Efficient Treatment Effect Estimation in Observational Studies under Heterogeneous Partial Interference

Ruoxuan Xiong

Emory University

Joint work with Zhaonan Qu (Stanford), Jizhou Liu (UChicago), Guido Imbens (Stanford GSB)

Introduction

Interference

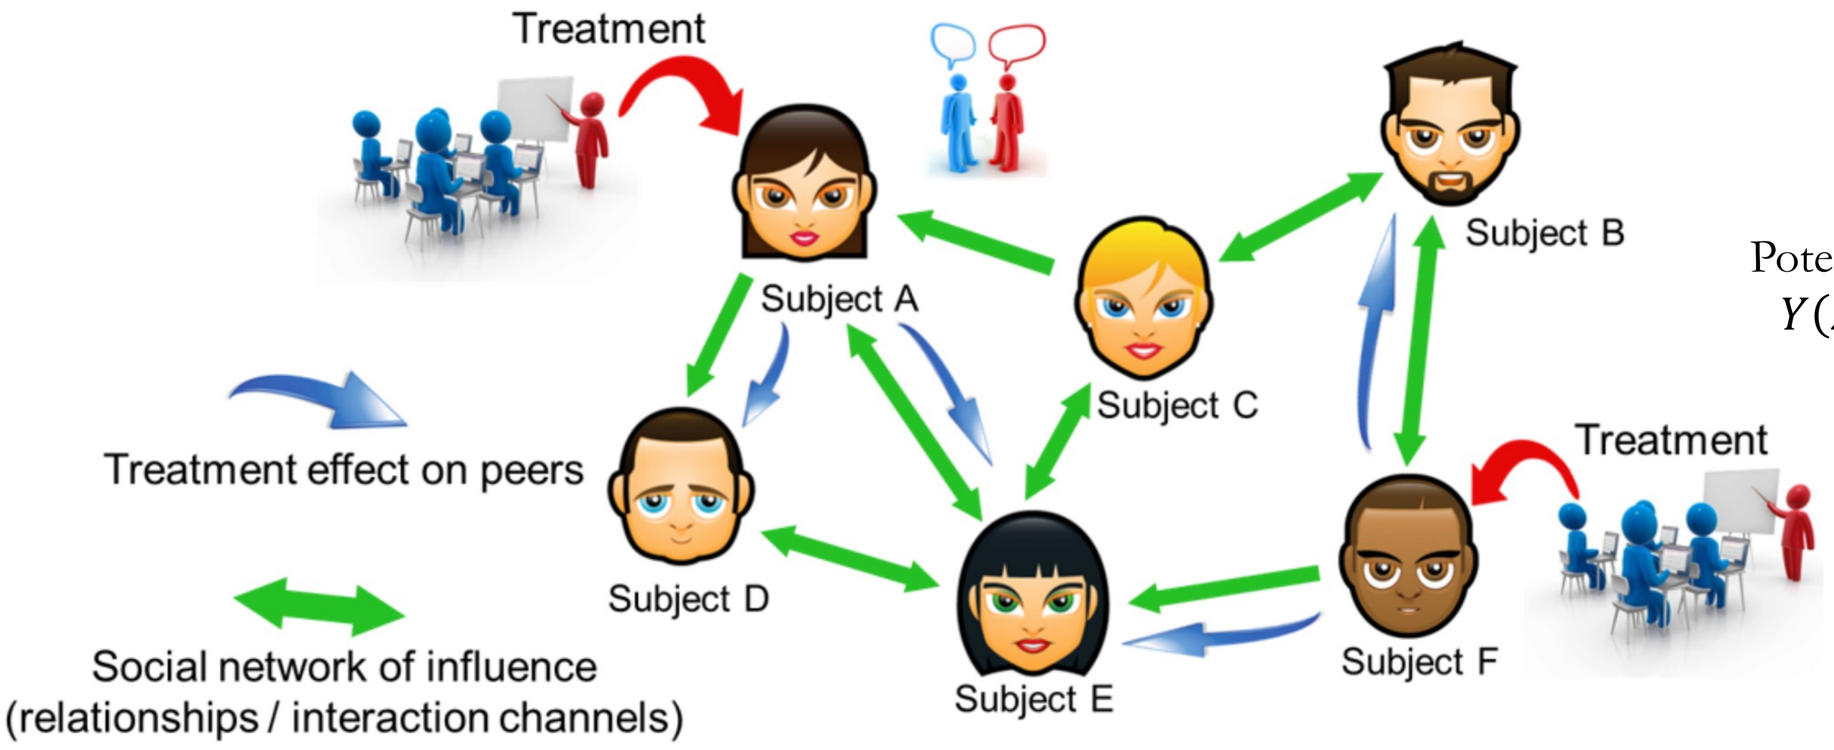

Potential outcomes:
$Y(z) \rightarrow Y(z, \boldsymbol{z}_{neighbors})$

Figure 1-1 from Kao (2017)

Spillover effects from influencers may differ from spillover effects from regular subjects

- This paper: Heterogeneous interference

Motivating example: Add Health data

- **Add Health data**: National longitudinal study of adolescent health

- **Interference** in the Add Health data
  - The treatment on one unit has an effect on the response of another unit (Cox, 1958)
    - E.g., risky behavior (e.g., smoking/drug use) may have spillover effects on friends' academic performance
    - See Bramoullé et al. (2009), Goldsmith-Pinkham and Imbens (2013), Forastiere, Airoldi, and Mealli (2020)

This paper: Heterogeneous interference

- **Two types of heterogeneous interference**
  1. *Spillover effect* varies with the types of units and treated neighbors

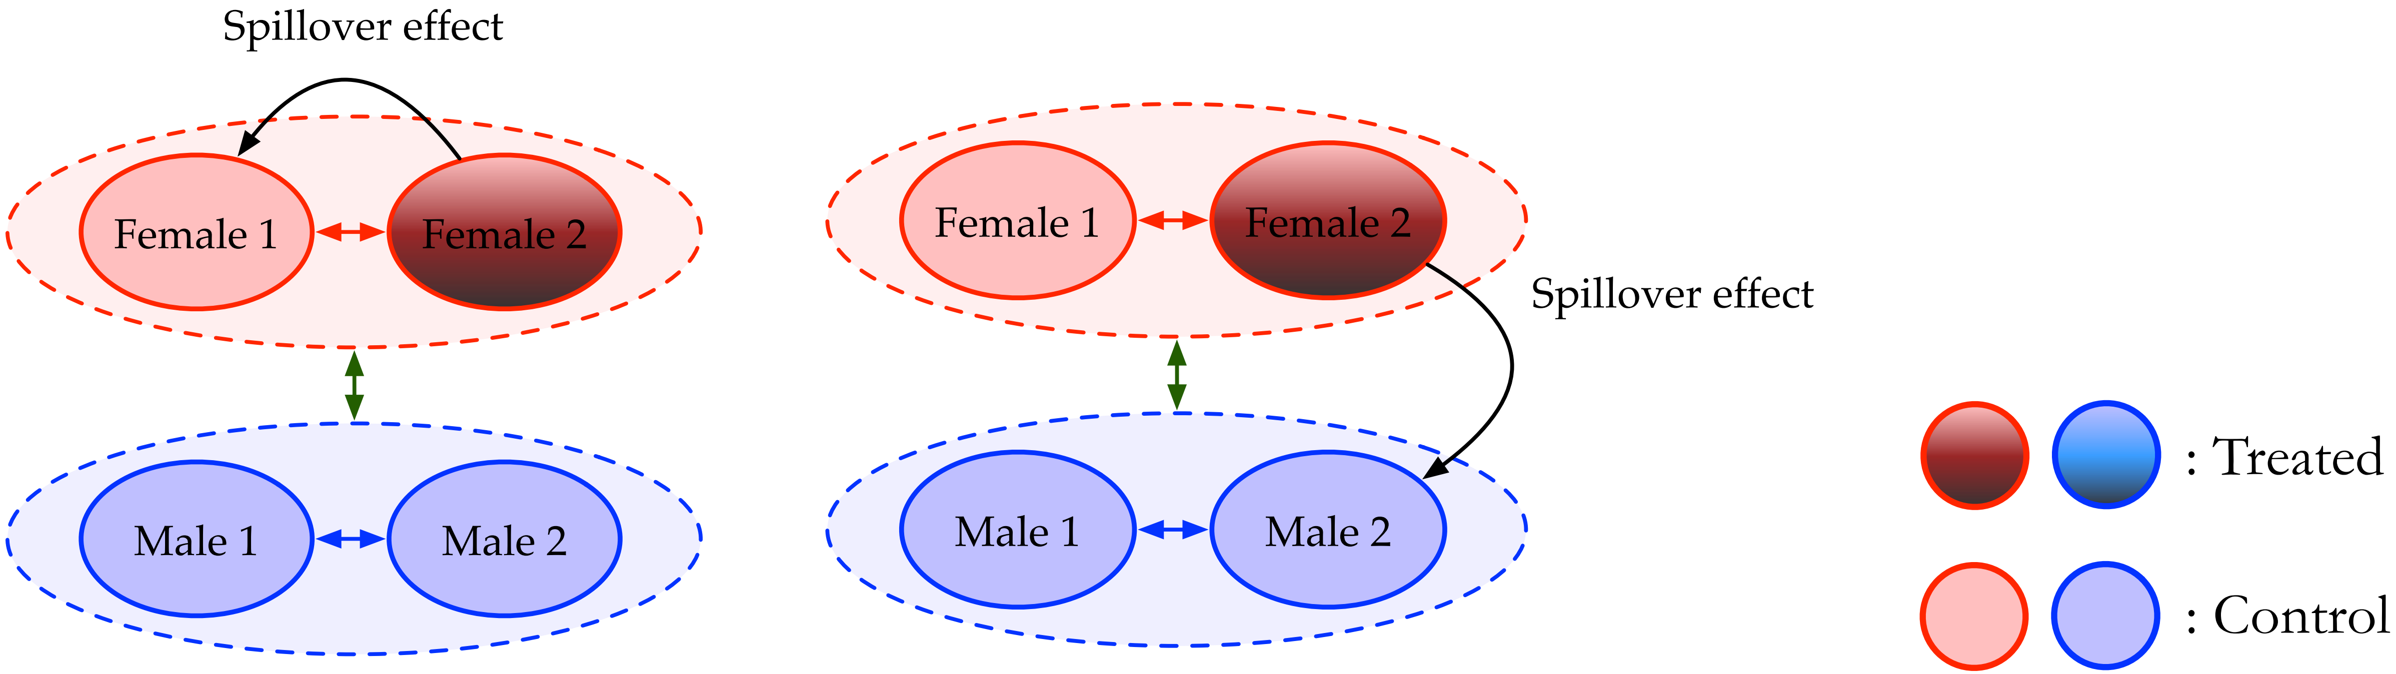

This paper: Heterogeneous interference

- **Two types of heterogeneous interference**
  1. *Spillover effect* varies with the types of units and treated neighbors
  2. *Direct treatment effect* varies with the types of units and treated neighbors

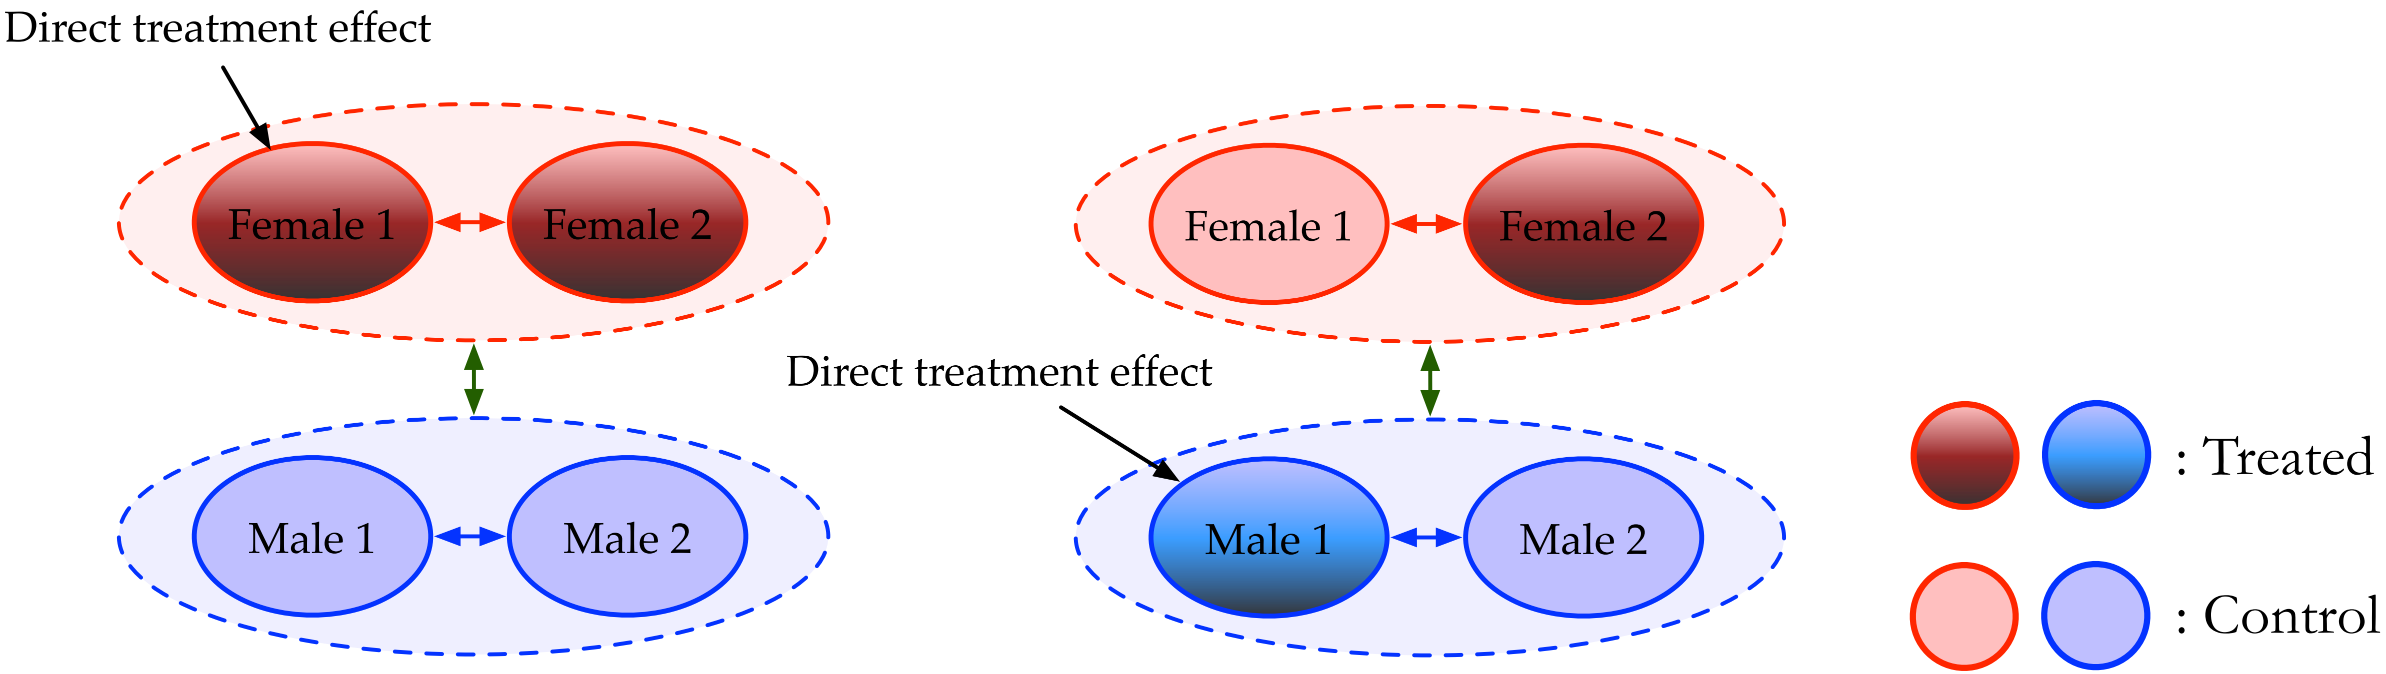

Some evidence of heterogeneous interference

- Study the effect of smoking on academic performance
- *Direct and spillover effects* seem to vary with gender on the Add Health data

Estimated *spillover effects* of two smoking friends

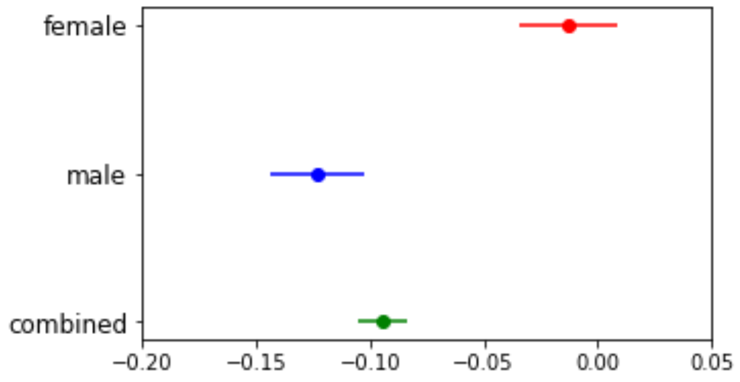

Estimated *direct effects* given a treated friend of opposite gender

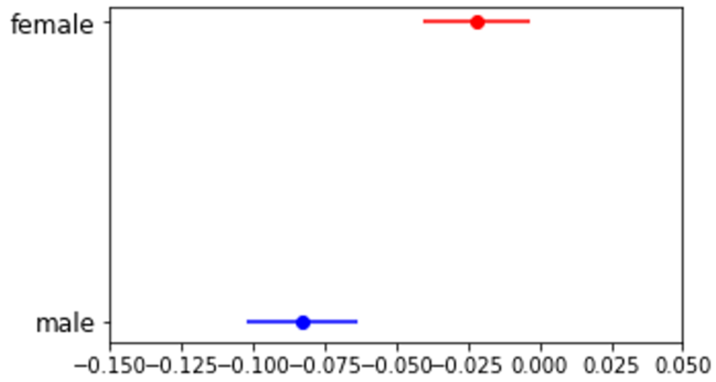

Importance of heterogeneous interference

1. The estimation of treatment effects can be biased if heterogeneity in interference is neglected

2. Treatment targeting/policy learning
   - E.g., Banerjee et al. (2013)

Estimation of average treatment effect

- In the presence of heterogeneous interference, the estimation of ATE can be biased under the assumption of no interference or homogeneous interference
  - *Intuition for bias*: Subpopulations are not properly weighted in the presence of selection bias

Estimation error of ATE

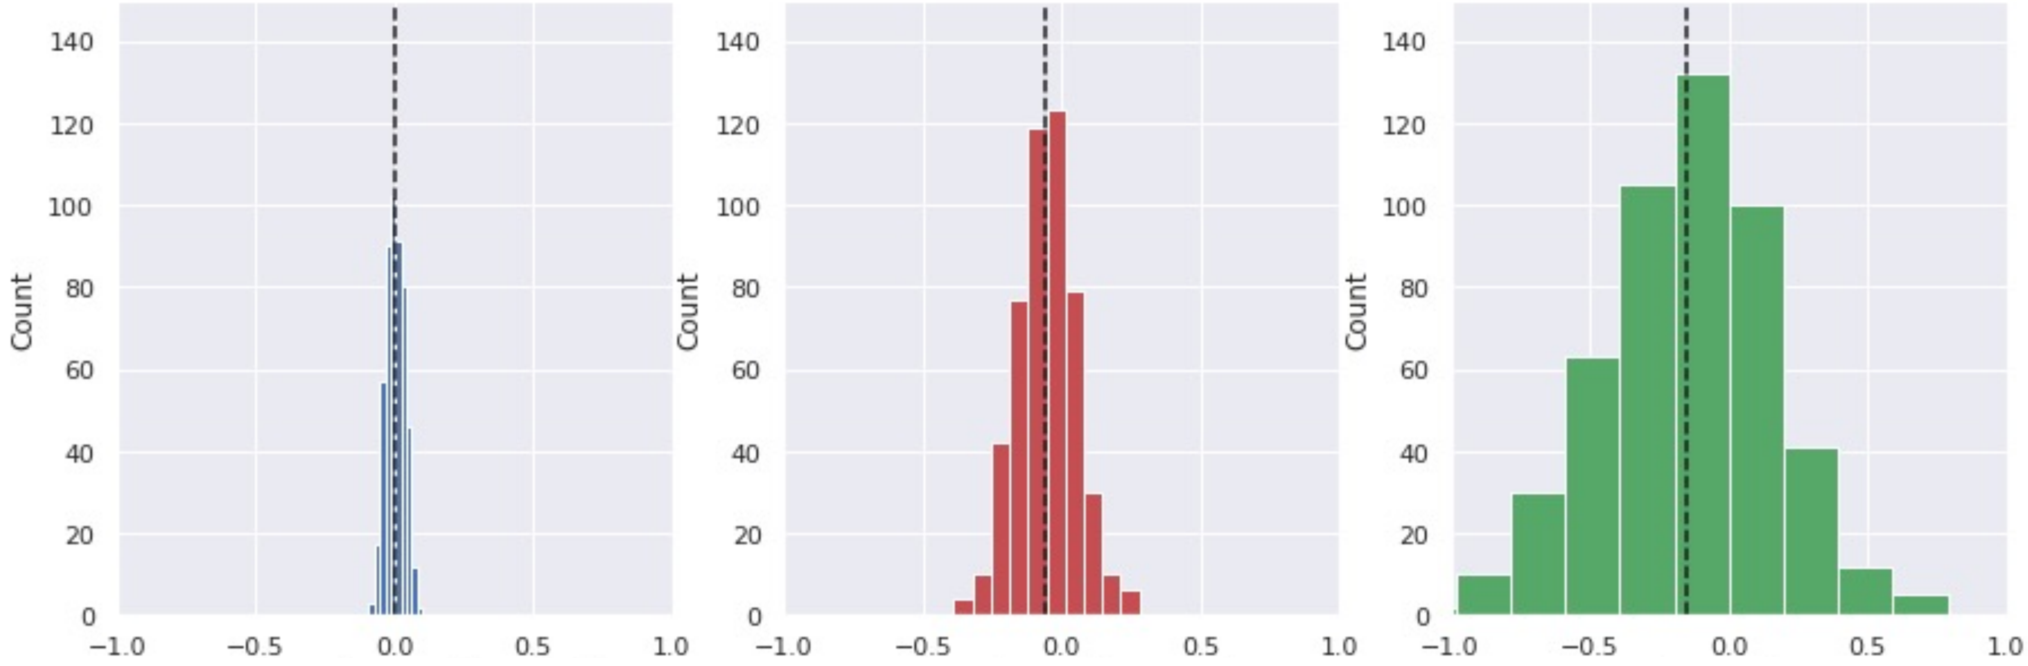

Heterogeneous interference — Homogeneous interference — No interference

Importance of heterogeneous interference

1. The estimation of treatment effects can be biased if heterogeneity in interference is neglected

2. Treatment targeting/policy learning
   - E.g., Banerjee, Duflo, Glennerster, and Kinnan (2013)

Contributions

- **This paper**: Treatment effect estimation in observational studies under partial interference
- **Methodology and theory**
  - A "conditional exchangeability" framework for heterogeneous interference
  - Generalized AIPW estimators for *direct and spillover effects*
  - Show that the generalized AIPW estimators are doubly-robust, asymptotically normal, semiparametric efficient
  - Show a bias-variance trade-off between robustness to interference and estimation efficiency
  - Debiased, matching-based feasible variance estimators
- **Empirical study**
  - **Add Health data**: Study direct and spillover effects of smoking on academic performance

Related work

- **Treatment effect estimation with interference**
  - **Observational data**
    - *Our framework* is closest to Forastiere, Airoldi, and Mealli (2020) (full exchangeability)
    - α *treatment allocation strategy*
      - *IPW*: Liu et al. (2016), Barkley et al. (2020)
      - *Doubly robust estimators*: Liu et al. (2019), Park and Kang (2020)
  - **Experimental data**
    - Interference as nuisance
      - E.g., Rosenbaum (2007), Hudgens and Halloran (2008), Tchetgen and VanderWeele (2012)
    - Interference-based causal effect estimation
      - E.g., Aronow and Samii (2017), Leung (2016, 2021)

Model setup and assumptions

Setup

- N units in an undirected network
- M disjoint clusters in this network
- For simplicity of presentation, clusters are of equal size n
  - Our paper: Extension to varying cluster size

- Cluster c:
  - $\boldsymbol{Y}_c \in \mathbb{R}^n$: Outcomes
  - $\boldsymbol{X}_c \in \mathbb{R}^{n \times k}$: Covariates
  - $\boldsymbol{Z}_c \in \{0,1\}^n$: Treatments

- Unit i in cluster c:
  - $Z_{c,i} \in \{0,1\}$: Treatment of unit i in cluster c
  - $\boldsymbol{Z}_{c,(i)} \in \{0,1\}^{n-1}$: Treatments of other units in cluster c
  - $\left(Y_{c,i}, \boldsymbol{Y}_{c,(i)}\right)$ and $\left(X_{c,i}, \boldsymbol{X}_{c,(i)}\right)$ are defined analogously

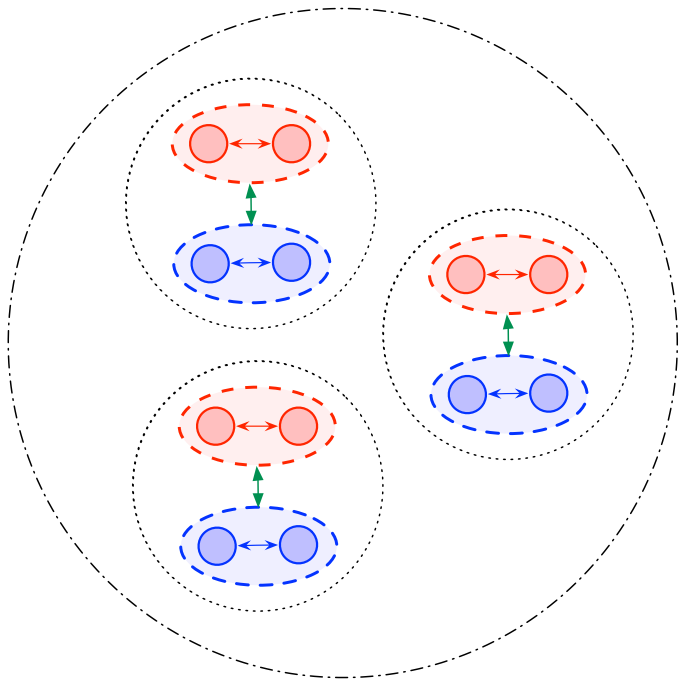

Partial interference

Assumption (Partial interference). Potential outcomes only depend on treatment assignments of units in the same cluster.

- Potential outcomes: $Y_{c,i}(Z_{c,i}, \boldsymbol{Z}_{c,(i)})$
  - E.g., a student's academic performance $Y_{c,i}$ depends on
    1. his/her smoking status $Z_{c,i}$
    2. his/her friends' smoking status $\boldsymbol{Z}_{c,(i)}$

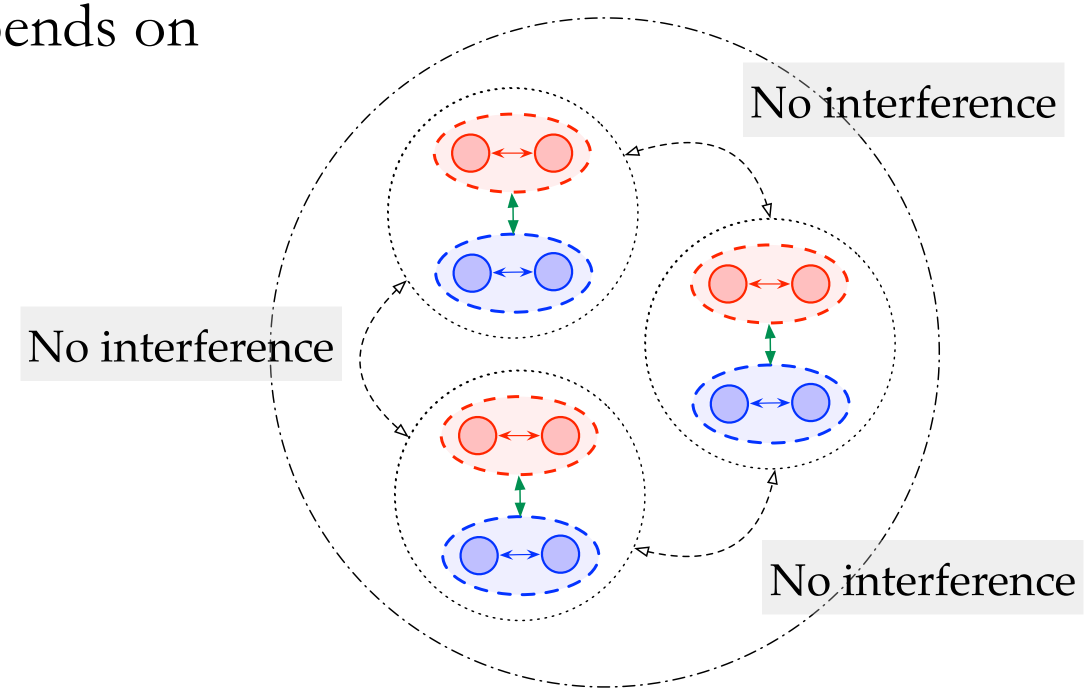

- See Sobel (2006)

Generalized unconfoundedness and overlap

Assumption (Generalized unconfoundedness). $Y_{c,i}\left(z_{c,i}, \boldsymbol{z}_{c,(i)}\right) \perp \left(Z_{c,i}, \boldsymbol{Z}_{c,(i)}\right) \mid \boldsymbol{X}_c$

Assumption (Overlap). $0 < \Pr\left(Z_{c,i}, \boldsymbol{Z}_{c,(i)} \mid \boldsymbol{X}_c\right) < 1$

- Generalization of Rosenbaum and Rubin (1983) to condition on $\boldsymbol{X}_c$
  - E.g., conditional on a student and his/her friends' characteristics, "as good as random assignment"
- Also imposed in Liu et al (2016), Forastiere, Airoldi, and Mealli (2020), and Park and Kang (2020)

Partial exchangeablity

Assumption (Partial exchangeability). Units in a cluster can be partitioned into m disjoint sets $I_1, \cdots, I_m$. Potential outcomes are exchangeable with respect to arbitrary permutations of units in the same subset.

- $m = 1$: Full exchangeability (Forastiere, Airoldi, and Mealli, 2020)
- For simplicity of presentation, $m = 2$
  - E.g., I_1: female friends; I_2: male friends
- Potential outcomes: $Y_{c,i}(Z_{c,i}, G_{c,(i),1}, G_{c,(i),2})$
  - $G_{c,(i),1}$: Number of treated units in subset I_1
  - E.g., $Y_{c,i}$ depends on
    1. his/her smoking status $Z_{c,i}$
    2. number of female smoking friends $G_{c,(i),1}$
    3. number of male smoking friends $G_{c,(i),2}$

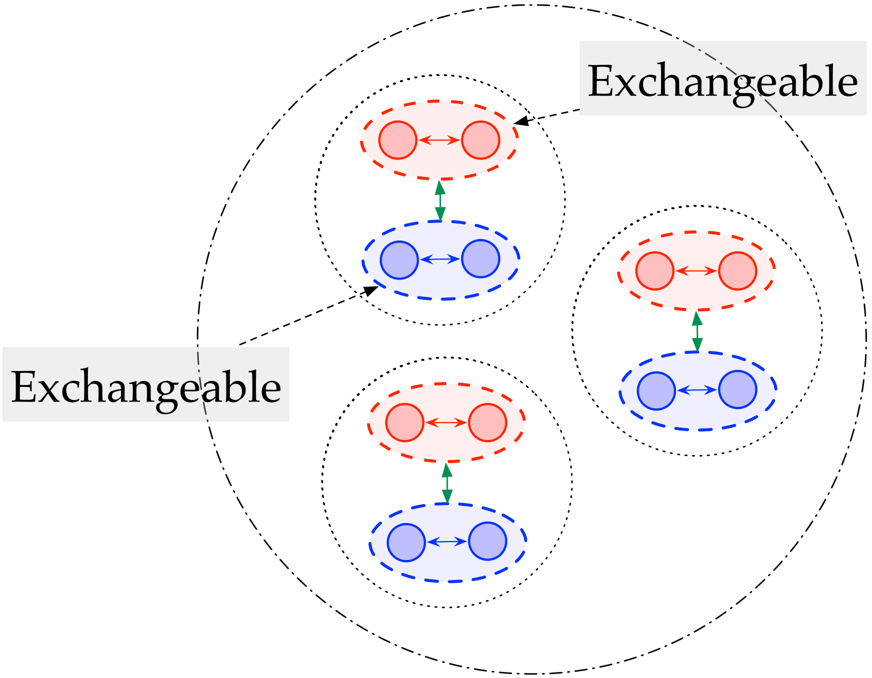

Estimation

Estimand: Direct treatment effect

$$\beta_j(g_1, g_2) = \frac{1}{|I_j|} \sum_{i \in I_j} E[Y_{c,i}(1, g_1, g_2) - Y_{c,i}(0, g_1, g_2)]$$

- Direct effect may vary with the subset of units I_j
  - E.g., direct effect on a female student $\beta_1(g_1, g_2)$ may differ from direct effect on a male student $\beta_2(g_1, g_2)$

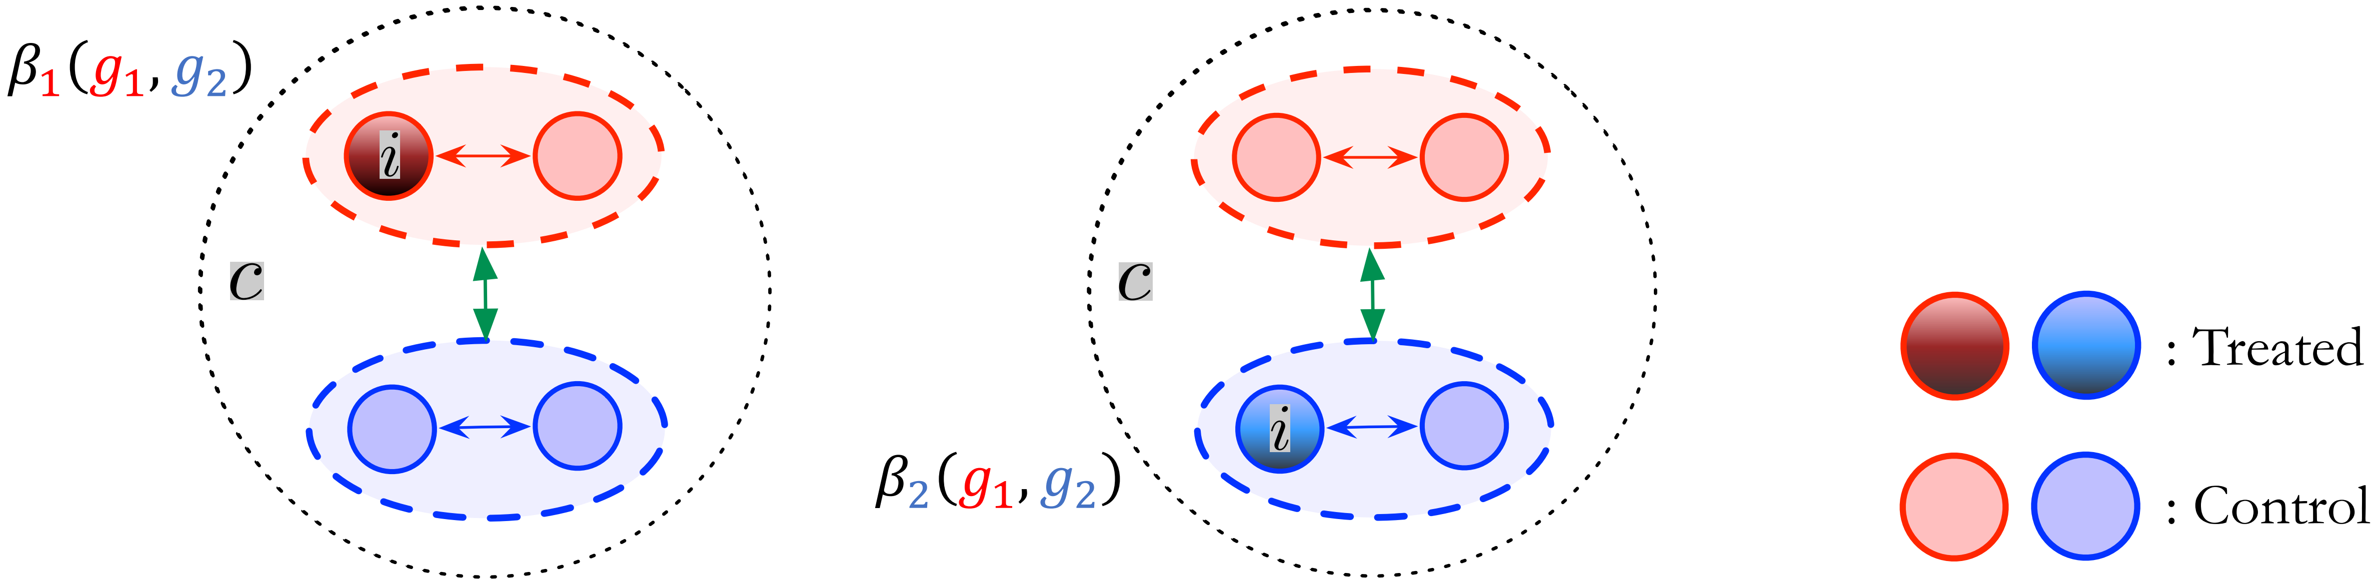

Estimand: Direct treatment effect

$$\beta_j(g_1, g_2) = \frac{1}{|I_j|} \sum_{i \in I_j} E[Y_{c,i}(1, g_1, g_2) - Y_{c,i}(0, g_1, g_2)]$$

- Direct effect may vary with number of and type of treated neighbors
  - E.g., direct effect of smoking may with whether the treated friend has the same gender

$\beta_1(1, 0)$ $\beta_1(0, 1)$

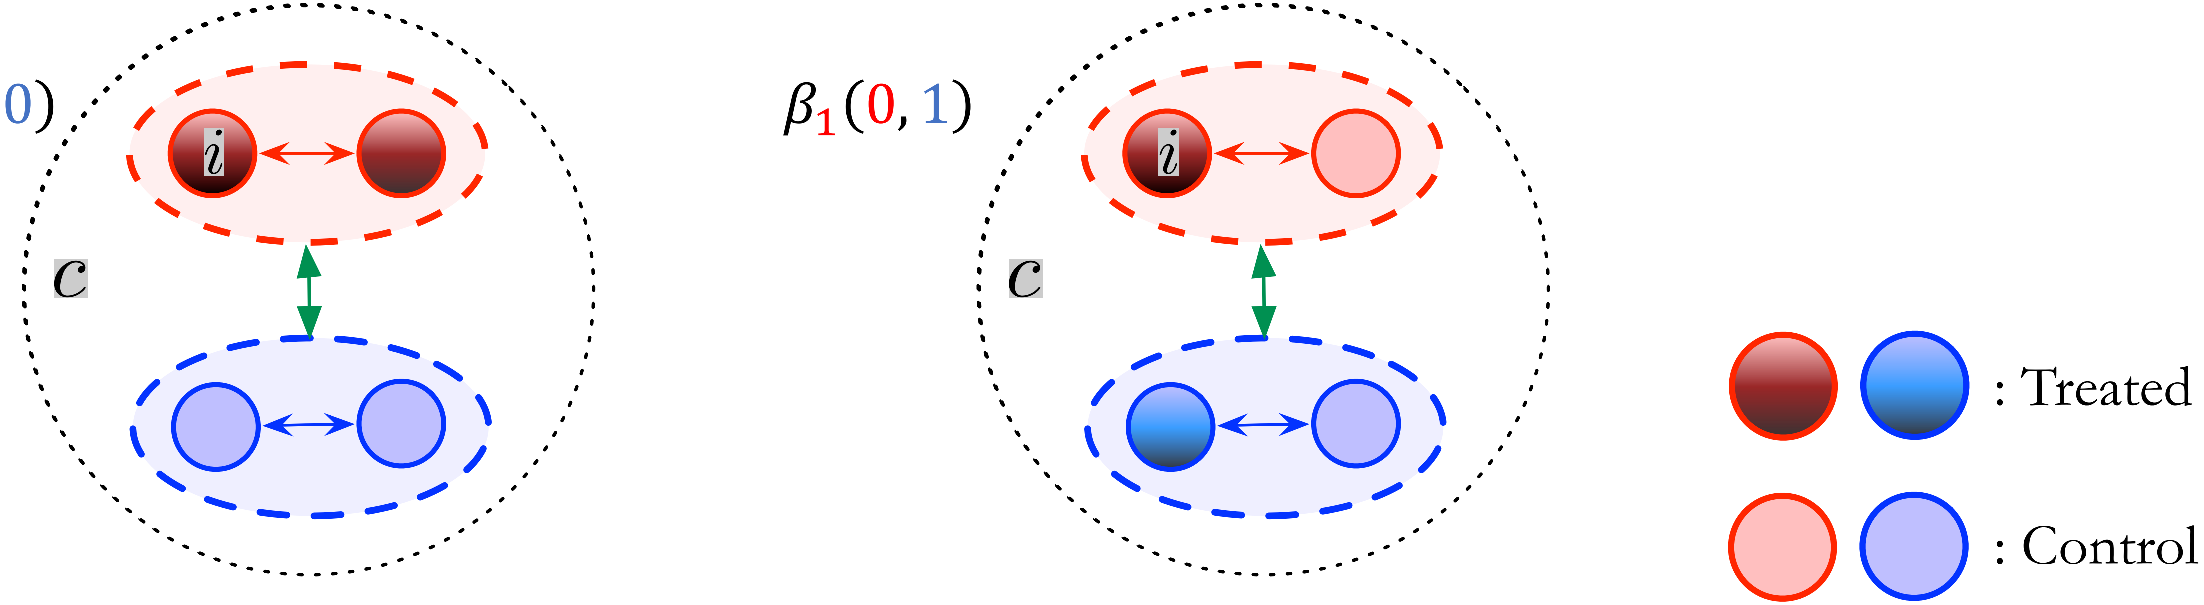

Estimand: Spillover effect

$$\tau_j(g_1, g_2) = \frac{1}{|I_j|} \sum_{i \in I_j} E[Y_{c,i}(0, g_1, g_2) - Y_{c,i}(0, 0, 0)]$$

- Spillover effect may vary with the subset of units I_j
- Spillover effect may vary with number of and type of treated neighbors

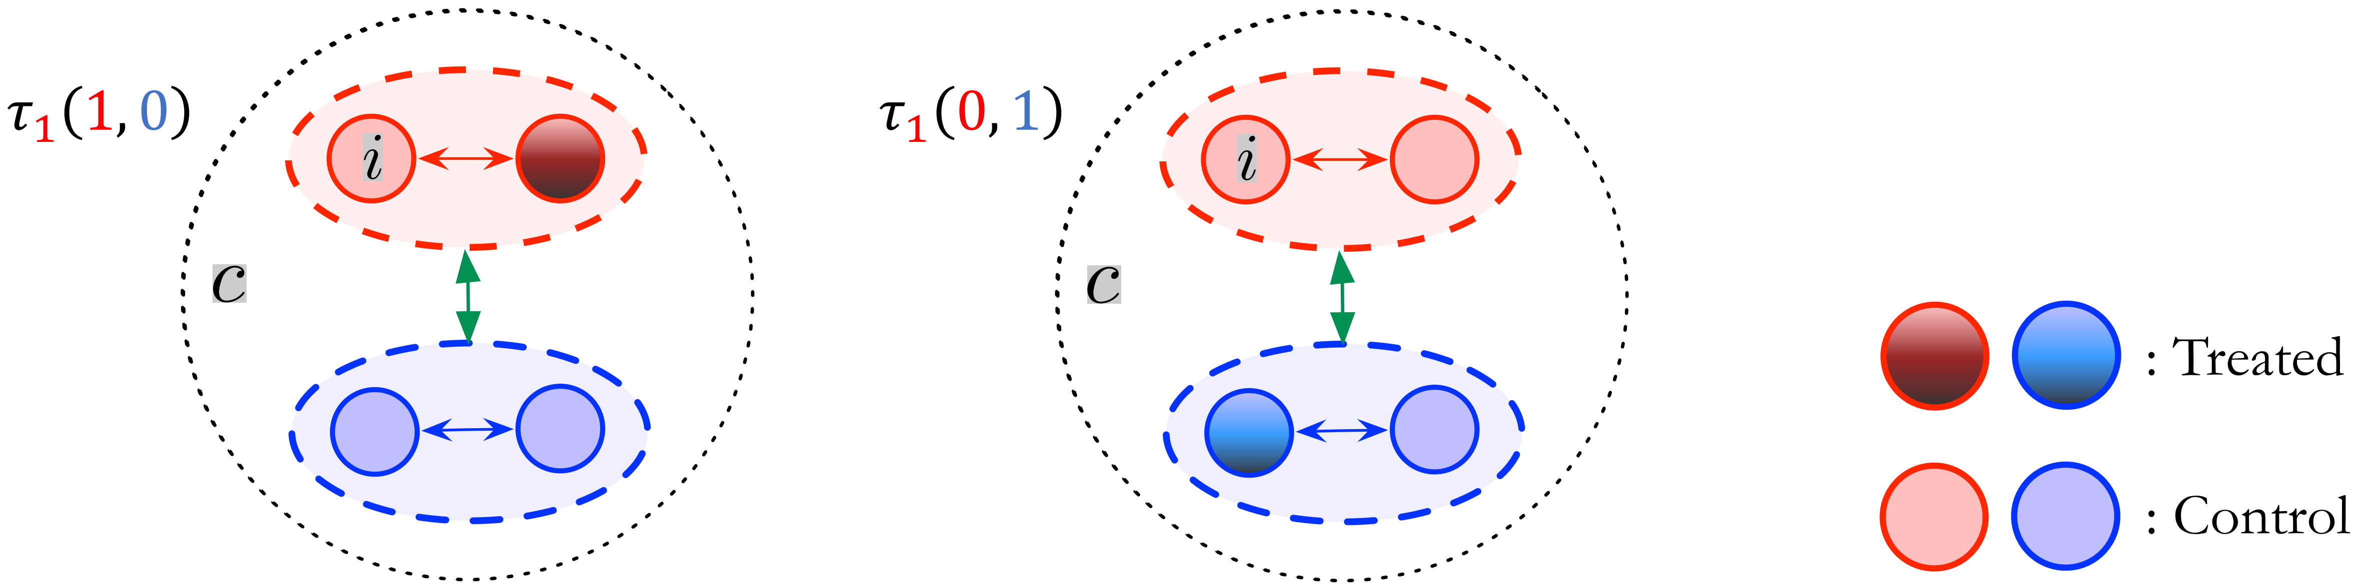

Estimator: Generalized AIPW

- Recall the conventional augmented inverse propensity weighted (AIPW) estimator for average treatment effect

$$\hat{\beta}_{ATE}^{aipw} = \frac{1}{N}\sum_{i}[\hat{\psi}_i(1) - \hat{\psi}_i(0)]$$

where the AIPW score is

$$\underbrace{\hat{\psi}_i(z) = \hat{\mu}_{(z)}(X_i)}_{\text{Estimated outcome model}} + \underbrace{\frac{1(Z_i = z)}{\hat{p}_{(z)}(X_i)}}_{\text{Inverse propensity weighting}} \cdot \underbrace{(Y_i - \hat{\mu}_{(z)}(X_i))}_{\text{Residual}}$$

- **Double robustness** property
  - If either the estimated outcome or propensity model is consistent, AIPW is consistent
- Achieve the **semiparametric efficiency bound**

Estimator: Generalized AIPW

- Generalized AIPW estimator for *direct treatment effects* of subset 1

$$\hat{\beta}_1^{aipw}(g_1, g_2) = \frac{1}{M \cdot |I_1|} \sum_{c=1}^{M} \sum_{i \in I_1} [\hat{\psi}_{c,i}(1, g_1, g_2) - \hat{\psi}_{c,i}(0, g_1, g_2)]$$

where the AIPW score is

$$\hat{\psi}_{c,i}(z, g_1, g_2) = \underbrace{\hat{\mu}_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)}_{\text{Estimated outcome model}} + \underbrace{\frac{1(Z_{c,i} = z, G_{c,(i),1} = g_1, G_{c,(i),2} = g_2)}{\hat{p}_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)}}_{\text{Inverse propensity weighting}} \cdot \underbrace{(Y_{c,i} - \hat{\mu}_{i,(z,g_1,g_2)}(\boldsymbol{X}_c))}_{\text{Residual}}$$

Estimator: Generalized AIPW

- Generalized AIPW for *spillover effects* is defined analogously
- How to estimate $\mu_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$ and $p_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$?
  - Our analysis is based on sieve estimator
  - Alternative: Machine learning estimators and cross-fitting

Main results

Asymptotic properties of generalized AIPW

Theorem 1 (Consistency and asymptotic normality, direct effects). As $M \to \infty$, if either $\hat{\mu}_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$ or $\hat{p}_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$ is uniformly consistent in $\boldsymbol{X}_c$, then $\hat{\beta}_1^{aipw}(g_1, g_2)$ is consistent and

$$\sqrt{M}\left(\hat{\beta}_1^{aipw}(g_1, g_2) - \beta_1(g_1, g_2)\right) \xrightarrow{\text{d}} N(0, V_{1,g_1,g_2})$$

where V_{1,g_1,g_2} is the semiparametric efficiency bound for $\beta_1(g_1, g_2)$.

- **Theorem 1** holds for AIPW estimators of spillover effects and for more than two subsets
  - E.g., $\hat{\beta}_2^{aipw}(g_1, g_2)$ and $\hat{\tau}_2^{aipw}(g_1, g_2)$
- Double robustness property
- Achieve the semiparametric efficiency bound
- Debiased, matching-biased feasible estimator for V_{1,g_1,g_2}

Robustness to interference and heterogeneity vs efficiency: A tradeoff

- A **bias-variance tradeoff** in the estimation of treatment effects
  - The estimator can be biased if the interference structure is (under) misspecified
  - The variance of the estimator is increased if the interference structure is more complex than necessary

Efficiency loss ↓ Reduced bias

No interference
Partial interference with **full** exchangeability (stratified interference)
Partial interference with **conditional** exchangeability
Network interference with **conditional** exchangeability
General interference

Three approaches to estimate ATE

- β_1: ATE for female students

- Three estimation approaches for β_1
  1. $\hat{\beta}_1^{n.i.}$: Use conventional approaches assuming **no interference**
  2. $\hat{\beta}_1^{p.i.}$: Based on $\sum \Pr(g_1 + g_2) \cdot \beta_1(g_1 + g_2)$, assuming **full** exchangeability (assume homogeneous interference)
  3. $\hat{\beta}_1^{p.i.c}$: Based on $\sum \Pr(g_1, g_2) \cdot \beta_1(g_1, g_2)$, assuming **conditional** exchangeability (allow for heterogeneous interference)

Bias-variance tradeoff between robustness and efficiency

Theorem 2 (Bias-variance tradeoff between robustness and efficiency). Suppose the assumptions in Theorem 1 holds.

1. $\hat{\beta}_1^{n.i.}$ is consistent only if there is no interference between units
2. $\hat{\beta}_1^{p.i.}$ is consistent only if units are fully exchangeable
3. If $\hat{\beta}_1^{n.i.}$, $\hat{\beta}_1^{p.i.}$ and $\hat{\beta}_1^{p.i.c}$ are consistent, then $\text{Avar}(\hat{\beta}_1^{n.i.}) \leq \text{Avar}(\hat{\beta}_1^{p.i.}) \leq \text{Avar}(\hat{\beta}_1^{p.i.c})$

- **Theorem 2** holds for other estimands and for more than two subsets

➢ Whenever possible, use the most parsimonious conditional exchangeability structure

✓ This paper provides hypothesis tests for heterogeneity of interference

- **Procedure**: Start with general interference structure, run hypothesis tests for heterogeneity, and potentially simplify interference structure

Variance of three estimation methods

Estimation of ATE based on different interference specifications (MSE in parentheses)

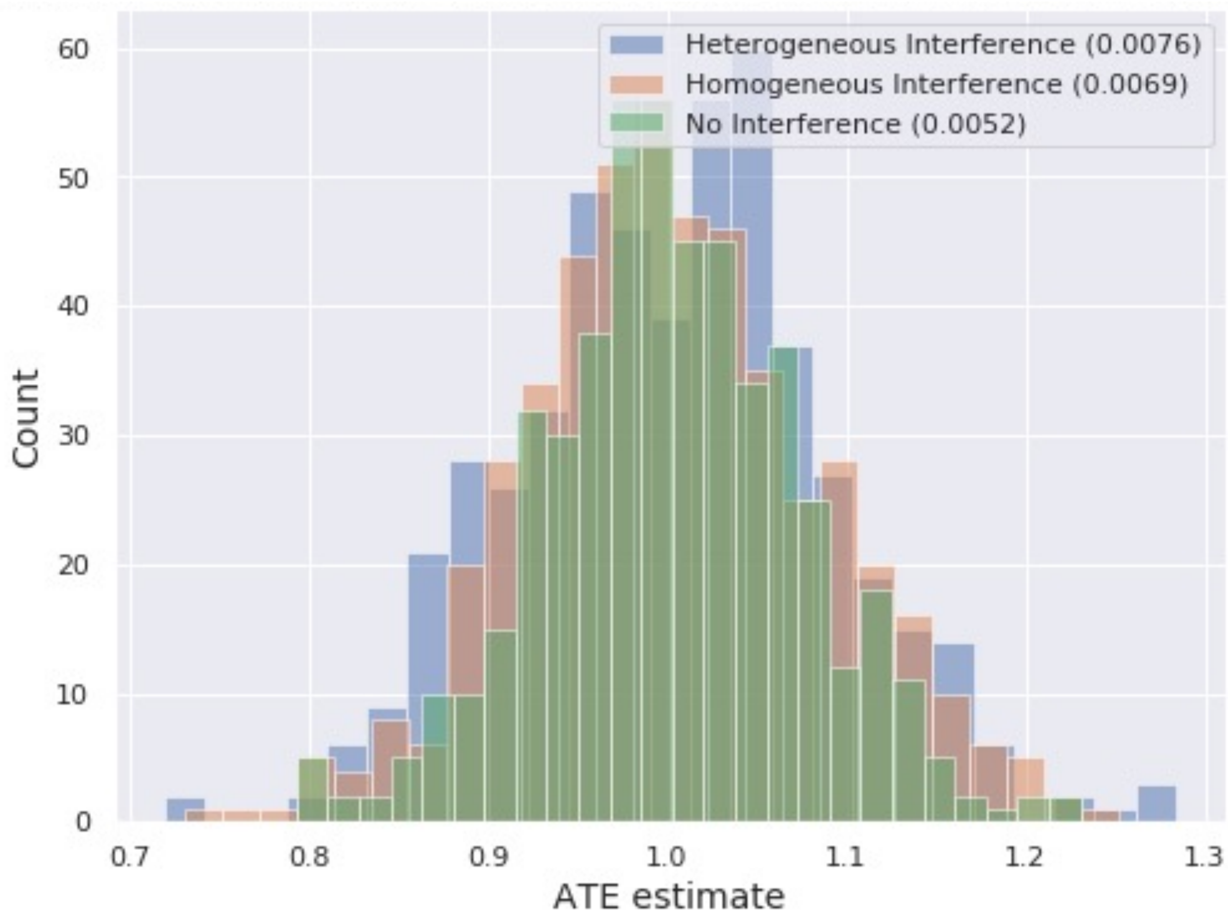

Comparison of three estimation methods on Add Health Data

- All three methods indicate that smoking has a negative impact on academic performance
- The method that is robust to heterogonous interference has a larger standard error than those assuming homogeneous inference or no interference

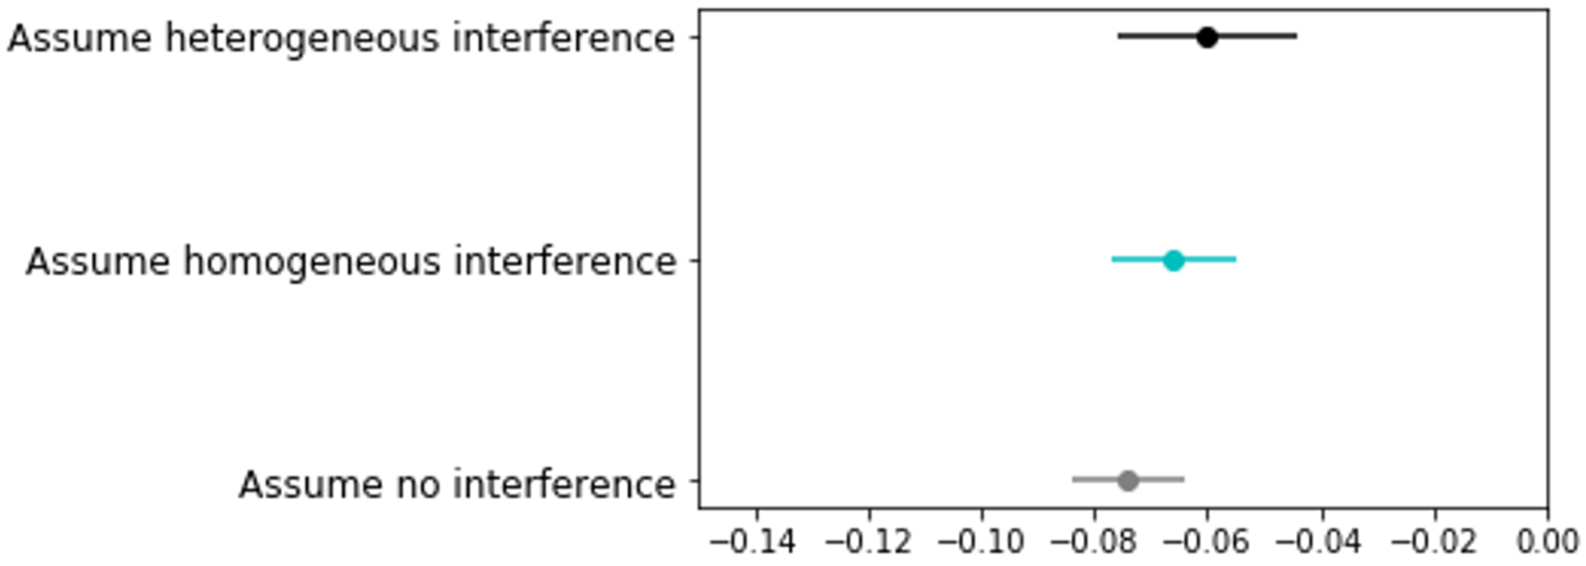

Extensions

Extensions

- Varying cluster size
  - Estimators for direct and spillover effects
  - Show consistency, asymptotic normality, and double robustness of the estimators
- Unit-dependent interference neighborhoods
- Weakly connected clusters
- Other estimands and corresponding estimators
  - ATT-type estimand: Direct and spillover effects on treated units
  - CATE-type estimand: Direct and spillover effects conditional on covariates
- Fraction of treated neighbors as continuous treatment index
  - Connection with Hudgens and Halloran (2008), Tchetgen and Vanderweele (2012)

Conclusion

- A "conditional exchangeability" framework for heterogeneous, partial interference
- Generalized AIPW estimators for direct and spillover effects in observational studies
- Show that generalized AIPW estimators are doubly-robust, asymptotically normal, semiparametric efficient
- Debiased, matching-based feasible variance estimators
- Show a bias-variance trade-off between robustness to interference and estimation efficiency
- Application to Add Health data: Demonstrate the practical relevance of our framework and methods

Supplementary slides

A toy example for heterogeneity interference

$\Pr(\boldsymbol{X}_c = \mathbf{1})$
0.5

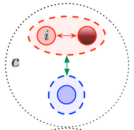

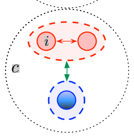

(g_1, g_2)	$p_{1,(z,g_1,g_2)}(\boldsymbol{X}_c)^*$		$\beta_1(g_1, g_2 \mid \boldsymbol{X}_c)$	
	$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$	$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$
$(1, 0)$	$1/12$	$1/6$	9	0
$(0, 1)$	$1/6$	$1/12$	0	3

$\beta_1(g_1, g_2)$
4.5
1.5

ATE
$\beta_1 = 3$

- $\beta_1(1, 0) = 9 \cdot \Pr(\boldsymbol{X}_c = \mathbf{1}) + 0 \cdot \Pr(\boldsymbol{X}_c = \mathbf{0}) = 4.5$
- $\beta_1^{p.i.h} = 4.5 \cdot \Pr(0, 1) + 1.5 \cdot \Pr(0, 1) = 3$

$*$: holds for $z = 1$ and $z = 0$

If we neglect the heterogeneity…

$\Pr(\boldsymbol{X}_c = \mathbf{1})$
0.5

	(g_1, g_2)	$p_{1,(z,g_1,g_2)}(\boldsymbol{X}_c)^*$		$\beta_1(g_1, g_2 \mid \boldsymbol{X}_c)$		$\beta_1(g_1, g_2)$
		$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$	$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$	
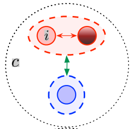	$(1, 0)$	$1/12$	$1/6$	9	0	4.5
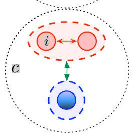	$(0, 1)$	$1/6$	$1/12$	0	3	1.5
	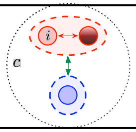 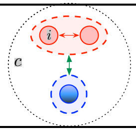			3	1	

ATE $\beta_1 = 3$

$\hat{\beta}_1^{n.i.} = \hat{\beta}_1^{p.i.} = 2$

Estimated ATE is incorrect!

- $\beta_1(\boldsymbol{X}_c = \mathbf{1}) = 9 \cdot \frac{1}{3} + 0 \cdot \frac{2}{3} = 3$
- $\hat{\beta}_1^{n.i.} = \hat{\beta}_1^{p.i.} = 3 \cdot \Pr(\boldsymbol{X}_c = \mathbf{1}) + 1 \cdot \Pr(\boldsymbol{X}_c = \mathbf{0}) = 2$

$*$: holds for $z = 1$ and $z = 0$

If we neglect the heterogeneity…

$\Pr(\boldsymbol{X}_c = \mathbf{1})$
0.5

(g_1, g_2)	$p_{1,(z,g_1,g_2)}(\boldsymbol{X}_c)^*$		$\beta_1(g_1, g_2 \mid \boldsymbol{X}_c)$	
	$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$	$\boldsymbol{X}_c = \mathbf{1}$	$\boldsymbol{X}_c = \mathbf{0}$
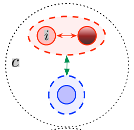 $(1, 0)$	$1/12$	$1/6$	9	0
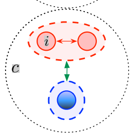 $(0, 1)$	$1/6$	$1/12$	0	3
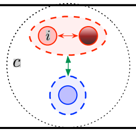 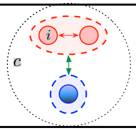			3	1

$\beta_1(g_1, g_2)$
4.5
1.5

ATE
$\beta_1 = 3$

$\hat{\beta}_1^{n.i.} = \hat{\beta}_1^{p.i.} = 2$

Estimated ATE is incorrect!

- Heterogeneous direct effects are not correctly weighted
  - E.g., the weight for $\beta_1\,(1, 0 \mid \boldsymbol{X}_c = \mathbf{1})$ is $\frac{1/12}{1/12+1/6} \cdot \Pr(\boldsymbol{X}_c = \mathbf{1}) = \frac{1}{6}$, but the correct weight is $\Pr(\boldsymbol{X}_c = \mathbf{1}) \cdot \Pr(1, 0) = 0.25$

Feasible variance estimator

- In Theorem 1, the asymptotic variance takes the form of

$$
\begin{aligned}
&V_{1,g_1,g_2} \\
&= \frac{1}{|I_1|^2} \sum_{i \in I_1} E\left[\frac{\sigma^2_{i,(1,g_1,g_2)}(\boldsymbol{X}_c)}{p_{i,(1,g_1,g_2)}(\boldsymbol{X}_c)} + \frac{\sigma^2_{i,(0,g_1,g_2)}(\boldsymbol{X}_c)}{p_{i,(0,g_1,g_2)}(\boldsymbol{X}_c)}\right] \\
&+ \frac{1}{|I_1|^2} \sum_{i,i' \in I_1} E[(\beta_i(g_1, g_2 \mid \boldsymbol{X}_c) - \beta_i(g_1, g_2)) \cdot (\beta_{i'}(g_1, g_2 \mid \boldsymbol{X}_c) - \beta_{i'}(g_1, g_2))]
\end{aligned}
$$

Analogous to the conventional term (come from observational noise)

Covariance between two units in a subset

- The plug-in estimator for V_{1,g_1,g_2} using the estimated $\beta_i(g_1, g_2 \mid \boldsymbol{X}_c)$ and $\sigma^2_{i,(1,g_1,g_2)}(\boldsymbol{X}_c)$ is **inconsistent**!
  - Bias-correction is need!

Matching-based feasible variance estimator

- For unit i in cluster c, to estimate V_{1,g_1,g_2}, among the units with

1. $\mathcal{J}_z = \left\{(c', i') : i' \in I_1, Z_{c',i'} = z, G_{c',(i'),1} = g_1, G_{c',(i'),2} = g_2\right\}$
2. In each of $\mathcal{J}_1$ and $\mathcal{J}_0$, find ℓ units whose $\left(X_{c',i'}, \boldsymbol{X}_{c',(i')}\right)$ are closest to $\left(X_{c,i}, \boldsymbol{X}_{c,(i)}\right)$, denoted by $\mathcal{J}_{1,\ell}$ and $\mathcal{J}_{0,\ell}$
3. Use $\mathcal{J}_{1,\ell}$ and $\mathcal{J}_{0,\ell}$ to estimate $\beta_i\left(g_1, g_2 \mid \boldsymbol{X}_c\right)$ and $\sigma^2_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$
4. Plug the estimators in Step 3 into $V_{1,g_1,g_2} - \frac{1}{\ell}\frac{1}{|I_1|^2 \cdot M}\sum_{z,c,i\in I_1} \sigma^2_{i,(z,g_1,g_2)}(\boldsymbol{X}_c)$

Matching-based feasible variance estimator

Theorem 3 (Consistency of matching-based feasible variance estimator). Suppose the assumptions in Theorem 1 holds. Under mild regularity conditions, the matching-based feasible variance estimator is consistent.

- Use the feasible variance estimator to construct test statistics of direct and spillover effects, or heterogeneity in interference