditional exchangeability, the standardized mean equals the counterfactual mean had all individuals in population received treatment a
- Thus causal effects can be identified under conditional exchangeability, e.g., the casual RR equals

$$\frac{\Pr[Y^{a=1}=1]}{\Pr[Y^{a=0}=1]} = \frac{\sum_l E[Y=1|A=1,L=l]\Pr[L=l]}{\sum_l E[Y=1|A=0,L=l]\Pr[L=l]}$$

- In the usual scenario where we are not ignoring sampling variability, this suggests a consistent estimator based on plugging-in observed proportions

$$\frac{\sum_l \hat{\Pr}[Y=1|A=1,L=l]\hat{\Pr}[L=l]}{\sum_l \hat{\Pr}[Y=1|A=0,L=l]\hat{\Pr}[L=l]}$$

- Similarly for the causal RD and OR

Inverse Probability Weighting

- Consider the following Horwitz-Thompson-type estimator of $E[Y^a]$

$$\frac{1}{n}\sum_{i=1}^{n}\frac{I[A_i = a]Y_i}{\Pr[A_i = a|L_i]}$$

- This estimator is unbiased when $\Pr[A = a|L = l]$ is known for all a, l (as in a conditional randomized experiment) and conditional exchangeability holds
- $\Pr[A = 1|L = l]$ is the *propensity score*
- Notes: (i) notation $\Pr[A = a|L]$ is different from notation of HR (Technical Point 2.2); (ii) $\Pr[A = a|L]$ is a random variable whereas $\Pr[A = a|L = l]$ is a constant

Conditional Exchangeability (CE)

- Unfortunately, it is not possible to verify $Y^a \perp\!\!\!\perp A|L$ because, in order to do so, one would need to test

$$\Pr[Y^a = 1|A = a, L = l] = \Pr[Y^a = 1|A \neq a, L = l]$$

  But this is not possible because Y^a is never observed for individuals with $A \neq a$ (i.e., right side is not identifiable)

- Thus causal inference in observational studies often relies on expert knowledge to select L in order to ensure CE plausible [In Section 7, we will discuss how to use graphs to choose the set of covariates L to ensure CE holds]
- CE will not hold if there exist unmeasured independent predictors U of outcome such that probability of receiving treatment A depends on U within strata of L. For this reason, $Y^a \perp\!\!\!\perp A|L$ often referred to as the *no unmeasured confounders* assumption.

Positivity

- Positivity is defined to be

$$\Pr[A = a|L = l] > 0 \text{ for all } l \text{ such that } \Pr[L = l] > 0$$

- In Table 3.1 positivity holds because there are individuals at both levels of treatment ($A = 0$ and $A = 1$) for each level of the covariate L (0 and 1)
- Positivity would not hold if an individual with critical condition at baseline $L = 1$ always gets treatment, i.e., $\Pr[A = 0|L = 1] = 0$
- Unlike conditional exchangeability, positivity can sometimes be empirically verified

Well-defined interventions

- Up until now we have assumed there are only two versions of treatment, $a = 1$ and $a = 0$, and hence two potential outcomes, $Y^{a=0}$ and $Y^{a=1}$, per individual (recall this was part of SUTVA)
- However it may be there are different versions of treatment $a = 1$
- E.g., "heart transplant" might entail different surgeons, different pre-op procedures, etc
- These different versions of treatment could result in different potential outcomes

Notation

- N groups; n_i individuals in groups $i = 1, \ldots, N$
- $\mathbf{A}_i = (A_{1i}, \ldots, A_{1n_i})$ treatments received for n_i individuals in group i
  $A_{ij} = 0$ or 1 implies $\mathbf{A}_i$ can take on 2^{n_i} possible values
  $\mathbf{A}_{i,-j}$ is the $n_i - 1$ subvector of $\mathbf{A}_i$ with the j^{th} entry deleted
  $\mathbf{a}_i$ and a_{ij} denote possible values of $\mathbf{A}_i$ and A_{ij}
- Let $A(n)$ be the set of vectors of all possible exposure allocations of length n. e.g., $A(2) = \{(0,0),(0,1),(1,0),(1,1)\}$, $\mathbf{a}_i \in R^{n_i}$
- $A(n,k)$ denotes when exactly k individuals receive treatment 1 (i.e., completely randomized design)
- Let α be the proportion treated in a group

Assignment Mechanism

- $S_i = 1$ if the i^{th} group is assigned to α_1 and 0 otherwise
  $\mathbf{S} = (S_1, \ldots, S_N)$
  $C = \sum_i S_i$
- Parameterization for treatment assignment strategy
  - Complete randomized group assignment strategy if k_i number treated in block i , i.e., $\pi(a_i, \alpha) = I(a_i \in A(n_i, k_i))/\binom{n_i}{k_i}$
  - Bernoulli Allocation: $\pi(\mathbf{a}_i, \alpha) = \prod_{j=1}^{n_i} \alpha^{a_{ij}} (1-\alpha)^{1-a_{ij}}$

Potential Outcomes

- $y_{ij}(\mathbf{a}_i)$ is the potential outcome of individual j in group i under $\mathbf{a}_i$
- Allows for interference between individuals within group i
- Can write $y_{ij}(\mathbf{a}_i)$ as $y_{ij}(\mathbf{a}_{i,-j}, a_{ij} = a)$
- Have 2^{n_i} potential outcomes per individual, instead of 2 potential outcomes per individual in the absence of interference

Average Potential Outcomes

- Individual average potential outcome

$$\bar{y}_{ij}(a, \alpha) = \sum_{a_{i,-j} \in A(n_i - 1)} y_{ij}(\alpha, a_{ij} = a) \Pr(\mathbf{A}_{i,-j} = a_{i,-j} | A_{ij} = a)$$

- Group average potential outcome

$$\bar{y}_i(a, \alpha) = \frac{1}{n_i} \sum_{i}^{n_i} \bar{y}_{ij}(a, \alpha)$$

- Population average potential outcome

$$\bar{y}(a, \alpha) = \frac{1}{N} \sum_{i}^{N} \bar{y}_i(a, \alpha)$$

Causal Estimands: Direct Effects

- Individual direct causal effect of treatment 0 compared to treatment 1 for the individual j in group i by

$$CE_{ij}^{D}(\alpha) = y_{ij}(a_{ij} = 1, \alpha) - y_{ij}(a_{ij} = 0, \alpha)$$

- Individual average direct causal effect

$$\overline{CE}_{ij}^{D}(\alpha) = \bar{y}_{ij}(1, \alpha) - \bar{y}_{ij}(0, \alpha)$$

- Group average direct causal effect

$$\overline{CE}_{i}^{D}(\alpha) = \bar{y}_{i}(1, \alpha) - \bar{y}_{i}(0, \alpha)$$

- Population average direct causal effect

$$\overline{CE}^{D}(\alpha) = \bar{y}(1, \alpha) - \bar{y}(0, \alpha)$$

Causal Estimands: Indirect Effects

- Individual indirect causal effect of treatment programs α_1 compared with α_1 on individual j in group i by

$$CE^{I}_{ij}(\alpha_1, \alpha_0) = y_{ij}(\alpha_1, a_{ij} = 0) - y_{ij}(\alpha_0, a_{ij} = 0)$$

- Individual average indirect causal effect

$$\overline{CE}^{I}_{ij}(\alpha_1, \alpha_0) = \bar{y}_{ij}(0, \alpha_1) - \bar{y}_{ij}(0, \alpha_0)$$

- Group average indirect causal effect

$$\overline{CE}^{I}_{i}(\alpha_1, \alpha_0) = \bar{y}_{i}(0, \alpha_1) - \bar{y}_{i}(0, \alpha_1)$$

- Population average indirect causal effect

$$\overline{CE}^{I}(\alpha_1, \alpha_0) = \bar{y}(0, \alpha_1) - \bar{y}(0, \alpha_0)$$

Causal Estimands: Total and overall Effects

- Population average total causal effect

$$\overline{CE}^{T}(\alpha_1, \alpha_0) = \bar{y}(1, \alpha_1) - \bar{y}(0, \alpha_0)$$

- Population average overall causal effect

$$\overline{CE}^{O}(\alpha_1, \alpha_0) = \bar{y}(\alpha_1) - \bar{y}(\alpha_0)$$

Estimands

$$\textit{Direct effect}: \quad \overline{DE}(\alpha) = \overline{Y}(a = 0; \alpha) - \overline{Y}(a = 1; \alpha) \tag{1}$$

$$\textit{Indirect effect}: \quad \overline{IE}(\alpha, \alpha') = \overline{Y}(a = 0; \alpha) - \overline{Y}(a = 0; \alpha') \tag{2}$$

$$\textit{Total effect}: \quad \overline{TE}(\alpha, \alpha') = \overline{Y}(a = 0; \alpha) - \overline{Y}(a = 1; \alpha') \tag{3}$$

$$\textit{Overall effect}: \quad \overline{OE}(\alpha, \alpha') = \overline{Y}(\alpha) - \overline{Y}(\alpha') \tag{4}$$

coverage: $\alpha < \alpha'$

Estimators

- $\widehat{CE}^{D}(\alpha_1) = \hat{Y}(1, \alpha_1) - \hat{Y}(1, \alpha_1)$
- $\widehat{CE}^{I}(\alpha_1, \alpha_0) = \hat{Y}(0, \alpha_1) - \hat{Y}(0, \alpha_0)$
- $\widehat{CE}^{T}(\alpha_1, \alpha_0) = \hat{Y}(1, \alpha_1) - \hat{Y}(0, \alpha_0)$

Overall Estimators

- $\widehat{Y}_i(\alpha) = \frac{\sum_j Y_{ij}}{n_i}$
- $\widehat{Y}(\alpha) = \frac{\sum_i \hat{Y}_i(\alpha) I[S_i=1]}{\sum_i I[S_i=1]}$
- Under assumption 1, $E\{\widehat{Y}(\alpha)\} = \bar{y}(\alpha)$
- Unbiased estimator: $\widehat{CE}^{O}(\alpha_1, \alpha_0) = \widehat{Y}(\alpha_1) - \widehat{Y}(\alpha_0)$

Variance

- Unbiased estimators of the variance of the estimators does not exist without further assumptions
- Stratified Interference (SI): Only matters how many were treated in group or cluster, and does not matter who was treated
- For a given $a_{ij} = a$, individual j in group i has
  1 potential outcome assuming no interference
  n_i potential outcomes assuming stratified interference
  $2^{n_i - 1}$ potential outcomes under no assumptions
- Under SI, simple random sampling and two stage cluster sampling yield unbiased estimators of variance of $\hat{Y}_i(0, \alpha)$ and $\hat{Y}(0, \alpha_1)$
- Variance estimators are unbiased when effect is additive, positively biased otherwise

Notation

- Consider finite population of n individuals with ids $i = 1, \ldots, n$
- A_i self-selected binary exposure ($A_i = 1$ if exposed and 0, otherwise)
- Z_i denote the vector of covariates for participant i.
- Observed network represented by a binary adjacency matrix $g = [g_{ij}]_{i,j=1}^{n} \in \{0,1\}^{n \times n}$.
- Nearest neighborhood of participant i is $\mathcal{N}_i = \{j : g_{ij} = 1\}$.
- Vector of exposures for neighbors of i is $A_{\mathcal{N}_i} = \{A_{ij}, g_{ij} = 1\}$.
- The observed outcome of i, Y_i, is random and depends not only on A_i, but also on $A_{\mathcal{N}_i}$, $Y_i = y_i(A_i, A_{\mathcal{N}_i})$.

Assumptions

- We assume Bernoulli counterfactual treatment allocation strategy with coverage α ($\sim$ participants in $\mathcal{N}_i$ are exposed with prob. α).
- Let $\pi(a_{\mathcal{N}_i}; \alpha) = \alpha^{\sum a_{\mathcal{N}_i}} (1-\alpha)^{|\mathcal{N}_i| - \sum a_{\mathcal{N}_i}}$ denote the probability of the nearest neighborhood for an individual i receiving treatment $A_{\mathcal{N}_i}$ under allocation strategy α.
- Define $\bar{y}_i(a, \alpha) = \sum_{a_{\mathcal{N}_i}} y_i(a_i = a, a_{\mathcal{N}_i}) \pi(a_{\mathcal{N}_i}; \alpha)$ to be the average potential outcome for individual i under allocation strategy α.
- With additional assumption that dissemination (i.e., interference) within $\mathcal{N}_i$ is invariant to which particular subset of neighbors is treated and depends only on number exposed in $\mathcal{N}_i \rightarrow y_i(a_i, \alpha) = y_i(a_i, f(a_{\mathcal{N}_i}))$, where $f(a_{\mathcal{N}_i}) = \sum_{\mathcal{N}_i} a_i / d_i$.

Models and Estimation

Inverse probability weighting (IPW) can be used to replicate two-stage randomized design with observational study (Tchetgen Tchetgen and VanderWeele, 2012).

$$\hat{Y}_i^{ipw}(a, \alpha) = \sum_{j=1}^{n} \frac{y_i(A_i, A_{\mathcal{N}_i})\pi_i(A_{\mathcal{N}_i}; \alpha)I(A_i = a)}{nf(A_i, \mathcal{N}_i | z_i, z_{\mathcal{N}_i})}$$