Staggered Rollout Designs Enable Causal Inference without Graph Knowledge

Mayleen Cortez-Rodriguez [1] Matthew Eichhorn [1] Christina Lee Yu [2]

[1]Cornell University Center for Applied Mathematics [2]Cornell University Operations Research and Information Engineering

The Problem

- A company runs an experiment to estimate the effectiveness of a national ad campaign
- The **Total Treatment Effect (TTE)** estimand measures the change in the average individual's behavior when everyone sees the ad versus when no one does

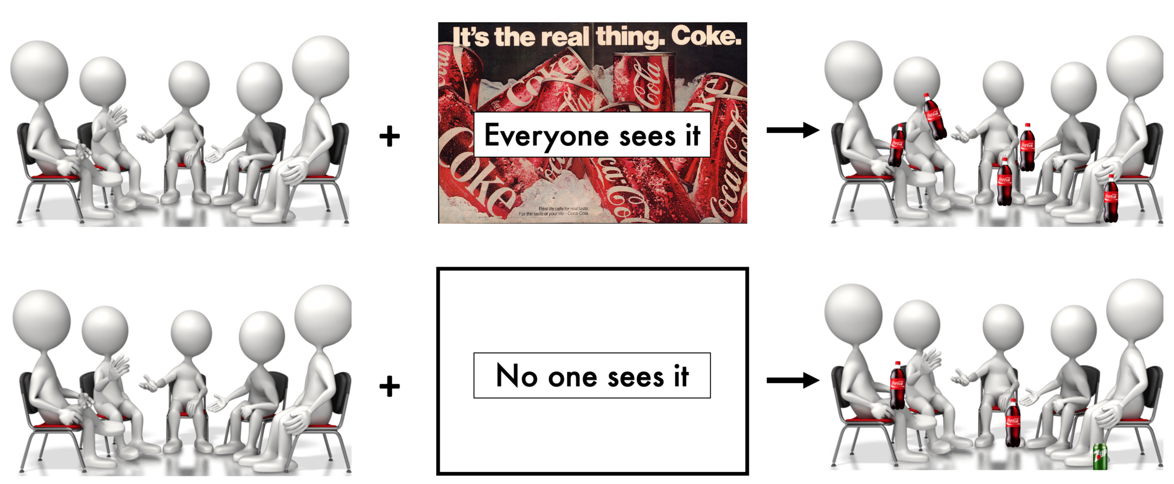

- **Network Interference:** Word-of-mouth spreads advertiser's message beyond direct viewers
- Interference violates the SUTVA assumption and introduces bias to classic estimators

Formalizing the Problem

- **Population:** Directed graph on n nodes, edges encode interference
- **Treatment:** Indicated by $\mathbf{z} \in \{0,1\}^n$
- **Outcomes:** $Y_i(\mathbf{z})$ for each person i
- Edges may be **unknown**

$$\mathrm{TTE} \triangleq \frac{1}{n}\sum_i \Big(Y_i(\mathbf{1}) - Y_i(\mathbf{0})\Big)$$

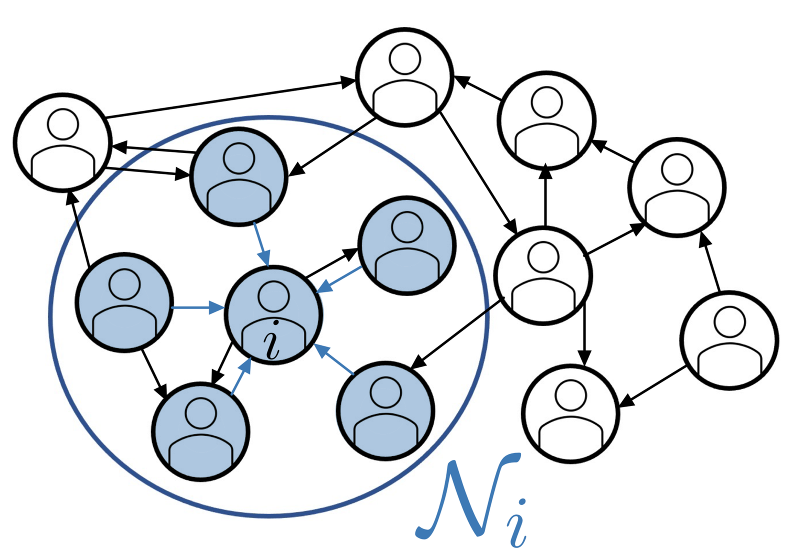

Research Question

Can we design an unbiased TTE estimator under the assumptions listed below that has a reasonable bound on its variance?

Assumptions

1. **Neighborhood Interference:** Individual i's outcome Y_i is a function only of in-neighbors' treatments $\{z_j\}_{j\in\mathcal{N}_i}$
2. **β-Order Interactions:** Only small subsets of *treated* neighbors affect i's outcome

$$Y_i(\mathbf{z}) = \underbrace{\sum_{\substack{S\subseteq\mathcal{N}_i \\ |S|\le\beta}}}_{\text{Subsets of } i\text{'s neighborhood}} \overbrace{\boxed{c_{i,S}}}^{\text{Effect of treated } S \text{ on } i} \prod_{j\in S} z_j$$

3. **Bounded Effects:** For each individual i, $\sum_{\substack{S\subseteq\mathcal{N}_i \\ |S|\le\beta}} |c_{i,S}| = \mathcal{O}(1)$
4. **Randomized Experiment Setting:** Randomly assign some individuals to treatment, some to control, and observe their outcomes

Staggered Rollout Experiment Design

Using a staggered rollout experiment design, we compensate for lack of network information by taking multiple outcome measurements.

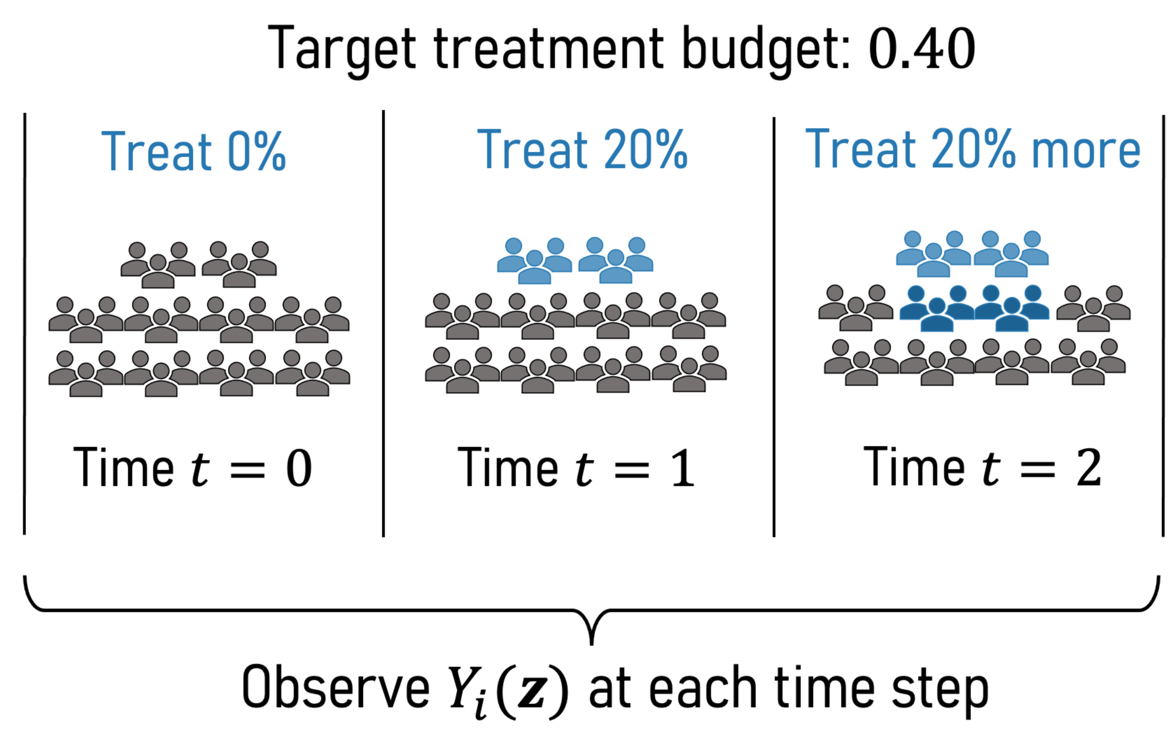

Theoretical Results

Consider the function $F(p) = \mathop{\mathbb{E}}_{\mathbf{z}\sim\mathrm{Bern}(p)}\left[\frac{1}{n}\sum_{i=1}^n Y_i(\mathbf{z})\right]$ and note:

- $\mathrm{TTE} = F(1) - F(0)$
- F is a polynomial in p with degree $\le \beta$
- Computing the average of $\{Y_i(\mathbf{z})\}_{i=1}^n$ with $\mathbf{z}\sim\mathrm{Bern}(p)$ gives an unbiased estimate of $F(p)$

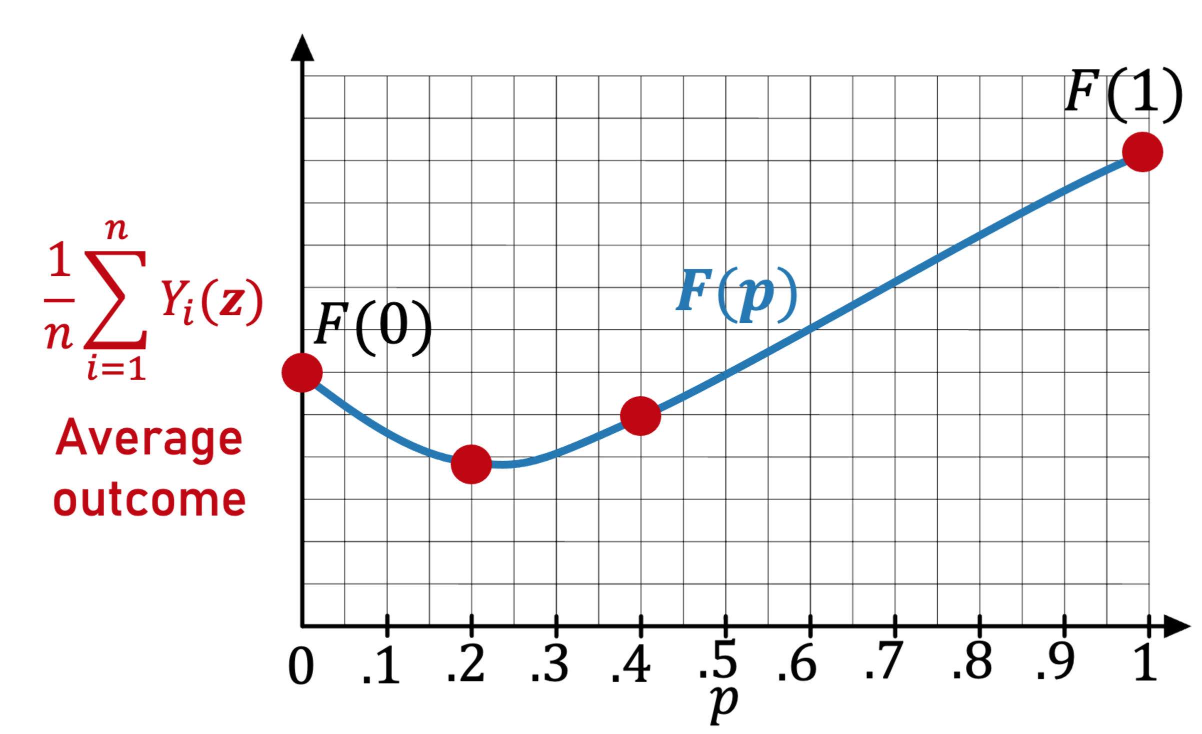

Recast TTE estimation as polynomial extrapolation after $\beta+1$ rounds:

Sample independent $u_i \sim \mathrm{Unif}(0,1)$ and define $\{\mathbf{z}^t\}_{t=0}^{\beta}$ with $z_i^t = \mathbb{I}\left(u_i \le \frac{tp}{\beta}\right)$.

Then,

$$\widehat{\mathrm{TTE}}_{\mathrm{PI}} := \frac{1}{n}\sum_{i=1}^{n}\sum_{t=0}^{\beta}\Big(\ell_t(1) - \ell_t(0)\Big)\cdot Y_i(\mathbf{z}^t), \qquad \ell_t(p) = \prod_{\substack{s=0 \\ s\ne t}}^{\beta}\frac{\beta-s}{t-s}$$

is an unbiased estimator for TTE with variance $\mathcal{O}\left(\frac{d^2\beta^2}{n}\cdot\left(\frac{\beta}{p}\right)^{2\beta}\right)$.

We also have results for when…

- Observations of $Y_i(\mathbf{z})$ are perturbed by mean-0 independent Gaussian noise
- $\mathbf{z}$ is obtained by uniformly sampling a subset of k individuals to treat

Experiments

- Configuration model network with in-degrees distributed as a power law with exponent 2.5
- Parameter r governs the strength of interference effects
- Parameter p is the treatment budget
- Compare against difference-in-means (DM) and least-squares (LS) estimators
- **Observation**: Under a β-order potential outcomes model, our estimator PI(p) is unbiased with lower variance than the other estimators
- **Key Point**: Our estimator uses no graph knowledge

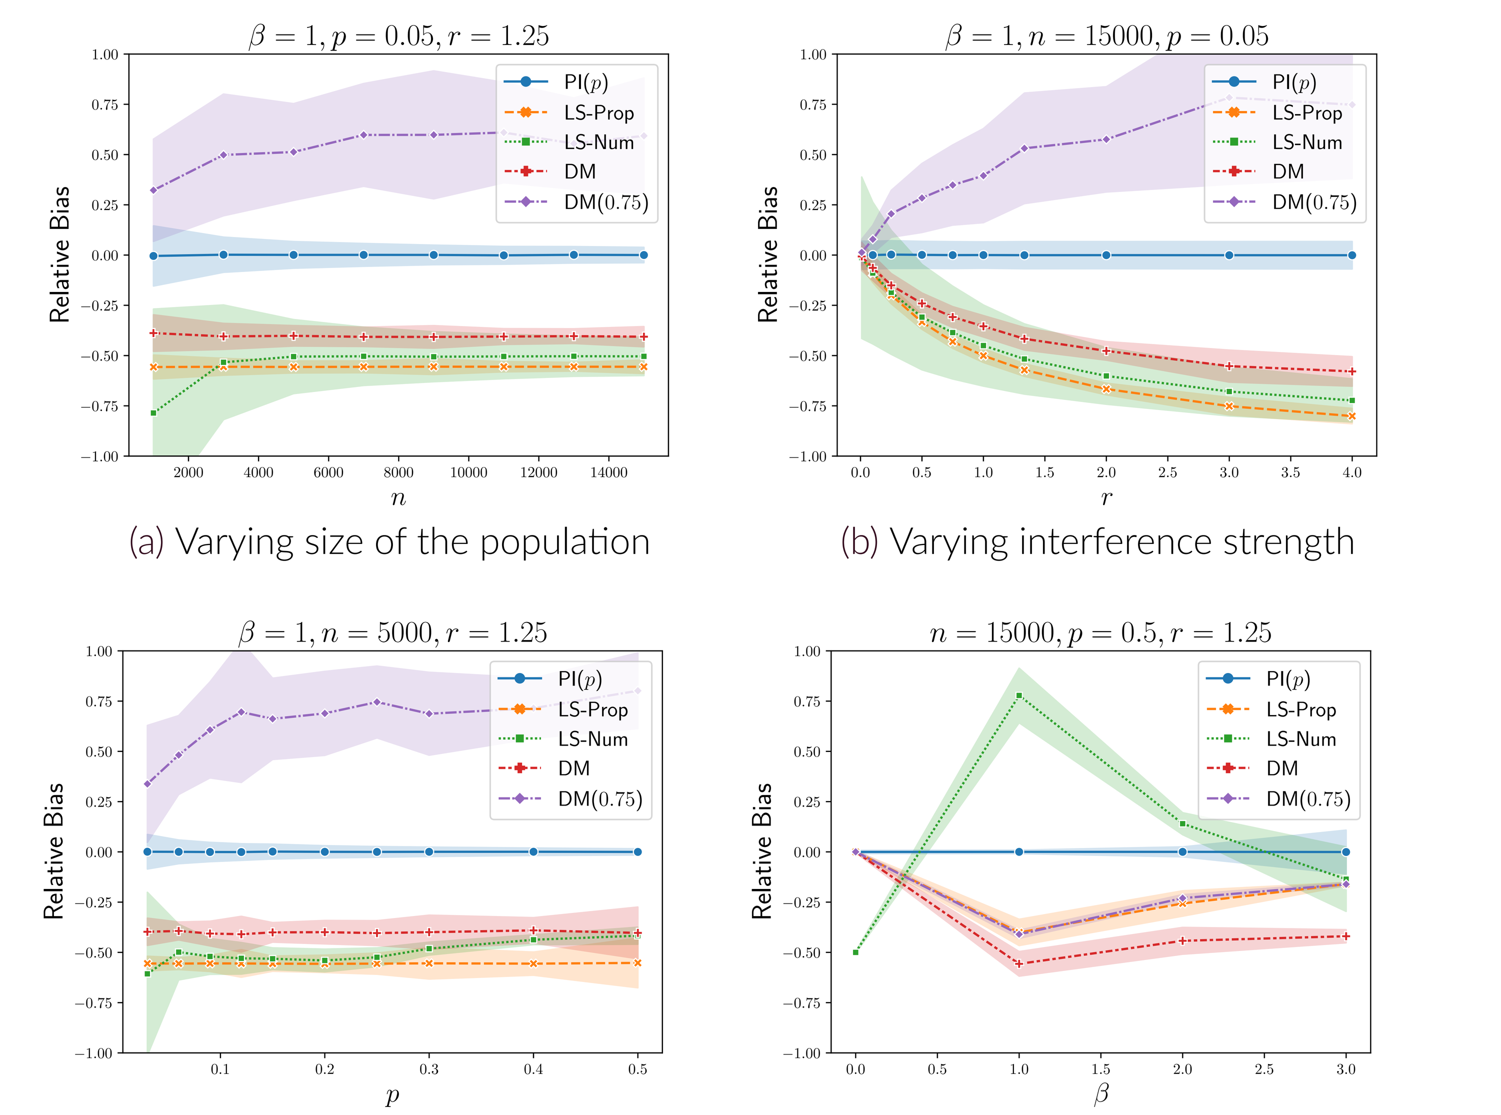

(a) Varying size of the population

(b) Varying interference strength

(c) Varying treatment budget

(d) Varying the model degree

Future Research Directions

- Run experiments on real-world data
- Allow for time-varying effects or time-varying networks
- Bias-variance trade off results when β is unknown

Acknowledgements

We gratefully acknowledge financial support from the National Science Foundation grants CCF-1948256 and CNS-1955997 and the Sloan Foundation grant 90855. Dr. Yu is also supported by an Intel Rising Stars award and a JPMorgan Faculty Research award. Mayleen is also supported by the National Science Foundation Graduate Research Fellowship Program under grant DGE-1650441.

References

[1] Mayleen Cortez, Matthew Eichhorn, and Christina Lee Yu. Staggered rollout designs enable causal inference under interference without network knowledge. *arXiv preprint arXiv:2205.14552*, 2022.