Network Interference in Micro-Randomized Trials

Shuangning Li

Stanford University

Nathan's Group Meeting

Joint work with

Stefan Wager

Li, S., & Wager, S. (2022). **Network Interference in Micro-Randomized Trials.** *arXiv preprint arXiv:2202.05356*.

Motivation

- Effect of motivational message on level of physical activities

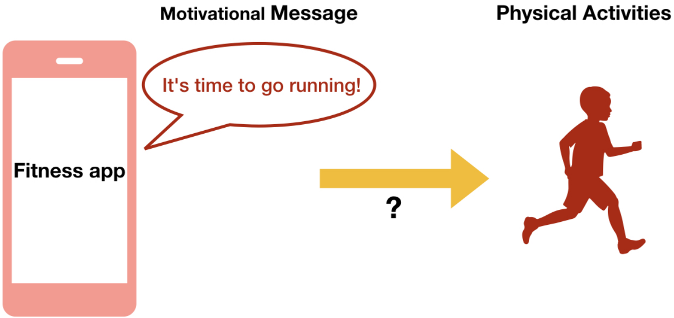

Motivation

- Effect of motivational message on level of physical activities

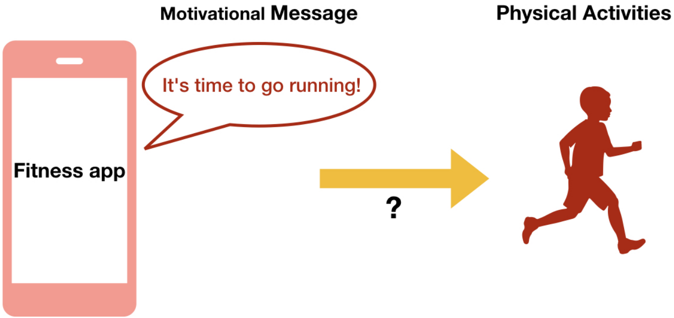

- Seems to be easy: Run RCTs?

Interference

- An individual's outcome could depend on other individuals' outcomes.

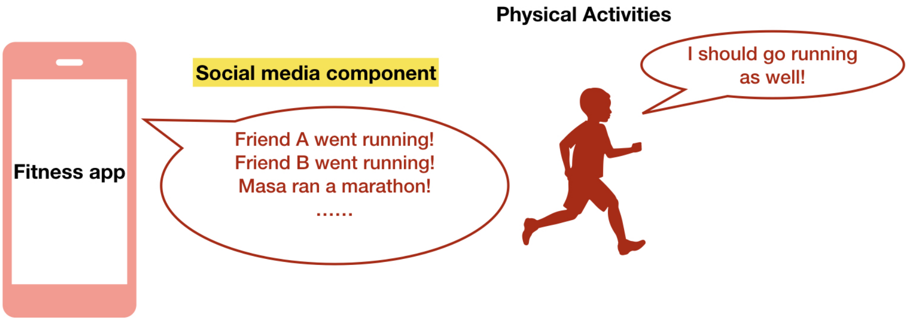

Dynamic Nature

- An additional time axis.

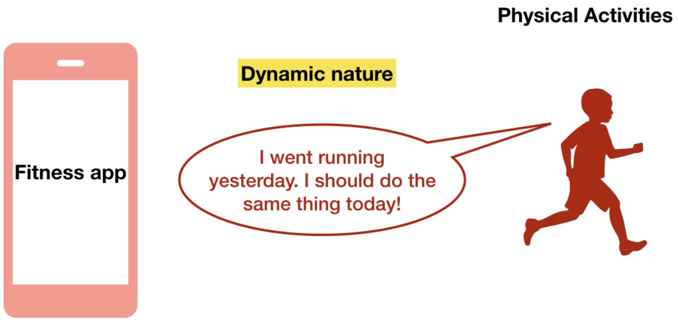

Micro-Randomized Trials

- The micro-randomized trial (MRT) is an experimental design that is often used to help evaluate and optimize dynamic interventions.
- We focus on Bernoulli treatments.

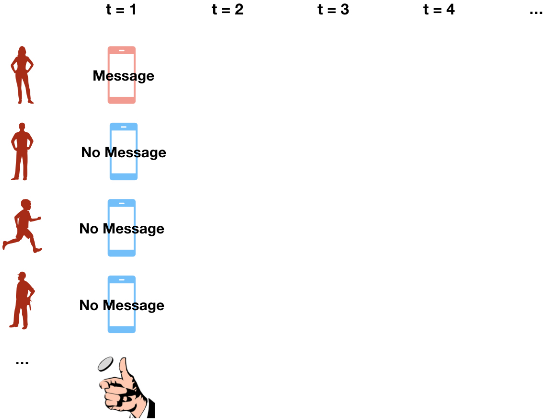

Micro-Randomized Trials

- The micro-randomized trial (MRT) is an experimental design that is often used to help evaluate and optimize dynamic interventions.
- We focus on Bernoulli treatments.

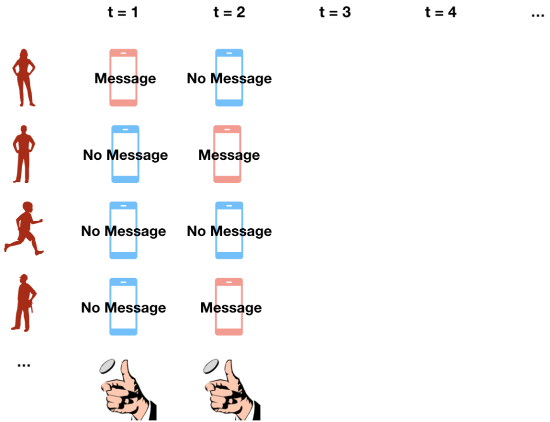

Micro-Randomized Trials

- The micro-randomized trial (MRT) is an experimental design that is often used to help evaluate and optimize dynamic interventions.
- We focus on Bernoulli treatments.

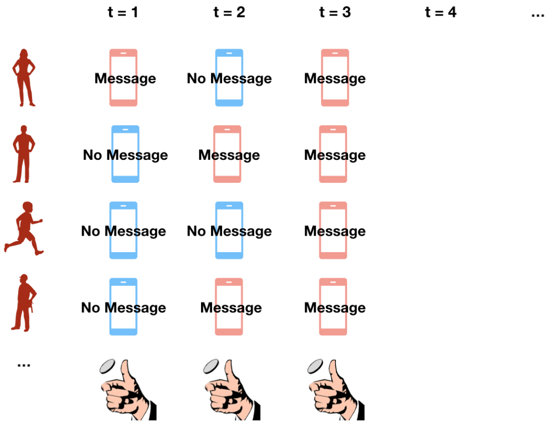

Micro-Randomized Trials

- The micro-randomized trial (MRT) is an experimental design that is often used to help evaluate and optimize dynamic interventions.
- We focus on Bernoulli treatments.

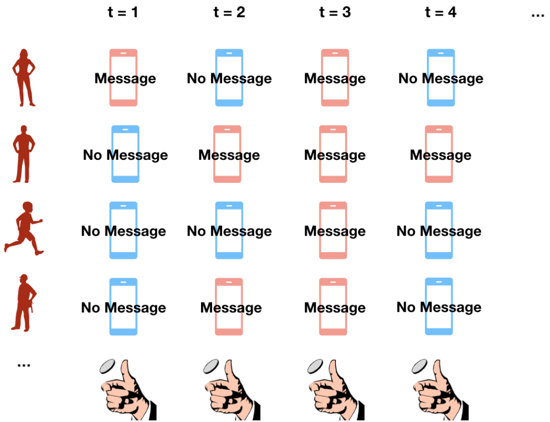

Problem Setup

- There are n subjects of interests indexed by $i = 1, \ldots, n$.
- Assume that there is an undirected graph with indices corresponding to the n subject. We call the undirected graph the *interference network* or the *interference graph*. We use $\{E_{ij}\}$ to denote the edge set of the graph.
- Let $\mathcal{N}_i = \{j : E_{ij} = 1\}$ be the set of neighbors of subject i.
- Let $Y_{it} \in \{0, 1\}$ denote the outcome of interest at time t.
- Let $W_{it} \in \{0, 1\}$ be the treatment at time t.

Assumptions

Assumption 1 (MDP with Network Interference)
Each unit $i = 1, \ldots, n$ is characterized by an activation function $f_i(\cdot)$ such that, conditionally on $Y_{1:t}$ and $W_{1:t}$,

$$Y_{i(t+1)} \sim \text{Ber}(f_i(Y_{it}, W_{it}, Z_{it})) \text{ independently,}$$

where $Z_{it} = \sum_{j \in \mathcal{N}_i} Y_{jt}$, and $0 < f_i(y, w, z) < 1$ for all $y, w \in \{0, 1\}$ and $z \in \mathbb{R}_+$.

- Fitness app example: The probability of an individual goes running tomorrow depends on
  - whether she went running today (Y_{it}),
  - whether she receives any encouraging message from the app (W_{it}),
  - the total number of her friends that went running today (Z_{it}),
  - her individual characteristics (f_i).

Assumptions

Assumption 1 (MDP with Network Interference)
Each unit $i = 1, \dots, n$ is characterized by an activation function $f_i(\cdot)$ such that, conditionally on $Y_{1:t}$ and $W_{1:t}$,

$$Y_{i(t+1)} \sim \text{Ber}(f_i(Y_{it}, W_{it}, Z_{it})) \text{ independently,}$$

where $Z_{it} = \sum_{j \in \mathcal{N}_i} Y_{jt}$, and $0 < f_i(y, w, z) < 1$ for all $y, w \in \{0, 1\}$ and $z \in \mathbb{R}_+$.

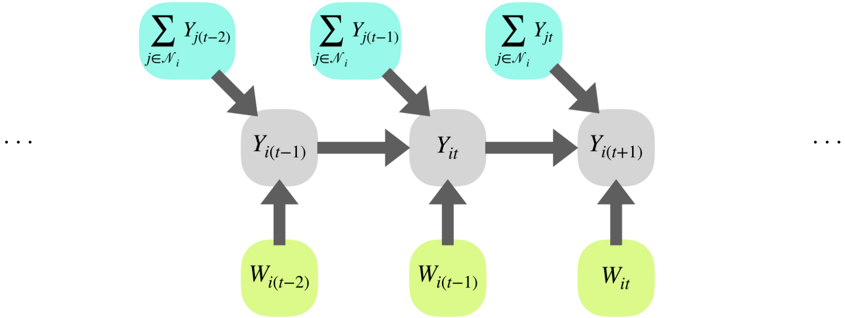

Assumptions

Assumption 2 (Bernoulli treatment)
The treatments $W_{it} \sim \text{Ber}(\pi_i)$ independently for each i and each t.

Causal Estimands

- **Short-term direct effect.** The short-term direct effect quantifies the immediate effect of the unit's treatment on its own outcome.

$$\tau_{\mathrm{SDE},t} = \frac{1}{n}\sum_{i=1}^{n} f_i\left(Y_{it}, 1, Z_{it}\right) - f_i\left(Y_{it}, 0, Z_{it}\right).$$

- **Long-term direct effect.** The long-term direct effect captures the long-term effect of the unit's treatment on its own outcome, averaged over units.

$$\tau_{\mathrm{LDE}}(\gamma_1, \gamma_2) = \frac{1}{n}\sum_{i=1}^{n} \left(\mathbb{E}_{\mu(\pi_i=\gamma_1, \pi_{-i})}\left[Y_i\right] - \mathbb{E}_{\mu(\pi_i=\gamma_2, \pi_{-i})}\left[Y_i\right]\right).$$

Here $\mu(\pi_i = \gamma, \pi_{-i})$ stands for the stationary distribution of the MDP when the treatment probability of the i-th unit has been changed to γ.

Causal Estimands

- **Long-term total effect.**
  It measures the effect of changing the entire treatment vector from π_2 to π_1 on the expected average outcome under the stationary distribution.

$$\tau_{\mathsf{LTE}}(\pi_1, \pi_2) = \frac{1}{n} \sum_{i=1}^{n} \left(\mathbb{E}_{\mu(\pi_1)}[Y_i] - \mathbb{E}_{\mu(\pi_2)}[Y_i] \right).$$

Existence of a stationary distribution

Since both Y_{it} and W_{it} are binary random variables, the function f_i can be decomposed into four terms:

$$f_i(y, w, z) = a_i(z) + b_i(z)w + c_i(z)y + d_i(z)wy.$$

Additional assumptions:

- **Boundedness and Lipschitzness**. Each function f_i is Lipschitz with Lipschitz constant L_n in its third argument. The functions c_i and d_i satisfy $|c_i(z) + d_i(z)w| \leq B$ for any $z \in \mathbb{R}_+$ and any $w \in \{0, 1\}$.
- **Node Degree**. The largest node degree of the interference graph is bounded by D_n.
- **Contraction** The constants B, L_n and D_n satisfy $B + L_n D_n \leq C < 1$.

Example: Erdős-Rényi
Each edge is included in the interference graph with probability ρ_n, independently from every other edge, i.e., $E_{ij} \sim \text{Bernoulli}(\rho_n)$ independently.

$$L_n \sim \frac{1}{n\rho_n}, \qquad D_n \sim n\rho_n.$$

Existence of a stationary distribution

Proposition (Stationary distribution)

1. *Under the above Assumptions, there* **exists** *a stationary distribution $\mu(\pi)$, such that if $Y_0 \sim \mu(\pi)$, then $Y_t \sim \mu(\pi)$ for any $t \geq 0$.*
2. *Furthermore, the Markov chain (induced by the Bernoulli policy) is ergodic, i.e., the stationary distribution $\mu(\pi)$ is* **unique**. *For any initial distribution of Y_0, we have that $Y_t \Rightarrow \mu(\pi)$ as $t \to \infty$.*

Mean-Field Characterization

- Difficulties:
  - Under the stationary distribution μ, outcome Y_i's are not independent.
  - The state space is of size 2^n.
- Consider a dynamical system—a "mean-field" version of the Markov Chain.
- Let $P_t = (P_{1t}, P_{2t}, \ldots, P_{nt}) \in [0,1]^n$ be the state of the dynamical system at time t. The evolution rule of the system is the following:

$$\begin{aligned} P_{i(t+1)} &= f_i(P_{it}, \pi_i, Q_{it}) \\ &= a_i(Q_{it}) + b_i(Q_{it})\pi_i + c_i(Q_{it})P_{it} + d_i(Q_{it})\pi_i P_{it}, \end{aligned}$$

  where $Q_{it} = \sum_{j \in \mathcal{N}_i} P_{jt}$.
- We can interpret $P_{i,t}$ as the **probability** that $Y_{i,t} = 1$ given past information. Here, the state of the i-th user depends directly on its neighbors probabilities rather than their realized outcomes.

Mean-Field Characterization

Markov chain

$$Y_{i(t+1)} \sim \text{Ber}\left(f_i(Y_{it}, W_{it}, \sum_{j \in \mathcal{N}_i} Y_{jt})\right)$$

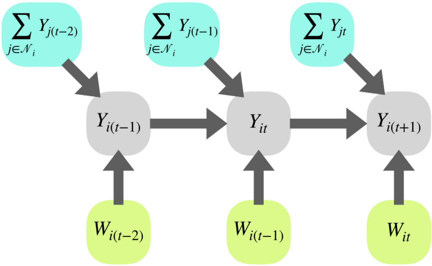

Dynamical system

$$P_{i(t+1)} = f_i(P_{it}, \pi_i, \sum_{j \in \mathcal{N}_i} P_{jt})$$

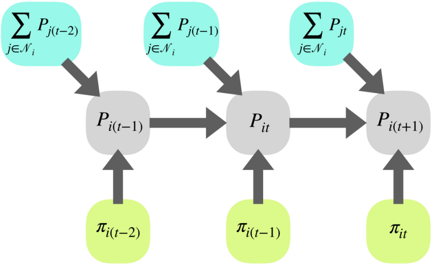

Mean-Field Characterization

- Easier to study
  - The probabilities are non-random numbers, thus if the outcomes $Y_{it} \sim \text{Bern}(P_{it})$ independently, then the Y_{it}'s are independent.
  - The fixed point of the dynamical system can be characterized by a vector of length n.

Proposition (Fixed point)

1. *There* **exists** *a fixed point $P^\star \in [0,1]^n$ of the dynamical system, i.e., if $P_t = P^\star$, then $P_{t+1} = P^\star$.*
2. *The fixed point is* **unique**, *and for any value of P_0, we have $P_t \to P^\star$ as $t \to \infty$.*

How good is the mean-field approximation?

- $Y \sim \mu(\pi)$, where $\mu(\pi)$ is the stationary distribution of the Markov chain.
- $Y_i^\star \sim \mathrm{Ber}(P_i^\star)$ independently, where $P^\star$ is the fixed point of the dynamical system.
- Is Y close to $Y^\star$ in distribution?
- Define an L_1-Wasserstein distance between two laws ν_1 and ν_2

$$W_{L_1}(\nu_1, \nu_2) = \inf\left\{\mathbb{E}\left[\|X_1 - X_2\|_1\right] : \mathcal{L}(X_1) = \nu_1, \mathcal{L}(X_2) = \nu_2\right\}.$$

Theorem (Mean field approximation)

$$W_{L_1}(\mu, \mathcal{L}(Y^\star)) \le n\sqrt{L_n}\sqrt{C}/(2(1-C))$$

Short-term direct effect

- Recall that

$$\tau_{\mathsf{SDE},t} = \frac{1}{n}\sum_{i=1}^{n} f_i\left(Y_{it}, 1, Z_{it}\right) - f_i\left(Y_{it}, 0, Z_{it}\right).$$

- A natural estimator to use here is the inverse propensity weighted (IPW) estimator:

$$\hat{\tau}_{\mathsf{IPW},t} = \frac{1}{n}\sum_{i=1}^{n} Y_{i(t+1)}\left(\frac{W_t}{\pi_i} - \frac{1-W_t}{1-\pi_i}\right).$$

- The IPW estimator is consistent for the short-term direct effect.

Theorem (Short-term direct effect estimation)

$$\hat{\tau}_{\mathsf{IPW}} = \tau_{\mathsf{SDE},t} + \mathcal{O}_p\left(\frac{1}{\sqrt{n}}\right).$$

Long-term direct effect

- Recall that

$$\tau_{\text{LDE}}(\gamma_1, \gamma_2) = \frac{1}{n} \sum_{i=1}^{n} \left(\mathbb{E}_{\mu(\pi_i=\gamma_1, \pi_{-i})}\left[Y_i\right] - \mathbb{E}_{\mu(\pi_i=\gamma_2, \pi_{-i})}\left[Y_i\right] \right).$$

- Under this stationary distribution,

$$\begin{aligned}
&\mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[Y_i\right] \\
&\quad = \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[f_i(Y_i, W_i, Z_i)\right] \\
&\quad = \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[a_i(Z_i) + b_i(Z_i)W_i + c_i(Z_i)Y_i + d_i(Z_i)W_iY_i\right] \\
&\quad \approx \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[a_i(Z_i)\right] + \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[b_i(Z_i)\right]\gamma \\
&\qquad + \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[c_i(Z_i)\right]\mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[Y_i\right] + \mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[d_i(Z_i)\right]\gamma\mathbb{E}_{\mu(\pi_i=\gamma, \pi_{-i})}\left[Y_i\right],
\end{aligned}$$

if W_i, Y_i and Z_i are roughly independent.

Long-term direct effect

- Moving all terms involving $\mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[Y_i]$ to the left hand side, we get

$$\mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[Y_i] \approx \frac{\mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[a_i(Z_i)] + \mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[b_i(Z_i)]\,\gamma}{1 - \mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[c_i(Z_i)] - \mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[d_i(Z_i)]\,\gamma}.$$

- If we assume that Z_i's are not influenced too much by the treatment probability of unit i, then

$$\mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}[Y_i] \approx \frac{\mathbb{E}_{\mu(\pi)}[a_i(Z_i)] + \mathbb{E}_{\mu(\pi)}[b_i(Z_i)]\,\gamma}{1 - \mathbb{E}_{\mu(\pi)}[c_i(Z_i)] - \mathbb{E}_{\mu(\pi)}[d_i(Z_i)]\,\gamma}.$$

Long-term direct effect

The above heuristic in fact correctly recovers the long-term direct effect.

Theorem (Long-term direct effect characterization)

$$\mathbb{E}_{\mu(\pi_i=\gamma,\pi_{-i})}\left[Y_i\right] = \frac{\mathbb{E}_{\mu(\pi)}\left[a_i(Z_i) + b_i(Z_i)\gamma\right]}{\mathbb{E}_{\mu(\pi)}\left[1 - c_i(Z_i) - d_i(Z_i)\gamma\right]} + \mathcal{O}\left(\sqrt{L_n}\right).$$

- It suffices to estimate $\mathbb{E}_{\mu(\pi)}\left[a_i(Z_{it})\right], \ldots, \mathbb{E}_{\mu(\pi)}\left[d_i(Z_{it})\right]$ for each i.

Long-term direct effect

- It suffices to estimate $\mathbb{E}_{\mu(\pi)}[a_i(Z_{it})], \ldots, \mathbb{E}_{\mu(\pi)}[d_i(Z_{it})]$ for each i.
- Recall that

$$f_i(y, w, z) = a_i(z) + b_i(z)w + c_i(z)y + d_i(z)wy.$$

- Take $\mathbb{E}_{\mu(\pi)}[a_i(Z_{it})]$ as an example. Let

$$\hat{a}_i = \frac{\frac{1}{T}\sum_{t=1}^{T} Y_{i(t+1)}(1 - W_{it})(1 - Y_{it})}{\frac{1}{T}\sum_{t=1}^{T}(1 - W_{it})(1 - Y_{it})}.$$

Proposition

$$\mathbb{E}\left[\left(\hat{a}_i - \mathbb{E}_{\mu(\pi)}[a_i(Z_{it})]\right)^2\right] \leq C_1\left(\frac{1}{T} + L_n\right),$$

for some constant C_1 not depending on i or n.

- Challenge: Y_{it}'s are not independent.

Long-term direct effect

- Let

$$\hat{\tau}_{\mathrm{LDE}}(\gamma_1, \gamma_2) = \frac{1}{n}\sum_{i=1}^{n} \frac{\hat{a}_i + \hat{b}_i\gamma_1}{1 - \hat{c}_i - \hat{d}_i\gamma_1} - \frac{1}{n}\sum_{i=1}^{n} \frac{\hat{a}_i + \hat{b}_i\gamma_2}{1 - \hat{c}_i - \hat{d}_i\gamma_2}.$$

Corollary (Long-term direct effect estimation)

$$\hat{\tau}_{\mathrm{LDE}}(\gamma_1, \gamma_2) - \tau_{\mathrm{LDE}}(\gamma_1, \gamma_2) = \mathcal{O}_p\left(\frac{1}{\sqrt{T}} + \sqrt{L_n}\right).$$

Long-term total effect

- Recall that

$$\tau_{\mathrm{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) = \frac{1}{n} \sum_{i=1}^{n} \left(\mathbb{E}_{\mu(\pi+\Delta \boldsymbol{v})} \left[Y_i \right] - \mathbb{E}_{\mu(\pi)} \left[Y_i \right] \right).$$

  where $\|\boldsymbol{v}\| = \sqrt{n}$.
- Hard to analyze the stationary distribution: start with fixed point of the dynamical system
- When Δ is small,

$$P_i^\star(\pi + \Delta \boldsymbol{v}) - P_i^\star(\pi) \approx \Delta \frac{d}{d\Delta} P_i^\star(\pi + \Delta \boldsymbol{v}) \approx \Delta \left(\nabla_\pi P_i^\star(\pi)^\top \boldsymbol{v} \right).$$

- Recall that P^* satisfies

$$P_i^\star(\pi) = a_i \left(Q_i^\star(\pi) \right) + b_i \left(Q_i^\star(\pi) \right) \pi_i + c_i \left(Q_i^\star(\pi) \right) P_i^\star(\pi) + d_i \left(Q_i^\star(\pi) \right) \pi_i P_i^\star(\pi),$$

  where $Q_i^\star(\pi) = \sum_{j \in \mathcal{N}_i} P_j^\star(\pi)$.

Long-term total effect

- $P_i^\star(\pi) = a_i\left(Q_i^\star(\pi)\right) + b_i\left(Q_i^\star(\pi)\right)\pi_i + c_i\left(Q_i^\star(\pi)\right)P_i^\star(\pi) + d_i\left(Q_i^\star(\pi)\right)\pi_i P_i^\star(\pi)$,
- Let $p_i^\star(\pi) = \nabla_\pi P_i^\star(\pi)^\intercal \boldsymbol{v}$. Taking derivative on both hand side gives:

$$\begin{aligned} p_i^\star(\pi) = &\, b_i\left(Q_i^\star(\pi)\right)v_i + d_i\left(Q_i^\star(\pi)\right)P_i^\star(\pi)v_i + \left[c_i\left(Q_i^\star(\pi)\right) + d_i\left(Q_i^\star(\pi)\right)\pi_i\right] p_i^\star(\pi) + \\ &\left[a_i'\left(Q_i^\star(\pi)\right) + b_i'\left(Q_i^\star(\pi)\right)\pi_i + c_i'\left(Q_i^\star(\pi)\right)P_i^\star(\pi) + d_i'\left(Q_i^\star(\pi)\right)\pi_i P_i^\star(\pi)\right]\sum_{j\in\mathcal{N}_i} p_j^\star(\pi). \end{aligned}$$

- We have a set of linear equations for $p_i^\star(\pi)$'s.
- Solving for $p_i^\star(\pi)$ gives

$$p^\star(\pi) = (I - DA - W)^{-1}\boldsymbol{u},$$

where A is the adjacency matrix of the interference graph,
$D = \operatorname{diag}\left(a_i'\left(Q_i^\star(\pi)\right) + b_i'\left(Q_i^\star(\pi)\right)\pi_i + c_i'\left(Q_i^\star(\pi)\right)P_i^\star(\pi) + d_i'\left(Q_i^\star(\pi)\right)\pi_i P_i^\star(\pi)\right)$,
$W = \operatorname{diag}\left(c_i\left(Q_i^\star(\pi)\right) + d_i\left(Q_i^\star(\pi)\right)\pi_i\right)$, and $\boldsymbol{u} = \operatorname{vec}\left(v_i(b_i(Q_i^\star(\pi)) + d_i\left(Q_i^\star(\pi)\right)P_i^\star(\pi))\right)$.

Long-term total effect

- Summarizing the above findings, we have

$$\widetilde{\tau}_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) = \frac{1}{n}\sum_{i=1}^{n}\left(P_i^\star(\pi + \Delta \boldsymbol{v}) - P_i^\star(\pi)\right) \approx \frac{\Delta}{n}\sum_{i=1}^{n}\left(\nabla_\pi P_i^\star(\pi)^\top \boldsymbol{v}\right) = \frac{\Delta}{n}\sum_{i=1}^{n} p_i^\star(\pi)$$
$$= \frac{\Delta}{n}\mathbf{1}^\top p^\star(\pi) = \frac{\Delta}{n}\mathbf{1}^\top (I - DA - W)^{-1}\boldsymbol{u}.$$

- Questions remaining:
  - Is $\tau_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi)$ close to $\widetilde{\tau}_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi)$?
  - How good is the above approximation?

Long-term total effect

Theorem

$$\tau_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) = \widetilde{\tau}_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) + \mathcal{O}\left(\sqrt{L_n}\right).$$

Theorem (Long-term total effect characterization)

Under some additional smoothness assumption,

$$\frac{\tau_{\mathsf{LTE}}(\pi + \Delta_n \boldsymbol{v}, \pi)}{\Delta_n} = \frac{1}{n}\mathbf{1}^\top (I - DA - W)^{-1}\boldsymbol{u} + \mathcal{O}\left(\frac{\sqrt{L_n}}{\Delta_n} + \Delta_n\right).$$

Long-term total effect

- The above shows that

$$\tau_{\mathrm{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) \approx \frac{\Delta}{n} \mathbf{1}^{\intercal}(I - DA - W)^{-1}\boldsymbol{u}.$$

- The estimation strategy is to estimate the components D, W and $\boldsymbol{u}$ separately.
- This involves estimating $a_i'(Q_i^\star(\pi))$, $b_i'(Q_i^\star(\pi))$, $\ldots$.
- We estimate $a_i'(Q_i^\star(\pi))$ by running a regression of $Y_{i(t+1)}$ on Z_{it} conditioning on $Y_{it} = 0$ and $W_{it} = 0$ using data from different time points.
- Challenges:
  - a_i is not linear,
  - data points are not i.i.d., and thus standard results from linear regression do not apply directly.

Long-term total effect

- We define

$$\hat{\tau}_{\mathsf{LTE}}(\pi + \Delta \boldsymbol{v}, \pi) = \frac{\Delta}{n} \mathbf{1}^\top (I - \hat{D}A - \hat{W})^{-1} \hat{\boldsymbol{u}}.$$

Theorem (Long-term total effect estimation)

Assume that Δ varies with n. We write Δ_n to emphasize such dependency.

$$\frac{\hat{\tau}_{\mathsf{LTE}}(\pi + \Delta_n \boldsymbol{v}, \pi)}{\Delta_n} = \frac{\tau_{\mathsf{LTE}}(\pi + \Delta_n \boldsymbol{v}, \pi)}{\Delta_n} + \widetilde{\mathcal{O}}_p\left(\frac{1}{\Delta_n \sqrt{D_n}} + \Delta_n + \left(\frac{D_n}{\sqrt{T}} + \frac{1}{\sqrt{D_n}}\right)\right).$$

Thank you!