Causal Inference when Intervention Units and Outcome Units Differ

Fabrizia Mealli

European University Institute and University of Florence
fabrizia.mealli@eui.eu

with Georgia Papadogeorgou (UF) and Guido Imbens (Stanford)

ESEM - August 27, 2024

Introduction

- Interconnectedness of units poses challenges to policy evaluation
- Complex settings require to go beyond the common unit-to-unit interference and usual notions of proximity
- We study settings with two distinct populations of units: *intervention* units and *outcome* units
- This setting was named *bipartite interference* (*Zigler, Papadogeorgou, 2021*)

Examples of bipartite settings

1. *Air Pollution Epidemiology* (Zigler, Mealli, Forastiere, 2023)
    - Intervention units: pollution emitters (e.g., power plants) on which a treatment (e.g., installing a filter on smokestacks) can be applied.
    - Outcome units: populations residing in a geographical area where outcomes (e.g., air quality, health) are measured.
2. *Economics of housing* (Stock, 1989)
    - Intervention units: Hazardous-waste disposal sites where treatments (cleaning up or removal) can be applied.
    - Outcome units: houses where outcomes (e.g., selling prices) are measured.
3. *Education economics* (Crema, 2022)
    - Intervention units: Neighborhoods which may be exposed to openings of charter schools (treatment).
    - Outcome units: Traditional public schools (TPS) where outcome (e.g., racial segregation) are measured.

Applied literature - examples

Applied works usually use restrictions on either the *bipartite graph* or *exposure mapping* and *potential outcomes*

Model-based approaches are used that typically require specification of an outcome model as a function of an *effective treatment* defined at the level of the outcome units *(Manski, 2013, Borusyak, Hull, 2023)*

Air pollution example:

- Intervention units are the 472 coal-burning power plants
- Outcome units are the $25,553$ ZIP codes in the US
- Treatment is the indicator of having scrubbers installed
- Outcome is the number of Ischemic Heart Disease (IHD) hospitalizations over Medicare beneficiaries

Air pollution example (cont.)

- The *bipartite graph* is characterized by a $472 \times 25,553$ *weighted* adjacency matrix A denoting how much air mass originating travels from each power plant to each ZIP code
- A *bivariate treatment* for the outcome units is defined as: the treatment of the key-associated (most influencial) power plant and a linear function of the treatment statuses of all other power plants, weighted by elements of A
- Identifying assumptions are specified at the level of outcome units, which is however not the level where treatment assignment takes place

Challenges

Fully acknowledging and exploting the bipartite structure in a design-based approach requires to address non trivial differences in the formulation of:

- SUTVA-type assumptions, limiting the potential outcomes (or restricting the *model space*, *Savje et al., (2024)*
- policy relevant causal estimands, associated with implementable policies
- testing, estimation and inference

Related literature and our contributions

We contribute to recent literature on causal inference under bipartite interference (*Zigler and Papadogeorgou, 2021, Zigler et al., 2023, Harshaw et al., 2023, Doudchenko et al., 2020, Pouget-Abadie et al., 2019, Brennan et al., 2022, Harshaw et al., 2022, Song and Papadogeorgou, 2024*)

- We first discuss Fisher randomization testing for sharp null hypotheses under bipartite interference
- We then establish causal estimands for bipartite settings under a variety of *actionable* policies
- We discuss design-based estimation and inference, being agnostic and avoiding stringent restrictions on potential outcomes, allowing TEH, non-linearity, non-additivity
- We also discuss possible optimal designs and optimal policies

Formalization of Bipartite Structure

- The first population is composed of N *intervention* units, $\mathcal{N} = \{1, 2, \ldots, N\}$, $n \in \mathcal{N}$, eligible to receive a (binary) treatment, $W_n \in \mathcal{W} = \{0, 1\}$, $\mathbf{W} \in \mathcal{W}^N$ denotes the N-vector of treatment assignments.
- The second population is composed of M *outcome* units, $\mathcal{M} = \{1, 2, \ldots, M\}$.
- The outcome unit m has 2^N potential outcomes $Y_m(w) \in \mathbb{R}$, for $w \in \mathcal{W}^N$.
- The M vector of potential outcomes under the N-component treatment vector w is denoted by the M component vector

$$Y(w) = (Y_1(w), Y_2(w), \ldots, Y_M(w)) \in \mathbb{R}^M$$

.

Restricting POs according to a partially known bipartite graph

- Suppose there exists a bipartite graph $\mathcal{G}$ with two sets of nodes, $\mathcal{N}$ and $\mathcal{M}$, and edges which can only exist between an intervention and an outcome unit

- $\mathcal{G} = \{\mathcal{N}, \mathcal{M}, A\}$ where A is an $N \times M$ binary matrix, $A_{nm} = 1$ if nodes n, m are connected with respect to $\mathcal{G}$, and $A_{nm} = 0$ otherwise

- $\mathcal{G}$ is a fixed quantity that characterizes the populations

- Partially known graph because we do not assume to know the strength of the connections between nodes

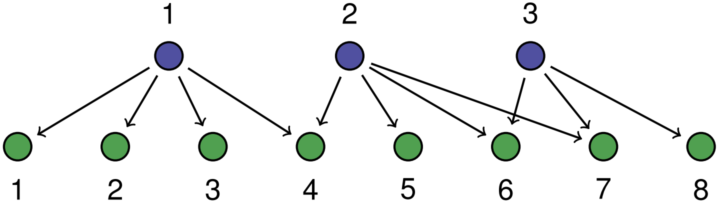

General bipartite interference graph with 3 interventional units and 8 outcome units

B-SUTVA: Bipartite Stable Unit Treatment Value Assumption

- Outcome Set for Intervention Unit n: $\mathcal{M}_n = \{m \in \mathcal{M} : A_{nm} = 1\}$
- Intervention Set for Outcome Unit m: $\mathcal{N}_m = \{n \in \mathcal{N} : A_{nm} = 1\}$
- Consistent Experience for Outcome Unit m: given w, w', an outcome unit m has a consistent experience if $w_i = w'_i$ for all all n such that $A_{nm} = 1$; that is, if $w_{\mathcal{N}_m} = w'_{\mathcal{N}_m}$

B-SUTVA: The potential outcomes for outcome unit m agree with the bipartite graph $\mathcal{G}$ if it holds: if $w, w' \in \{0, 1\}^N$ provide unit m a *consistent experience*, then $Y_m(w) = Y_m(w')$.

If B-SUTVA holds potential outcomes for unit m are denoted as $Y_m(w_{\mathcal{N}_m})$ and $Y(w) = (Y_1(w_{\mathcal{N}_1}), Y_2(w_{\mathcal{N}_2}), \ldots, Y_N(w_{\mathcal{N}_N}))$

Examples

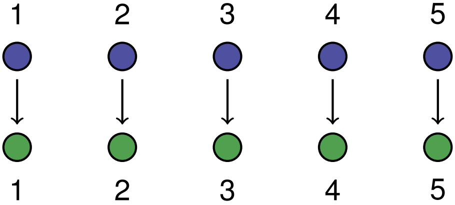

Special Case 1: $\mathcal{M} = \mathcal{N}$, A identity matrix

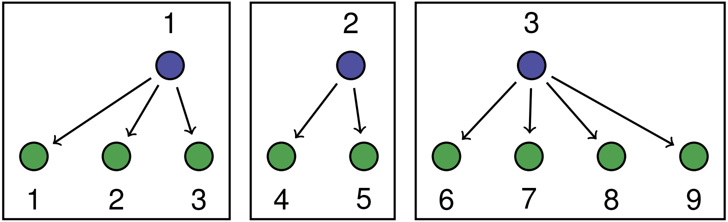

Special Case 2: $M \geq N$, $\sum_n A_{nm} = 1$ for all m (clusters)

Examples

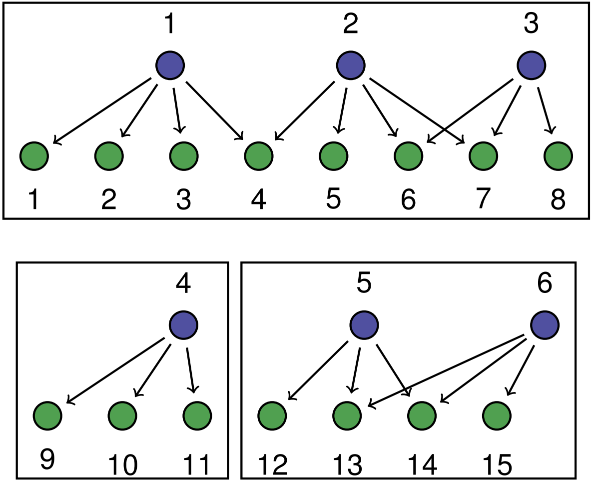

Special Case 3: Partial Bipartite Interference

Estimands

- The All Versus Nothing Effect: $\tau^{aon} = \overline{Y}(1) - \overline{Y}(0)$

$$\overline{Y}(0) = \frac{1}{M} \sum_{m=1}^{M} Y_m(0_N) \quad \text{and} \quad \overline{Y}(1) = \frac{1}{M} \sum_{m=1}^{M} Y_m(1_N)$$

- The Status Quo Effect for the Treated: $\tau_0^{sq} = \overline{Y}(W) - \overline{Y}(0)$
- The Status Quo Effect for the Controls: $\tau_1^{sq} = \overline{Y}(1) - \overline{Y}(W)$

Estimation and Inference for AoN, SQs

We consider weighting estimators,weight for each outcome unit involves the probability of the realized treatment for the unit intervention set $\mathcal{N}_m$, $\pi_{\mathcal{N}_m}$

$$\widehat{Y}(0) = \frac{1}{M} \sum_{m=1}^{M} \frac{I(W_{\mathcal{N}_m} = 0)}{\pi_{\mathcal{N}_m}(W_{\mathcal{N}_m})} Y_m$$

$$\widehat{Y}(1) = \frac{1}{M} \sum_{m=1}^{M} \frac{I(W_{\mathcal{N}_m} = 1)}{\pi_{\mathcal{N}_m}(W_{\mathcal{N}_m})} Y_m$$

Theorem

Assume that π is a treatment allocation strategy with $\pi_{\mathcal{N}_m}(\mathbf{a}) > 0$ for all m and for $\mathbf{a} = (a, a, \ldots, a)$ equal treatment level. Then, for $a = 0$ or 1,

- $\left[\widehat{Y}(a)\right] = \overline{Y}(a)$. *As a result, we also have unbiased estimators for τ and τ^{sq}.*
- $Var\left[\widehat{Y}(a)\right]$ *is equal to*

$$\frac{1}{M^2}\left[\sum_{m=1}^{M} \pi_{\mathcal{N}_m}(\mathbf{a})\left(1 - \pi_{\mathcal{N}_m}(\mathbf{a})\right)\left[\frac{Y_m(a)}{\pi_{\mathcal{N}_m}(\mathbf{a})}\right]^2 + \sum_{m \neq m'} \left(\pi_{\mathcal{N}_{m,m'}}(\mathbf{a}) - \pi_{\mathcal{N}_m}(\mathbf{a})\pi_{\mathcal{N}_{m'}}(\mathbf{a})\right) \frac{Y_m(a)}{\pi_{\mathcal{N}_m}(\mathbf{a})} \frac{Y_{m'}(a)}{\pi_{\mathcal{N}_{m'}}(\mathbf{a})}\right],$$

where $\mathcal{N}_{m,m'} = \mathcal{N}_m \cup \mathcal{N}_{m'}$ is the intervention set for outcome units m, m'.

Consistency and CLT

- Condition 1: the maximum size of the intervention set of all outcome units must not grow too fast with the number of outcome units to ensure that the positivity assumption for each outcome unit holds
- Condition 2: overlap of outcome units' interventional sets: the maximum number of common intervential units across all outcome units pairs must not grow too fast to restrict the pairs of outcome units with"correlated" treatment levels
- In general, *N* cannot grow too fast with *M*, otherwise positivity will be violated. But also *N* cannot grow too slowly with *M*, otherwise we will have too many outcome units that share interventional units, so we do not see enough variability in their exposures

Stochastic intervention effects

Alternative estimands may be easier to estimate than others, those corresponding to *smaller* changes in the distribution of assignments relative to the AoN estimands.

- For outcome unit $m \in \mathcal{M}$, we hypothesize a distribution $h_{m,\alpha} : \mathcal{W}^{N_m} \to [0, 1]$ with parameter α over the treatment assignment of its interventional set, $\mathcal{N}_m$

- $\overline{Y}_{m,h_{m,\alpha}} = \sum_{w_{\mathcal{N}_m} \in \mathcal{W}^{N_m}} Y_m(w_{\mathcal{N}_m}) h_{m,\alpha}(w)$

- For the collective stochastic intervention $h_\alpha = (h_{1,\alpha}, h_{2,\alpha}, \ldots, h_{M,\alpha})$ define:$\overline{Y}_{h_\alpha} = \frac{1}{M} \sum_{m=1}^{M} \overline{Y}_{m,h_{m,\alpha}}$

- The causal effect of switching the hypothetical distribution from parameter α to α' is $\tau(\alpha, \alpha') = \overline{Y}_{h_{\alpha'}} - \overline{Y}_{h_\alpha}$

.

Stochastic interventions: special cases

- Bernoulli interventions (*Hudgens, Halloran, 2008)*
- Completely randomized interventions
- Interventions on key-associated interventional units (*Zigler, Papdogeorgou, 2021)*
- Targeted exposures for each outcome units through exposure mapping (*Aronow, Samii, 2017; Forastiere et al., 2021)*
- The effect of treating one additional control intervention unit randomly chosen
- *Actionable* vs *Non-actionable* policies

Estimation and inference

- For stochastic intervention, h_α, define the estimator for the average potential outcome as

$$\widehat{Y}_{h_\alpha} = \frac{1}{M} \sum_{m=1}^{M} \frac{h_{m,\alpha}(W_{\mathcal{N}_m})}{\pi_{\mathcal{N}_m}(W_{\mathcal{N}_m})} Y_m$$

- Unbiasdness, Variance, Consistency under specific regimes of growth of the bipartite graph
- Compare variances of estimators of similar estimands under different restrictions on the potential outcomes

Concluding remarks and future work

- Causal inference in bipartite settings
  - ▷ Relevant for many interventions with a clear distinction between intervention and outcome units
  - ▷ Complex exposure patterns generate different types of interference, going beyond unit-to-unit or spatial interference
  - ▷ Introduce new types of causal questions and estimands
- Implications of our results: treatment effect heterogeneity and optimal policies
- Introduce optimal designs for different target estimands
- Extension to observational settings