Policy Design in Experiments with (Unknown) Interference

Davide Viviano
UC San Diego/Stanford GSB

NBER SI 2022 (Labor Studies)

July, 2022

Introductory Example

Ex. Informing farmers exposed to environmental disasters to increase insurance take-up.

Introductory Example

Ex. Informing farmers exposed to environmental disasters to increase insurance take-up.

How should we design information campaigns?

⇒ Choosing how many people (and whom) to treat.

(i) Spillovers among farmers in the same village (e.g., Cai et al., 2015);

(ii) Treatment can be costly: treating each individual is sub-optimal;

(iii) And... network data can be difficult or infeasible to collect.

Introductory Example

Ex. Informing farmers exposed to environmental disasters to increase insurance take-up.

How should we design information campaigns?

$\Rightarrow$ Choosing how many people (and whom) to treat.

(i) Spillovers among farmers in the same village (e.g., Cai et al., 2015);

(ii) Treatment can be costly: treating each individual is sub-optimal;

(iii) And... network data can be difficult or infeasible to collect.

Q1 Does what the policy maker is currently doing maximize benefits net of costs/can we conduct **inference** on policy-optimality?

Q2 Measure that indicates the direction for an improvement/can we **estimate** the best policy?

Introductory Example

Ex. Informing farmers exposed to environmental disasters to increase insurance take-up.

How should we design information campaigns?

⇒ Choosing how many people (and whom) to treat.

(i) Spillovers among farmers in the same village (e.g., Cai et al., 2015);

(ii) Treatment can be costly: treating each individual is sub-optimal;

(iii) And... network data can be difficult or infeasible to collect.

Q1 Does what the policy maker is currently doing maximize benefits net of costs/can we conduct **inference** on policy-optimality?

Q2 Measure that indicates the direction for an improvement/can we **estimate** the best policy?

Ex2 Cash transfers, health and welfare programs, etc.

This Paper: Welfare Maximization

Suppose we want to allocate treatments to maximize the average outcome in the population (welfare) taking **spillovers** into account.

This Paper: Welfare Maximization

Suppose we want to allocate treatments to maximize the average outcome in the population (welfare) taking **spillovers** into account.

Example: No costs, basic policy, simple model where $\mathcal{N}_i$ are neighbors of i.

$$\underbrace{Y_i}_{\text{outcome}} = \underbrace{D_i \times \gamma_1}_{\text{direct effect}} + \qquad .$$

This paper: welfare maximization

Suppose we want to allocate treatments to maximize the average outcome in the population (welfare) taking **spillovers** into account.

Example: No costs, basic policy, simple model where $\mathcal{N}_i$ are neighbors of i.

$$\underbrace{Y_i}_{\text{outcome}} = \underbrace{D_i \times \gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} \quad .$$

Treat *everybody* if $\gamma_1 + \gamma_2 > 0$;

This paper: welfare maximization

Suppose we want to allocate treatments to maximize the average outcome in the population (welfare) taking **spillovers** into account.

Example: No costs, basic policy, simple model where $\mathcal{N}_i$ are neighbors of i.

$$\underbrace{Y_i}_{\text{outcome}} = \underbrace{D_i \times \gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} - \underbrace{\Big(\frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3}_{\text{quadratic spillovers}} + \nu_i.$$

Treat *everybody* if $\gamma_1 + \gamma_2 > 0$; if $\gamma_3 \neq 0$, unclear.

This paper: welfare maximization

Suppose we want to allocate treatments to maximize the average outcome in the population (welfare) taking **spillovers** into account.

Example: No costs, basic policy, simple model where $\mathcal{N}_i$ are neighbors of i.

$$\underbrace{Y_i}_{\text{outcome}} = \underbrace{D_i \times \gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} - \underbrace{\Big(\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3}_{\text{quadratic spillovers}} + \nu_i.$$

Treat *everybody* if $\gamma_1 + \gamma_2 > 0$; if $\gamma_3 \neq 0$, unclear.

Policy with no information: $P(D_i = 1) = \beta$ (Example)

$$W(\beta) = \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}_\beta[Y_i], \quad \beta^* \in \arg\max_\beta W(\beta).$$

Some remarks to keep in mind

In the paper, I consider (illustration)

(A) More complex policies:
- ⇒ Targeted treatments on observables $P(D_i = 1|X_i = x) = \pi(x; \beta)$;
- ⇒ Constraints on the policy space.
- Ex. Treating differently people in remote areas, younger/older, etc.

(B) Unknown model for spillovers.

Throughout most of this talk: policy is choosing how many people to treat.

Basic Intuition Behind our Design

- Network is difficult to collect:
  - ⇒ Consider few unobserved networks/clusters (e.g., regions).

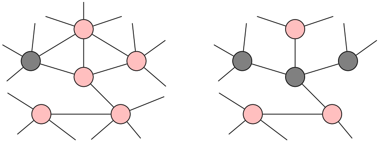

Basic Intuition Behind our Design

- Network is difficult to collect:
  - ⇒ Consider few unobserved networks/clusters (e.g., regions).

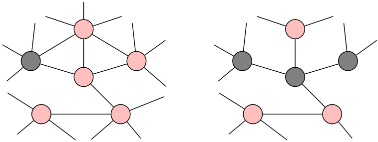

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

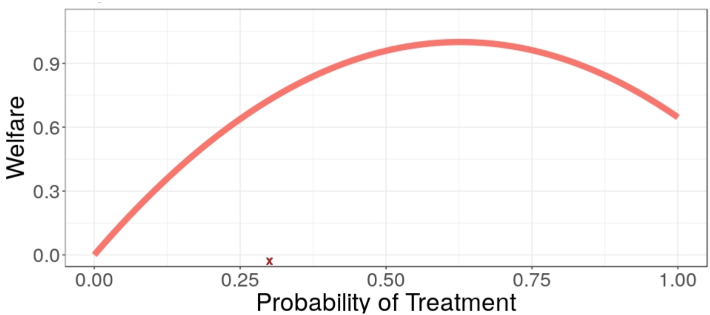

Basic Intuition Behind our Design

- Network is difficult to collect:
  - ⇒ Consider few unobserved networks/clusters (e.g., regions).

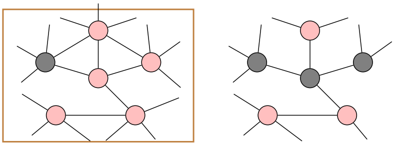

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

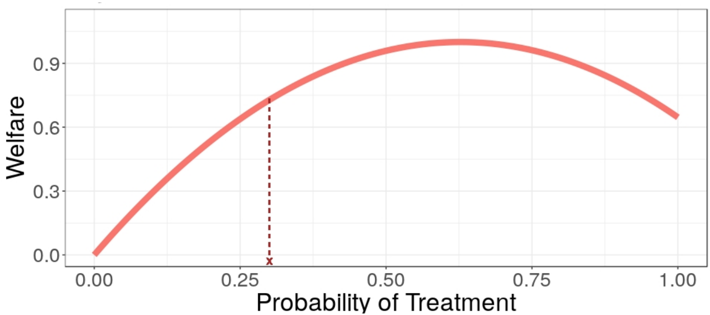

Basic Intuition Behind our Design

- Network is difficult to collect:
  - ⇒ Consider few unobserved networks/clusters (e.g., regions).

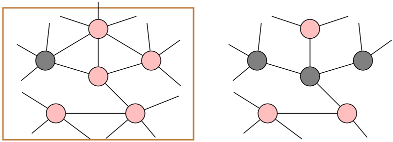

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

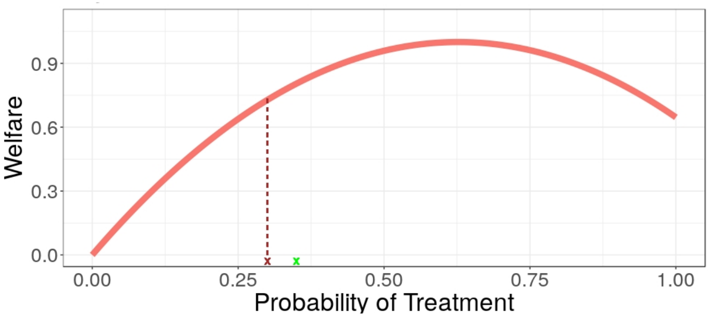

Basic Intuition Behind our Design

- Network is difficult to collect:
  - $\Rightarrow$ Consider few unobserved networks/clusters (e.g., regions).

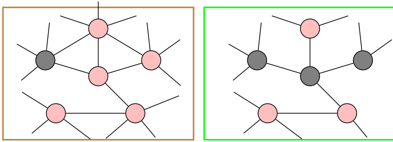

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

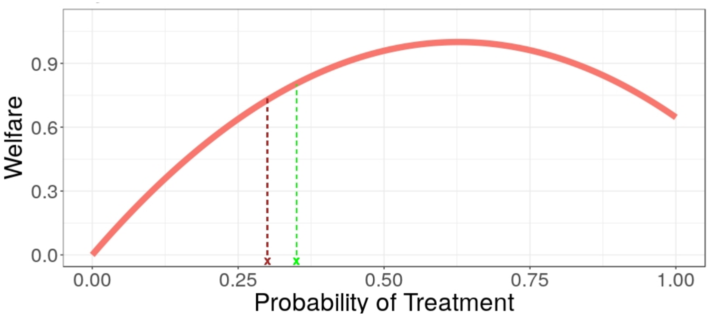

Basic Intuition Behind our Design

- Network is difficult to collect:
  - ⇒ Consider few unobserved networks/clusters (e.g., regions).

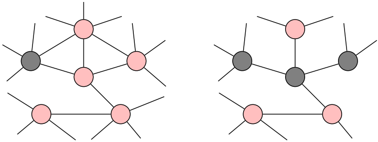

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

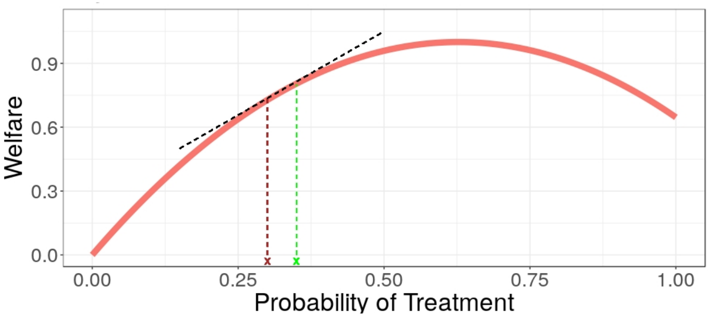

Basic Intuition Behind our Design

- Network is difficult to collect:
  - $\Rightarrow$ Consider few unobserved networks/clusters (e.g., regions).

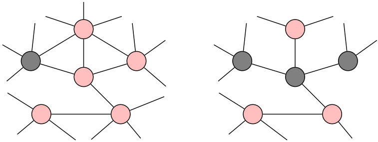

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

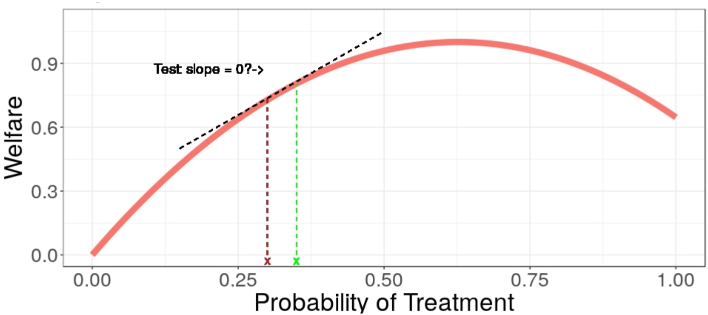

Basic Intuition Behind our Design

- Network is difficult to collect:
  - $\Rightarrow$ Consider few unobserved networks/clusters (e.g., regions).

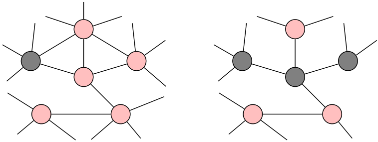

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

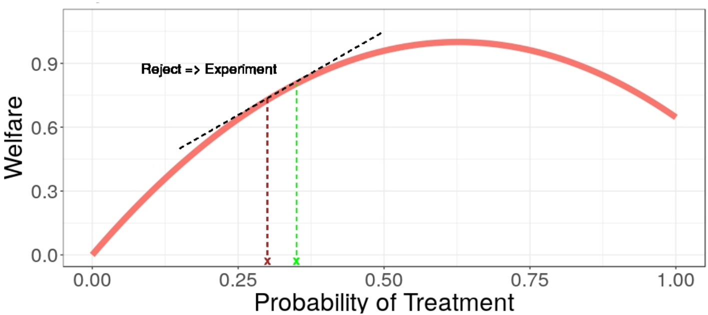

Basic Intuition Behind our Design

- Network is difficult to collect:
  - $\Rightarrow$ Consider few unobserved networks/clusters (e.g., regions).

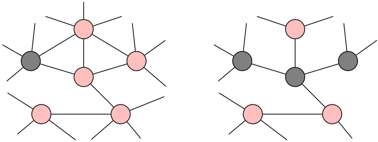

Illustration: $D_i \sim \mathrm{Bern}(\beta)$

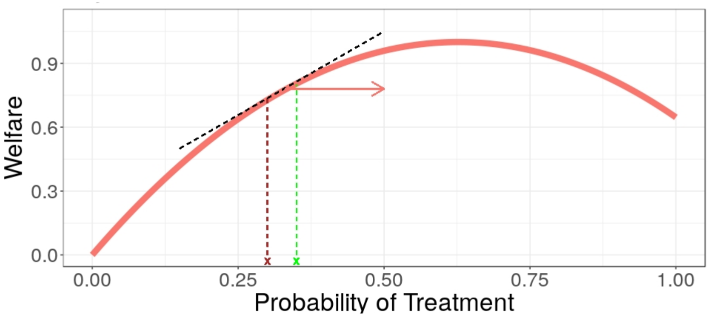

Overview and Contributions

- 40% of experimental papers in 2020 top-five mention spillovers. more
- Standard experiments with spillover effects are geared towards ATEs.

  $\Rightarrow$ Not sufficient for welfare maximization.

Overview and Contributions

- 40% of experimental papers in 2020 top-five mention spillovers. more
- Standard experiments with spillover effects are geared towards ATEs.
  $\Rightarrow$ Not sufficient for welfare maximization.

What we do:

1. Design a short (single-wave) experiment that
   - Identifies policy relevant estimands/direction for welfare improvement;
   - Also allows for inference on treatment and spillover effects.
2. Design a multi-wave/sequential experiment that
   - Efficiently recovers the best policy;
   - Also improves *participants'* welfare.

Overview and Contributions

- 40% of experimental papers in 2020 top-five mention spillovers. more
- Standard experiments with spillover effects are geared towards ATEs.
  $\Rightarrow$ Not sufficient for welfare maximization.

What we do:

1. Design a short (single-wave) experiment that
   - Identifies policy relevant estimands/direction for welfare improvement;
   - Also allows for inference on treatment and spillover effects.
2. Design a multi-wave/sequential experiment that
   - Efficiently recovers the best policy;
   - Also improves *participants'* welfare.

   $\Rightarrow$ Controls in-sample and out-of-sample regret at a fast rate in T and K:

$$\underbrace{W(\beta^*) - W(\hat{\beta})}_{\text{out-of-sample regret}} \leq \frac{C}{K}, \quad \text{In-sample regret} \leq \frac{C \log(K)}{K}.$$

Overview and Contributions

- 40% of experimental papers in 2020 top-five mention spillovers. more
- Standard experiments with spillover effects are geared towards ATEs.
  $\Rightarrow$ Not sufficient for welfare maximization.

What we do:

1. Design a short (single-wave) experiment that
   - Identifies policy relevant estimands/direction for welfare improvement;
   - Also allows for inference on treatment and spillover effects.
2. Design a multi-wave/sequential experiment that
   - Efficiently recovers the best policy;
   - Also improves *participants'* welfare.

   $\Rightarrow$ Controls in-sample and out-of-sample regret at a fast rate in T and K:

$$\underbrace{W(\beta^*) - W(\hat{\beta})}_{\text{out-of-sample regret}} \leq \frac{C}{K}, \quad \text{In-sample regret} \leq \frac{C \log(K)}{K}.$$

$\Rightarrow$ First framework for experimental design to maximize welfare with unobserved/unknown spillovers (see Viviano, 2019 for observed spill).

Related Literature

1. **Experimental design** for
   - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
   - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

Related Literature

1. **Experimental design** for
   - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
   - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

$\Rightarrow$ None studies policy design with network spillovers/partial interference.

Related Literature

1. **Experimental design** for
   - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
   - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

   $\Rightarrow$ None studies policy design with network spillovers/partial interference.

2. **Inference on networks/clusters**
   - Savje et al. (2020); Aronow and Samii (2017); Hudgens and Halloran (2008); Ibragimov and Mueller (2010, 2016); Goldsmith-Pinkham and Imbens (2013); ...

Related Literature

1. **Experimental design** for
   - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
   - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

   $\Rightarrow$ None studies policy design with network spillovers/partial interference.
2. **Inference on networks/clusters**
   - Savje et al. (2020); Aronow and Samii (2017); Hudgens and Halloran (2008); Ibragimov and Mueller (2010, 2016); Goldsmith-Pinkham and Imbens (2013); ...

   $\Rightarrow$ None studies policy/experimental design.

Related Literature

1. **Experimental design** for
   - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
   - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

   $\Rightarrow$ None studies policy design with network spillovers/partial interference.
2. **Inference on networks/clusters**
   - Savje et al. (2020); Aronow and Samii (2017); Hudgens and Halloran (2008); Ibragimov and Mueller (2010, 2016); Goldsmith-Pinkham and Imbens (2013); ...

   $\Rightarrow$ None studies policy/experimental design.
3. **Other literature**:
   - Treatment choice;
   - Seeding/centrality;
   - Peer groups' allocations;
   - Optimal taxation.

Related Literature

1. **Experimental design** for
    - Inference (with interference): Baird et al., 2018; Eckles et al., 2017; Basse and Airoldi, 2018; Johari et al. (2020); Viviano (2020)... ;
    - Optimization: *Pricing in two-sided markets* (Wager and Xu, 2021); *Bandits with iid* (Bubeck, 2012; Agarwal et al., 2010); *Adaptive randomizations without inference* (Kasy and Sattmann, 2020); ... .

   $\Rightarrow$ None studies policy design with network spillovers/partial interference.

2. **Inference on networks/clusters**
    - Savje et al. (2020); Aronow and Samii (2017); Hudgens and Halloran (2008); Ibragimov and Mueller (2010, 2016); Goldsmith-Pinkham and Imbens (2013); ...

   $\Rightarrow$ None studies policy/experimental design.

3. **Other literature**:
    - Treatment choice;
    - Seeding/centrality;
    - Peer groups' allocations;
    - Optimal taxation.

   $\Rightarrow$ None studies experimental design/adaptive assignments.

Content

1. Single wave experiment

2. Adaptive experiment

3. Additional results and conclusions

Basic setup

Asm Individuals are organized in relatively few large clusters with N units:

- Individuals interact with at most γ_N many other individuals;
- Each cluster may have different networks;
- Outcomes depend arbitrarily on neighbors' assignments.

Basic setup

Asm Individuals are organized in relatively few large clusters with N units:

- Individuals interact with at most γ_N many other individuals;
- Each cluster may have different networks;
- Outcomes depend arbitrarily on neighbors' assignments.

Exp: $D_i^k | X_i^k = x \sim \pi(x; ?)$

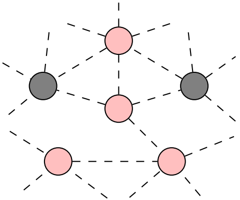

$\left(Y_{i,1}^k, Y_{i,0}^k, D_i^k, X_i^k\right)_{i=1}^n, n \leq N, k \leq K$

Target: $D_i | X_i = x \sim \pi(x; \hat{\beta})$

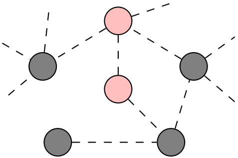

X

Sampling

For the moment, consider:

- $t = 1$ (one-wave experiment), one cluster
- No covariates: $D_i | \beta \sim \mathrm{Bern}(\beta)$.

Sampling

For the moment, consider:

- $t = 1$ (one-wave experiment), one cluster
- No covariates: $D_i|\beta \sim \mathrm{Bern}(\beta)$.

Network: (with some simplification, (general))

$$A^k_{i,j} = \underbrace{f(U^k_i, U^k_j)}_{\text{graphon}} \underbrace{1\{i \sim j\}}_{\text{latent space}}, \quad U^k_i \sim_{i.i.d.} F_U.$$

Ex Farmers a given *region* only interact in the same and nearby villages.

Def $\gamma_N = \sum_i 1\{i \sim j\}$.

Sampling

For the moment, consider:

- $t = 1$ (one-wave experiment), one cluster
- No covariates: $D_i|\beta \sim \mathrm{Bern}(\beta)$.

Network: (with some simplification, (general))

$$A_{i,j}^k = \underbrace{f(U_i^k, U_j^k)}_{\text{graphon}} \underbrace{1\{i \sim j\}}_{\text{latent space}}, \quad U_i^k \sim_{i.i.d.} F_U.$$

Ex Farmers a given *region* only interact in the same and nearby villages.

Def $\gamma_N = \sum_i 1\{i \sim j\}$.

Basic outcome model: In cluster k (dropping the superscript k), for unknown $r()$ het

$$Y_{i,1} = r\Big(D_i, D_{\mathcal{N}_i}, U_i, U_{\mathcal{N}_i}, |\mathcal{N}_i|, \nu_{i,1}\Big) + \tau_k, \quad Y_{i,0} : \text{baseline}.$$

Welfare and Marginal Effects

Lem We write (in cluster k), for $\nu_{i,1}$ i.i.d. exogenous unobservables

Welfare and Marginal Effects

Lem We write (in cluster k), for $\nu_{i,1}$ i.i.d. exogenous unobservables

$$\mathbb{E}_\beta[Y_{i,1}|D_i] = m(D_i, \beta) + \tau_k, \quad \text{for} \quad D_i \sim \mathrm{Bern}(\beta).$$

Welfare and Marginal Effects

Lem We write (in cluster k), for $\nu_{i,1}$ i.i.d. exogenous unobservables

$$\mathbb{E}_\beta[Y_{i,1}|D_i] = m(D_i, \beta) + \tau_k, \quad \text{for} \quad D_i \sim \mathrm{Bern}(\beta).$$

Def Welfare, marginal effect, direct effect: example

$$W(\beta) = \mathbb{E}_\beta[Y_{i,1}] - c\beta, \quad \underbrace{V(\beta) = \frac{\partial W(\beta)}{\partial \beta}}_{\text{marginal effect}}, \quad \underbrace{\Delta(\beta) = m(1,\beta) - m(0,\beta)}_{\text{Direct effect}}.$$

Welfare and Marginal Effects

Lem We write (in cluster k), for $\nu_{i,1}$ i.i.d. exogenous unobservables

$$\mathbb{E}_\beta[Y_{i,1}|D_i] = m(D_i, \beta) + \tau_k, \quad \text{for} \quad D_i \sim \mathrm{Bern}(\beta).$$

Def Welfare, marginal effect, direct effect: example

$$W(\beta) = \mathbb{E}_\beta[Y_{i,1}] - c\beta, \quad \underbrace{V(\beta) = \frac{\partial W(\beta)}{\partial \beta}}_{\text{marginal effect}}, \quad \underbrace{\Delta(\beta) = m(1, \beta) - m(0, \beta)}_{\text{Direct effect}}.$$

Remarks:

- $V(\beta)$ also captures spillovers (function of $\Delta(\beta)$ and $\frac{\partial m(d,\beta)}{\partial \beta}$, more);
- $V(\beta)$ provides a *direction* for a welfare improvement;
- $\beta = \beta^* \Rightarrow V(\beta) = 0$.

Estimation of marginal effects: one-wave experiment

⇒ Goal: estimate $V(\beta)$.

Estimation of marginal effects: one-wave experiment

⇒ Goal: estimate $V(\beta)$.

1. Matching clusters:

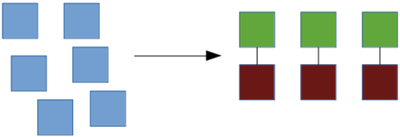

Estimation of marginal effects: one-wave experiment

⇒ Goal: estimate $V(\beta)$.

1. Matching clusters:

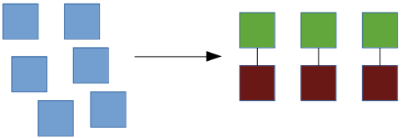

2. Small deviations:

$$D_i^k \sim \mathrm{Bern}(\beta_k), \quad \beta_k = \begin{cases} \beta - \eta_n \text{ if } k \text{ is odd;} \\ \beta + \eta_n \text{ otherwise} \end{cases}$$

Estimation of marginal effects: one-wave experiment

⇒ Goal: estimate $V(\beta)$.

1. Matching clusters:

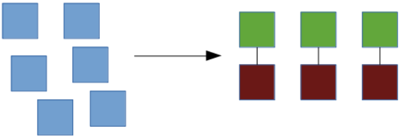

2. Small deviations:

$$D_i^k \sim \mathrm{Bern}(\beta_k), \quad \beta_k = \begin{cases} \beta - \eta_n & \text{if } k \text{ is odd;} \\ \beta + \eta_n & \text{otherwise} \end{cases}$$

3. Estimator: for pair $g = (k, k+1)$ illustration

$$\widehat{V}_g(\beta) = \frac{1}{2\eta_n}\left[\bar{Y}_1^{k+1} - \bar{Y}_0^{k+1}\right] - \frac{1}{2\eta_n}\left[\bar{Y}_1^{k} - \bar{Y}_0^{k}\right].$$

Some Guarantees

Under regularities more, with probability at least $1 - 1/n$

$$\left|\hat{V}_g - V(\beta)\right| = \tilde{\mathcal{O}}\Big(\underbrace{\sqrt{\frac{\gamma_N}{n \times \eta_n^2}}}_{var} + \underbrace{\eta_n}_{bias}\Big)$$

Some Guarantees

Under regularities more, with probability at least $1 - 1/n$

$$\left| \hat{V}_g - V(\beta) \right| = \tilde{\mathcal{O}}\Big(\underbrace{\sqrt{\frac{\gamma_N}{n \times \eta_n^2}}}_{var} + \underbrace{\eta_n}_{bias} \Big)$$

$\Rightarrow$ Bias-variance trade-off for tuning parameter (rule of thumb).

Some Guarantees

Under regularities more, with probability at least $1 - 1/n$

$$\left| \hat{V}_g - V(\beta) \right| = \tilde{\mathcal{O}}\Big(\underbrace{\sqrt{\frac{\gamma_N}{n \times \eta_n^2}}}_{var} + \underbrace{\eta_n}_{bias} \Big)$$

$\Rightarrow$ Bias-variance trade-off for tuning parameter (rule of thumb).

In the paper

$\Rightarrow$ Asymptotic normality of $\hat{V}_g$;

$\Rightarrow$ Test statistic to test $H_0 : V(\beta) = 0$ using clusters' pairs (details).

$\Rightarrow$ Finite cluster asymptotics ($K < \infty$) with pairing more, power sim.

Some Guarantees

Under regularities (more), with probability at least $1 - 1/n$

$$\left|\hat{V}_g - V(\beta)\right| = \tilde{\mathcal{O}}\Big(\underbrace{\sqrt{\frac{\gamma_N}{n \times \eta_n^2}}}_{var} + \underbrace{\eta_n}_{bias}\Big)$$

$\Rightarrow$ Bias-variance trade-off for tuning parameter (rule of thumb).

In the paper

- $\Rightarrow$ Asymptotic normality of $\hat{V}_g$;
- $\Rightarrow$ Test statistic to test $H_0 : V(\beta) = 0$ using clusters' pairs (details).
  - $\Rightarrow$ Finite cluster asymptotics ($K < \infty$) with pairing (more), (power sim).
- $\Rightarrow$ Guarantees also for treatment/spillover effects (more)

$$\mathbb{E}\Big[\frac{1}{K}\sum_{k=1}^{K}\hat{\Delta}_k\Big] = \Delta(\beta) + o(n^{-1/2}), \quad \frac{\partial m(0,\beta)}{\partial \beta}.$$

where $\hat{\Delta}^k$: Weighted difference treated/control in cluster k.

Practical implications

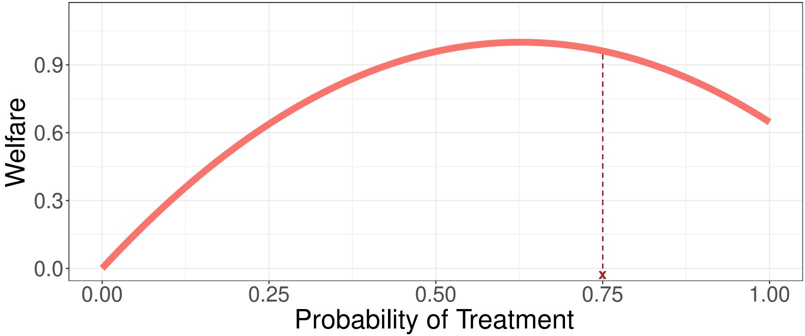

Practical implications

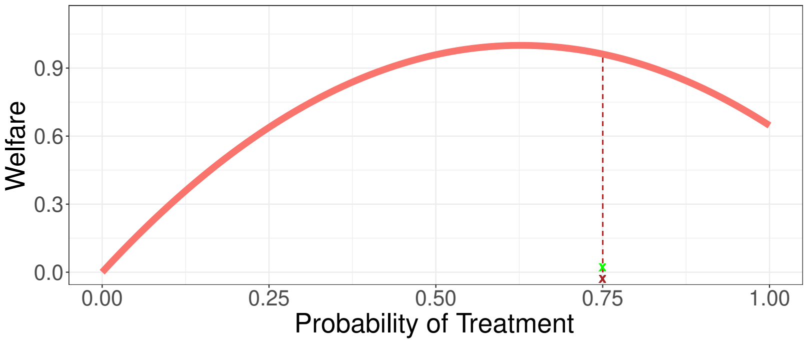

Practical implications

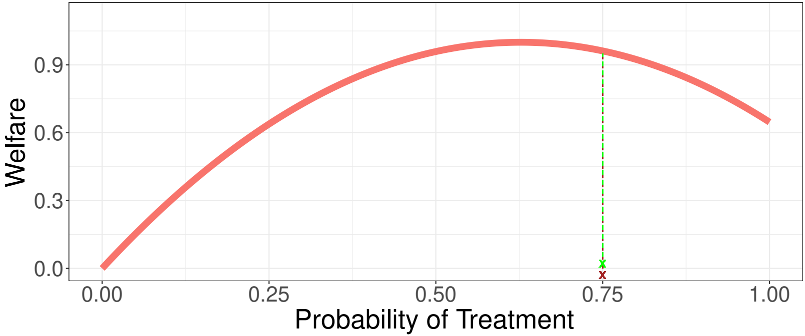

Practical implications

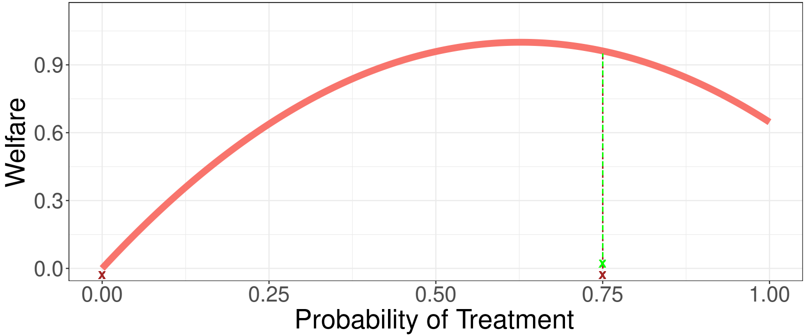

Practical implications

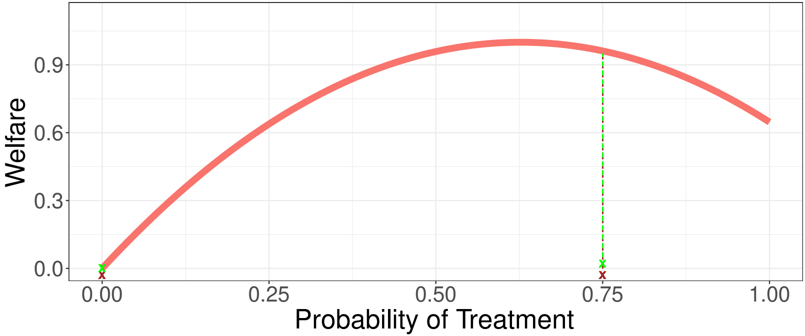

Practical implications

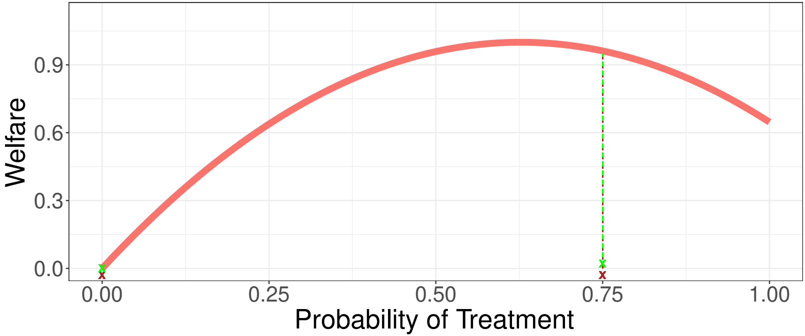

$\Rightarrow$ Estimate direct effect $\Delta(\beta)$ and $W(\beta) - W(0)$ by pooling observations

Practical implications

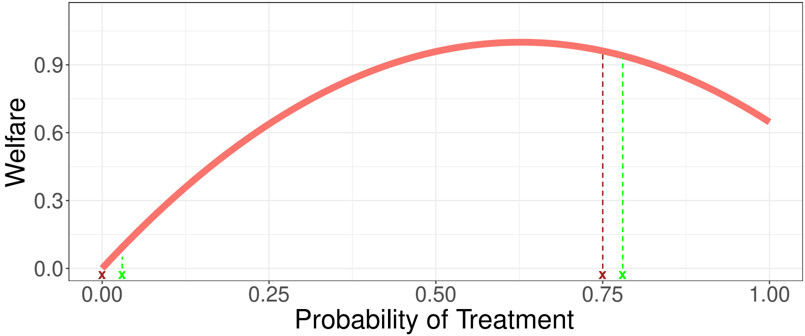

$\Rightarrow$ Estimate direct effect $\Delta(\beta)$ and $W(\beta) - W(0)$ by pooling observations

Practical implications

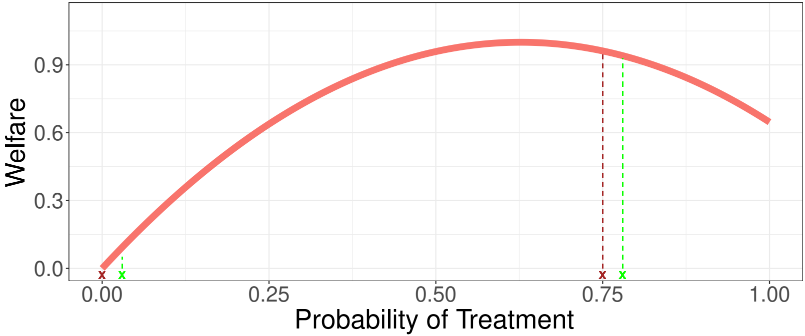

$\Rightarrow$ Estimate direct effect $\Delta(\beta)$ and $W(\beta) - W(0)$ by pooling observations up to a small bias (negligible for inference);

Practical implications

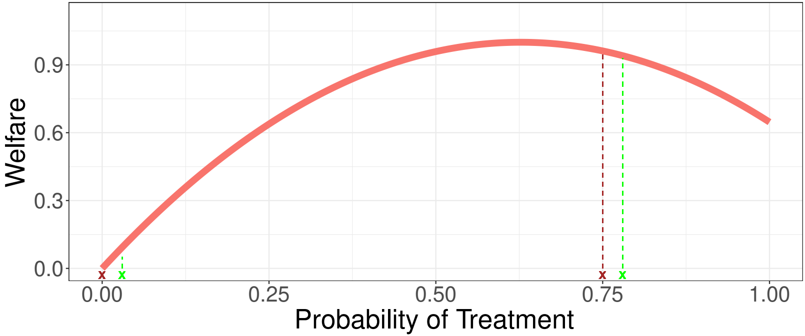

⇒ Estimate direct effect $\Delta(\beta)$ and $W(\beta) - W(0)$ by pooling observations up to a small bias (negligible for inference);

⇒ Estimate marginal effect $\frac{\partial W(\beta)}{\partial \beta}$;

Practical implications

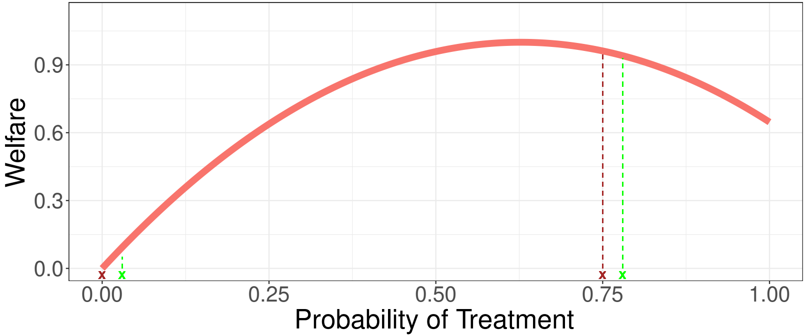

$\Rightarrow$ Estimate direct effect $\Delta(\beta)$ and $W(\beta) - W(0)$ by pooling observations up to a small bias (negligible for inference);

$\Rightarrow$ Estimate marginal effect $\frac{\partial W(\beta)}{\partial \beta}$;

$\Rightarrow$ Also estimate marginal spillover effects $\frac{\partial m(0,\beta)}{\partial \beta}$.

Content

1. Single wave experiment
2. Adaptive experiment
3. Additional results and conclusions

Adaptive experiment

"How do we estimate β^* with a sequential experiment?".

$\Rightarrow$ Treatments/outcomes collected sequentially $(D_{i,t}^k, Y_{i,t}^k)$ (more).

Adaptive experiment

"How do we estimate β^* with a sequential experiment?".

$\Rightarrow$ Treatments/outcomes collected sequentially $(D_{i,t}^k, Y_{i,t}^k)$ (more).

1. Initialization $(\beta_{1,0}, \beta_{2,0}, \cdots, \beta_{K,0}) = (\beta_0, \cdots, \beta_0)$.

Adaptive experiment

"How do we estimate β^* with a sequential experiment?".

$\Rightarrow$ Treatments/outcomes collected sequentially $(D_{i,t}^k, Y_{i,t}^k)$ (more).

1. Initialization $(\beta_{1,0}, \beta_{2,0}, \cdots, \beta_{K,0}) = (\beta_0, \cdots, \beta_0)$.
2. Pairing and circular cross-fitting:

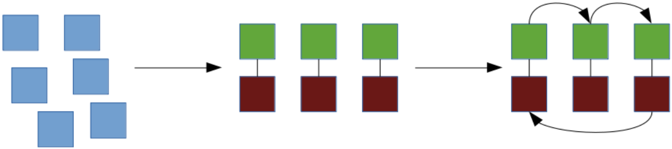

Adaptive experiment

"How do we estimate β^* with a sequential experiment?".

$\Rightarrow$ Treatments/outcomes collected sequentially $(D_{i,t}^k, Y_{i,t}^k)$ (more).

1. Initialization $(\beta_{1,0}, \beta_{2,0}, \cdots, \beta_{K,0}) = (\beta_0, \cdots, \beta_0)$.
2. Pairing and circular cross-fitting:

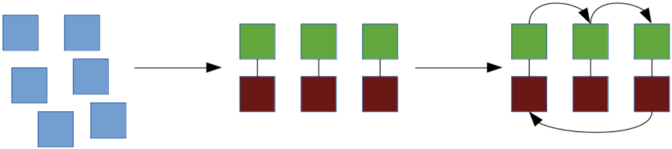

3. Policy update: $\beta_{k,t} = \beta_{k,t-1} + \alpha_{k,t} \hat{V}_{k+2,t-1}$ (learning rate)

Adaptive experiment

"How do we estimate β^* with a sequential experiment?".

$\Rightarrow$ Treatments/outcomes collected sequentially $(D^k_{i,t}, Y^k_{i,t})$ (more).

1. Initialization $(\beta_{1,0}, \beta_{2,0}, \cdots, \beta_{K,0}) = (\beta_0, \cdots, \beta_0)$.
2. Pairing and circular cross-fitting:

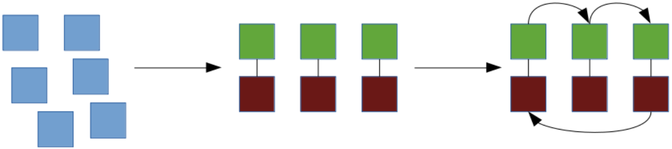

3. Policy update: $\beta_{k,t} = \beta_{k,t-1} + \alpha_{k,t} \hat{V}_{k+2,t-1}$ (learning rate)
4. Small deviations

$$D^k_{i,t} | \beta_{k,t} \sim \mathrm{Bern}(\beta_{k,t} \pm \eta_n)$$

Properties

Theorem: let $K \geq 2T$. Under regularity conditions (more), $p = \dim(\beta)$:

$$\underbrace{\max_k \frac{1}{T} \sum_{t=1}^{T} \Big[W(\beta^*) - W(\beta_{k,t}) \Big] \leq \frac{\bar{C} p^2 \log(T)}{T}}_{\text{in-sample}}, \quad \underbrace{W(\beta^*) - W(\hat{\beta}) \leq \frac{\bar{C} p^2}{T}}_{\text{out-of-sample}}$$

Properties

Theorem: let $K \geq 2T$. Under regularity conditions (more), $p = \dim(\beta)$:

$$\underbrace{\max_k \frac{1}{T} \sum_{t=1}^{T} \Big[W(\beta^*) - W(\beta_{k,t}) \Big] \leq \frac{\bar{C} p^2 \log(T)}{T}}_{\text{in-sample}}, \quad \underbrace{W(\beta^*) - W(\hat{\beta}) \leq \frac{\bar{C} p^2}{T}}_{\text{out-of-sample}}$$

- $\Rightarrow$ Rate $1/K$ for $K = 2T$;
- $\Rightarrow$ For $n \geq \bar{C} e^T \Rightarrow W(\beta^*) - W(\hat{\beta}) = \mathcal{O}(e^{-T/\kappa})$;
- $\Rightarrow$ Stronger than what you obtain with clusters as sampled units.
- $\Rightarrow$ 2-60% improvement over grid-search methods in simulations calibrated to information diffusion and cash-transfers. (more)

Properties

Theorem: let $K \geq 2T$. Under regularity conditions (more), $p = \dim(\beta)$:

$$\underbrace{\max_k \frac{1}{T}\sum_{t=1}^{T}\Big[W(\beta^*) - W(\beta_{k,t})\Big] \leq \frac{\bar{C}p^2\log(T)}{T}}_{\text{in-sample}}, \quad \underbrace{W(\beta^*) - W(\hat{\beta}) \leq \frac{\bar{C}p^2}{T}}_{\text{out-of-sample}}$$

$\Rightarrow$ Rate $1/K$ for $K = 2T$;

$\Rightarrow$ For $n \geq \bar{C}e^T \Rightarrow W(\beta^*) - W(\hat{\beta}) = \mathcal{O}(e^{-T/\kappa})$;

$\Rightarrow$ Stronger than what you obtain with clusters as sampled units.

$\Rightarrow$ 2-60% improvement over grid-search methods in simulations calibrated to information diffusion and cash-transfers. (more)

Why circular fitting?

$\Rightarrow$ I illustrate that with repeated sampling: $\mathbb{E}_{\beta_{k,t}}[Y_{i,t}^k] \neq \mathbb{E}_{\beta_{k,t}}[Y_{i,t}^k | \beta_{k,t}]$.

$\Rightarrow$ Circular fitting avoids bias and maximizes the number of clusters, and existing cross-fitting techniques would fail here for $T > 2$. (more)

Why circular fitting? Illustration

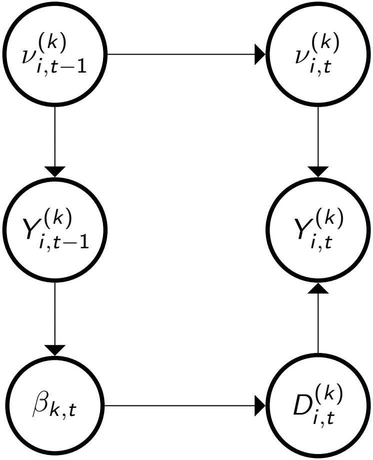

Experiment with repeated sampling

Why circular fitting? Illustration

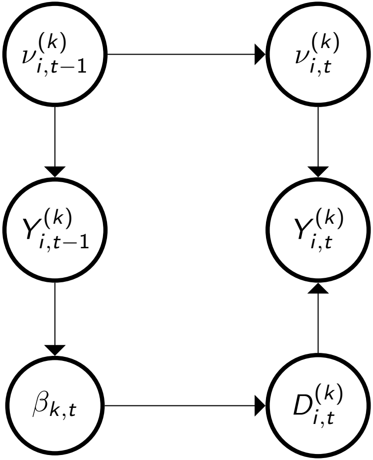

Experiment with repeated sampling

! $\mathbb{E}_{\beta_{k,t}}[Y_{i,t}^k] \neq \mathbb{E}_{\beta_{k,t}}[Y_{i,t}^k|\beta_{k,t}]$;

Why circular fitting? Illustration

! $\mathbb{E}_{\beta_{k,t}}[Y_{i,t}^{k}] \neq \mathbb{E}_{\beta_{k,t}}[Y_{i,t}^{k}|\beta_{k,t}]$;

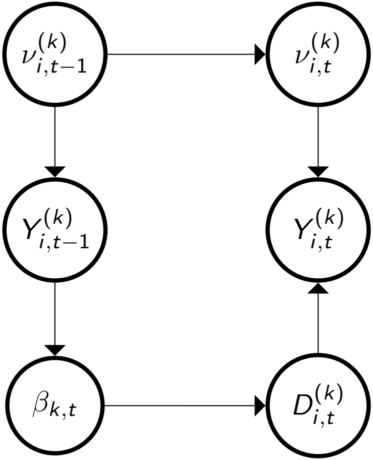

Experiment with repeated sampling

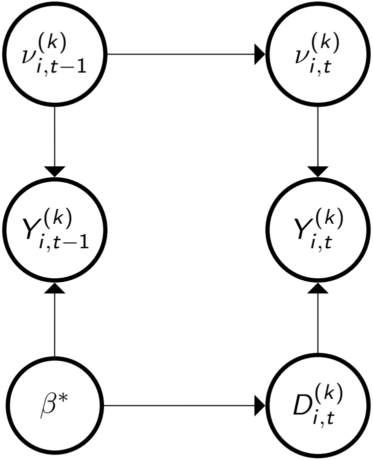

Policy on a new population

Content

1 Single wave experiment

2 Adaptive experiment

3 Additional results and conclusions

Compare to the policy that *does* observe the network

Difference betw largest welfare with observed and unobserved network?

Compare to the policy that *does* observe the network

Difference betw largest welfare with observed and unobserved network?

Asm1 Costs as opportunity costs without spillovers: $c = m(1, \beta) - m(0, \beta)$;

Asm2 Individuals depend on the share of treated friends (can be relaxed);

Asm3 Network is sufficiently dense.

Compare to the policy that *does* observe the network

Difference betw largest welfare with observed and unobserved network?

Asm1 Costs as opportunity costs without spillovers: $c = m(1, \beta) - m(0, \beta)$;

Asm2 Individuals depend on the share of treated friends (can be relaxed);

Asm3 Network is sufficiently dense.

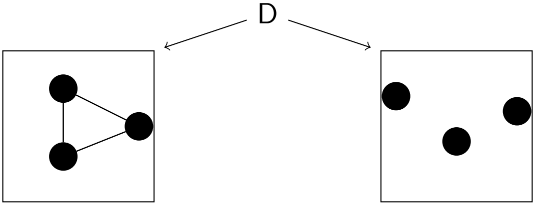

Compare to the policy that *does* observe the network

Difference betw largest welfare with observed and unobserved network?

Asm1 Costs as opportunity costs without spillovers: $c = m(1, \beta) - m(0, \beta)$;

Asm2 Individuals depend on the share of treated friends (can be relaxed);

Asm3 Network is sufficiently dense.

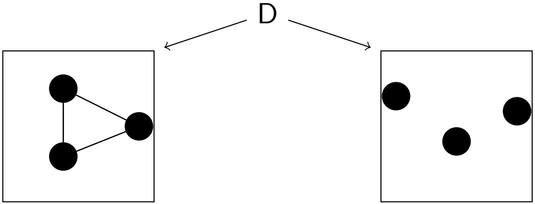

Thm The welfare with the best policy without observing the network converges to the largest welfare as we observe the network. more

Some extensions

In the paper:

- Matching with heterogeneity of covariates' distribution more;
- Inference/estimation with observed heterogeneity betw clusters more.
- Dynamic treatments/path of policies more;
- Treatments can be assigned only once;
- Strict quasi-concavity;
- Simple global interference mechanisms within cluster more;
- ...

Some extensions

In the paper:

- Matching with heterogeneity of covariates' distribution more;
- Inference/estimation with observed heterogeneity betw clusters more.
- Dynamic treatments/path of policies more;
- Treatments can be assigned only once;
- Strict quasi-concavity;
- Simple global interference mechanisms within cluster more;
- ...

In progress

- Application in collaboration with PxD/Chicago Lab (500k farmers).

Interesting future directions and related works

- Value of collecting network information in generic settings?
- unbounded degree with decaying dependence? (e.g., Theorem 3.1, more)
- network which also depends on the treatments? more

Interesting future directions and related works

- Value of collecting network information in generic settings?
- unbounded degree with decaying dependence? (e.g., Theorem 3.1, more)
- network which also depends on the treatments? more

Some related works

- Policy Targeting under Network Interference (Viviano, 2019):
  $\Rightarrow$ Policy choice using data from an existing experiment.
- Experimental Design under Network Interference (Viviano, 2020):
  $\Rightarrow$ Statistical framework for two-wave experiments with networks: select participants, and assign treatments for variance reduction.
- Fair Policy Targeting (Viviano and Bradic, 2020):
  $\Rightarrow$ Design fair and efficient treatment rules.

Conclusions

- I have introduced a statistical framework for **estimation** and **inference** for welfare-maximizing policies;
  - The framework allows for unobserved (and partial) interference;
  - The experiment consists of a matched-pair local and two-stage design.
- I have discussed an adaptive experiment for policy choice;
- I provide asymptotic properties and regret bounds of the procedure;
- I illustrate the method in a calibrated simulation.

Conclusions

- I have introduced a statistical framework for **estimation** and **inference** for welfare-maximizing policies;
    - The framework allows for unobserved (and partial) interference;
    - The experiment consists of a matched-pair local and two-stage design.
- I have discussed an adaptive experiment for policy choice;
- I provide asymptotic properties and regret bounds of the procedure;
- I illustrate the method in a calibrated simulation.

Questions? Thanks!

More at dviviano.github.io

Content

4 Appendix (not for presentation)

One-wave experiment

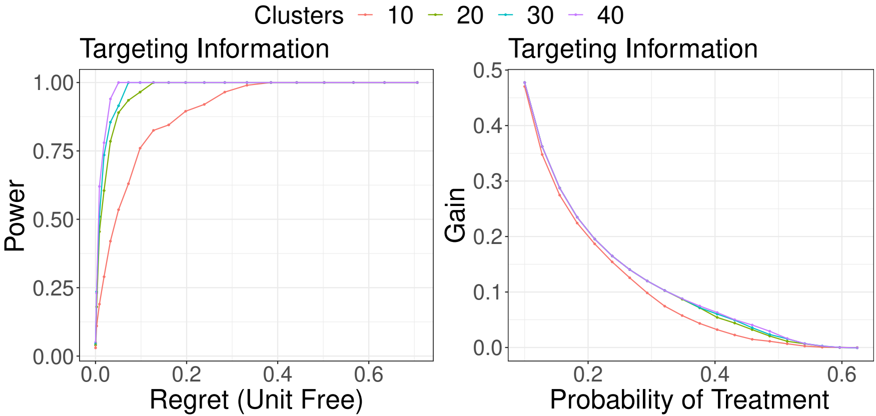

(back)

Why *circular*?

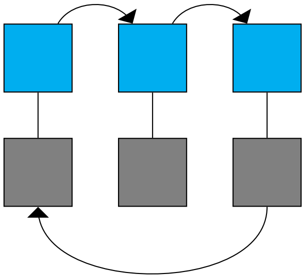

(i) Circular cross fitting guarantees exogeneity of the parameters over each iteration;

(ii) Existing sample-splitting/cross-fitting procedure would fail for $T > 2$.

Why *circular*?

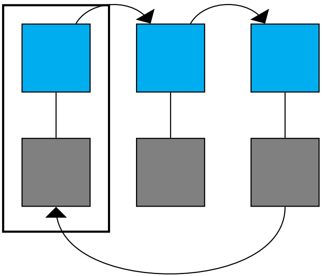

(i) Circular cross fitting guarantees exogeneity of the parameters over each iteration;

(ii) Existing sample-splitting/cross-fitting procedure would fail for $T > 2$.

Why *circular*?

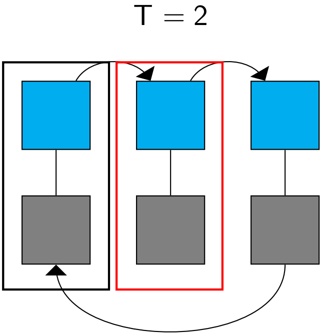

(i) Circular cross fitting guarantees exogeneity of the parameters over each iteration;

(ii) Existing sample-splitting/cross-fitting procedure would fail for $T > 2$.

Why *circular*?

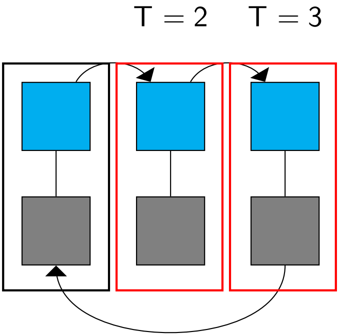

(i) Circular cross fitting guarantees exogeneity of the parameters over each iteration;

(ii) Existing sample-splitting/cross-fitting procedure would fail for $T > 2$.

(back)

Identification and estimation: DID

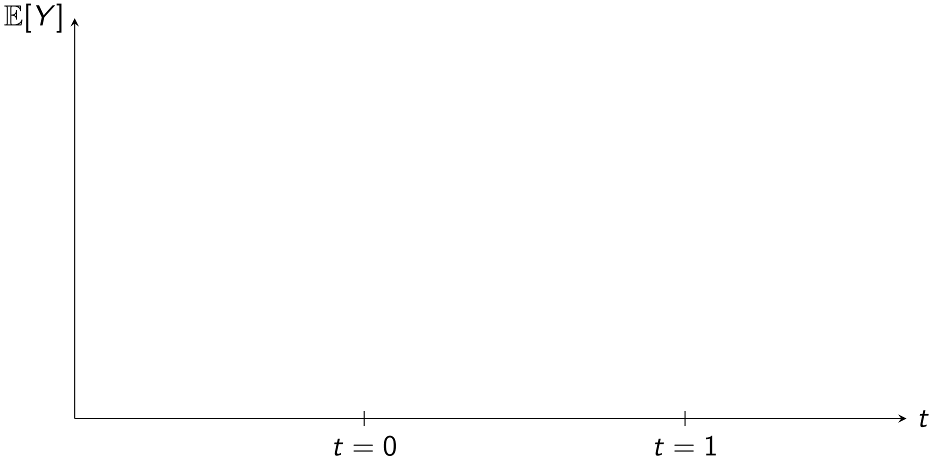

Identification and estimation: DID

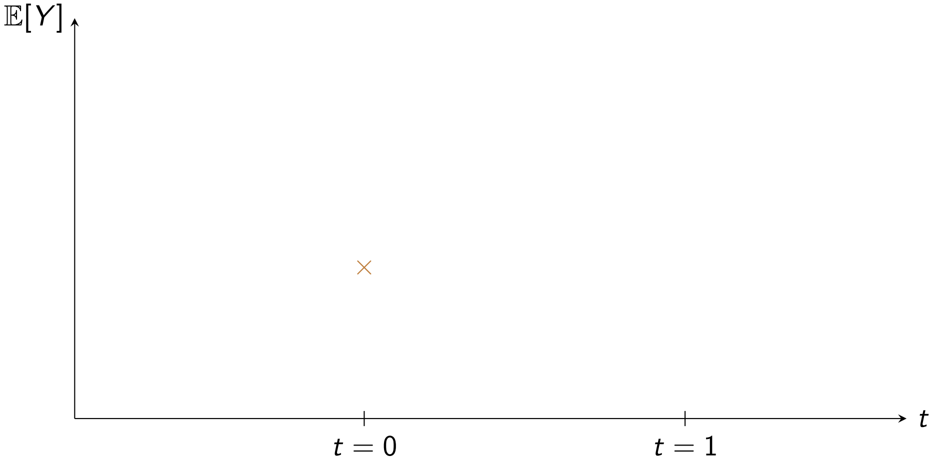

Identification and estimation: DID

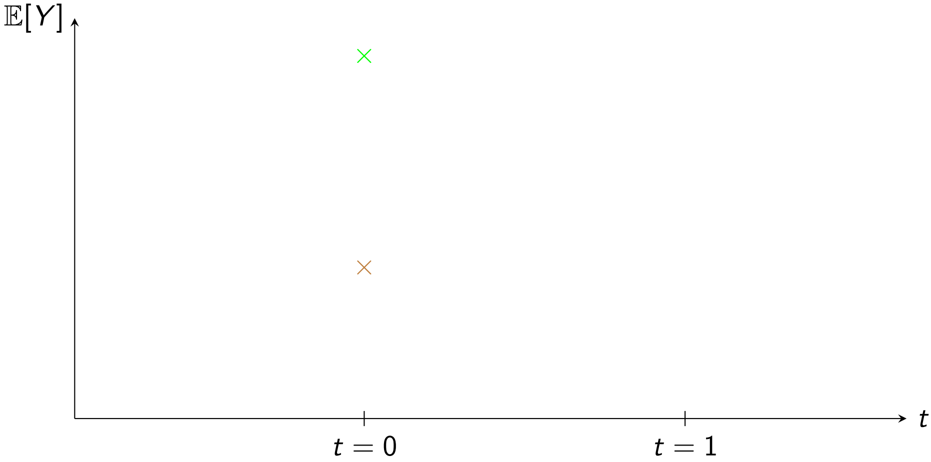

Identification and estimation: DID

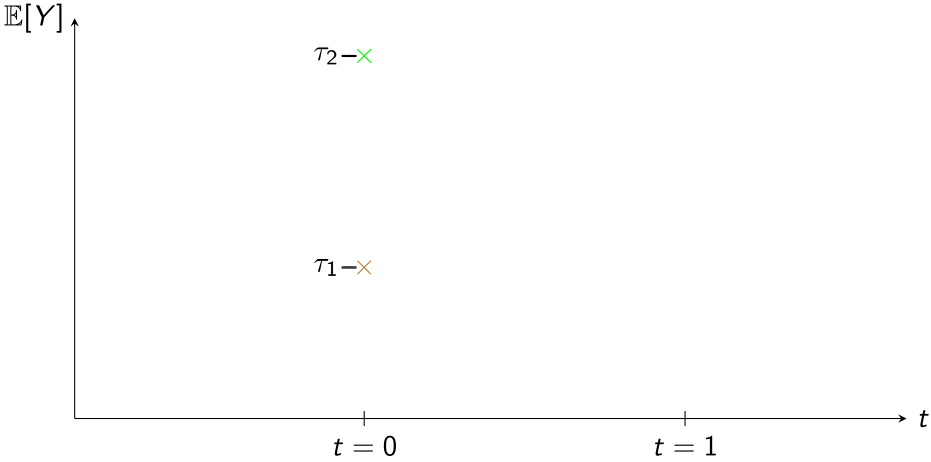

Identification and estimation: DID

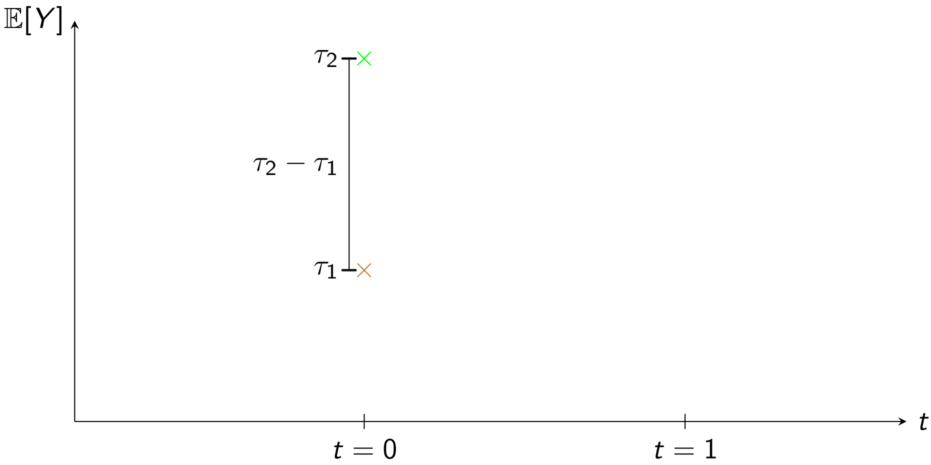

Identification and estimation: DID

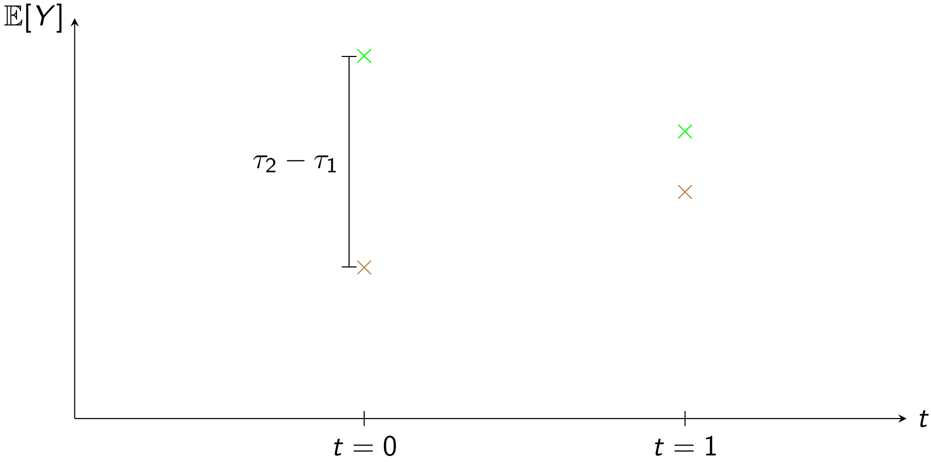

Identification and estimation: DID

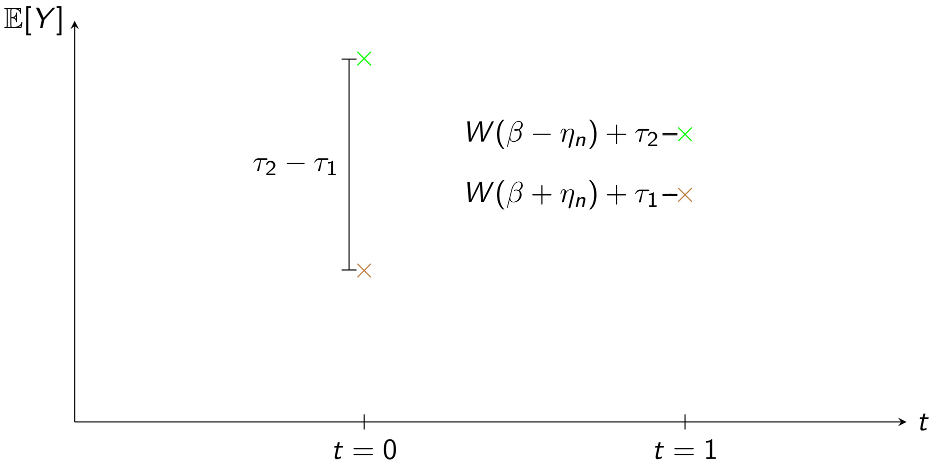

Identification and estimation: DID

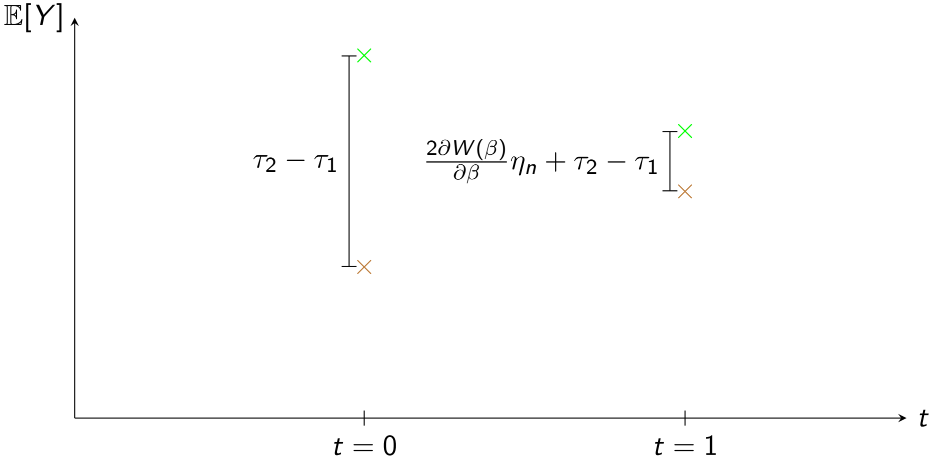

(back)

Adaptive Experiment: Comparisons

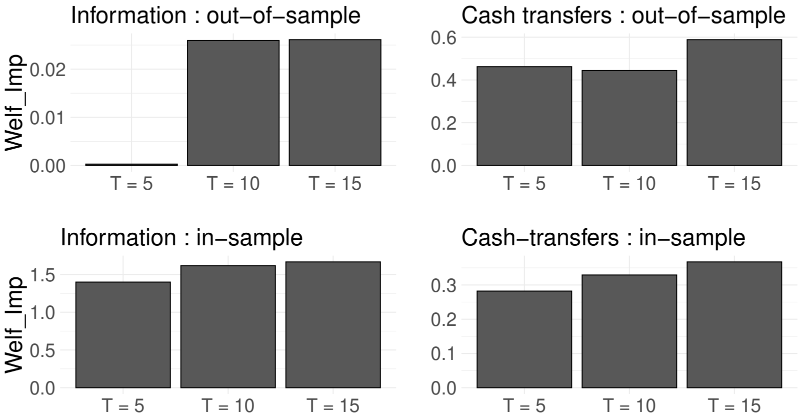

(back)

One simple example

$$Y_{i,1} = D_i \times \underbrace{\gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} - \underbrace{\Big(\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3}_{\text{quadratic spillovers}} + \nu_i.$$

One simple example

$$Y_{i,1} = D_i \times \underbrace{\gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} - \underbrace{\Big(\frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3}_{\text{quadratic spillovers}} + \nu_i.$$

Ex. Welfare is *approximately*: (details, back)

One simple example

$$Y_{i,1} = D_i \times \underbrace{\gamma_1}_{\text{direct effect}} + \underbrace{\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j \times \gamma_2}_{\text{linear spillovers}} - \underbrace{\Big(\frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3}_{\text{quadratic spillovers}} + \nu_i.$$

Ex. Welfare is *approximately*: (details, back)

$$W(\beta) \simeq \beta \underbrace{\gamma_1}_{\text{direct effect}} + \underbrace{\beta\gamma_2}_{\text{linear spillovers}} - \underbrace{\beta^2\gamma_3}_{\sim \text{ quadratic spillovers}} - \underbrace{c\beta}_{cost}$$

Note: this is just an example and the quadratic assumption/model is not required.

Example: details

Recall quadratic model

$$Y_{i,1} = D_{i,1}\gamma_1 + \frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j \gamma_2 - \Big(\frac{1}{|\mathcal{N}_i|}\sum_{j \in \mathcal{N}_i} D_j\Big)^2 \gamma_3 + \nu_{i,1}$$

Example: details

Recall quadratic model

$$Y_{i,1} = D_{i,1}\gamma_1 + \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\gamma_2 - \Big(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\Big)^2\gamma_3 + \nu_{i,1}$$

Taking expectations

$$\Rightarrow Y_{i,1} = \beta\gamma_1 + \beta\gamma_2 - Q(\beta)\gamma_3$$

Example: details

Recall quadratic model

$$Y_{i,1} = D_{i,1}\gamma_1 + \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\gamma_2 - \Big(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\Big)^2\gamma_3 + \nu_{i,1}$$

Taking expectations

$$\Rightarrow Y_{i,1} = \beta\gamma_1 + \beta\gamma_2 - Q(\beta)\gamma_3$$

We can write

$$Q(\beta) = \mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j\in\mathcal{N}_i} D_j^2\Big] + \mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j,h,j\neq h\in\mathcal{N}_i} D_j D_h\Big]$$

(back)

Example: details

Recall quadratic model

$$Y_{i,1} = D_{i,1}\gamma_1 + \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\gamma_2 - \Big(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j\Big)^2\gamma_3 + \nu_{i,1}$$

Taking expectations

$$\Rightarrow Y_{i,1} = \beta\gamma_1 + \beta\gamma_2 - Q(\beta)\gamma_3$$

We can write

$$Q(\beta) = \mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j\in\mathcal{N}_i} D_j^2\Big] + \mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j,h,j\neq h\in\mathcal{N}_i} D_j D_h\Big]$$
$$= \underbrace{\mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j\in\mathcal{N}_i} D_j\Big]}_{\beta\mathbb{E}[1/|\mathcal{N}_i|]\simeq 0} + \underbrace{\mathbb{E}\Big[\frac{1}{|\mathcal{N}_i|^2}\sum_{j,h,j\neq h\in\mathcal{N}_i} D_j D_h\Big]}_{(1-\mathbb{E}[1/|\mathcal{N}_i|])\beta^2\simeq\beta^2}$$

(back)

Marginal effects: a close look

Welfare: back

$$W(\beta) = \Big[\underbrace{\beta m(1, \beta) + (1 - \beta) m(0, \beta)}_{WELF} - \underbrace{c\beta}_{COST} \Big],$$

Marginal effects: a close look

Welfare: back

$$W(\beta) = \Big[\underbrace{\beta m(1,\beta) + (1-\beta)m(0,\beta)}_{WELF} - \underbrace{c\beta}_{COST}\Big],$$

Marginal effect:

$$V(\beta) = \Big[\underbrace{\beta\frac{\partial m(1,\beta)}{\partial\beta} + (1-\beta)\frac{\partial m(0,\beta)}{\partial\beta}}_{(S)} + \underbrace{(m(1,\beta) - m(0,\beta) - c)}_{(D)}\Big].$$

Marginal effects: a close look

Welfare: back

$$W(\beta) = \Big[\underbrace{\beta m(1,\beta) + (1-\beta) m(0,\beta)}_{WELF} - \underbrace{c\beta}_{COST}\Big],$$

Marginal effect:

$$V(\beta) = \Big[\underbrace{\beta \frac{\partial m(1,\beta)}{\partial \beta} + (1-\beta)\frac{\partial m(0,\beta)}{\partial \beta}}_{(S)} + \underbrace{\big(m(1,\beta) - m(0,\beta) - c\big)}_{(D)}\Big].$$

(S) Can be identified only with two clusters;
(D) Can be identified with a single cluster.

Marginal effects: a close look

Welfare: back

$$W(\beta) = \Big[\underbrace{\beta m(1,\beta) + (1-\beta)m(0,\beta)}_{WELF} - \underbrace{c\beta}_{COST}\Big],$$

Marginal effect:

$$V(\beta) = \Big[\underbrace{\beta\frac{\partial m(1,\beta)}{\partial\beta} + (1-\beta)\frac{\partial m(0,\beta)}{\partial\beta}}_{(S)} + \underbrace{\big(m(1,\beta) - m(0,\beta) - c\big)}_{(D)}\Big].$$

(S) Can be identified only with two clusters;

(D) Can be identified with a single cluster.

Ex Quadratic model: $W(\beta) = \beta\gamma_1 + \beta\gamma_2 - \beta^2\gamma_3 - c\beta$

Marginal effects: a close look

Welfare: back

$$W(\beta) = \Big[\underbrace{\beta m(1,\beta) + (1-\beta) m(0,\beta)}_{WELF} - \underbrace{c\beta}_{COST} \Big],$$

Marginal effect:

$$V(\beta) = \Big[\underbrace{\beta \frac{\partial m(1,\beta)}{\partial \beta} + (1-\beta)\frac{\partial m(0,\beta)}{\partial \beta}}_{(S)} + \underbrace{\big(m(1,\beta) - m(0,\beta) - c\big)}_{(D)} \Big].$$

(S) Can be identified only with two clusters;

(D) Can be identified with a single cluster.

Ex Quadratic model: $W(\beta) = \beta\gamma_1 + \beta\gamma_2 - \beta^2\gamma_3 - c\beta$

$$(D) = \gamma_1 - c, \quad \frac{\partial m(d,\beta)}{\partial \beta} \simeq \gamma_2 - 2\gamma_3\beta$$

Illustration for targeted treatments

Can only treat half of the population: (back)

⇒ Trade-off between treating people in remote/non remote areas.

Example

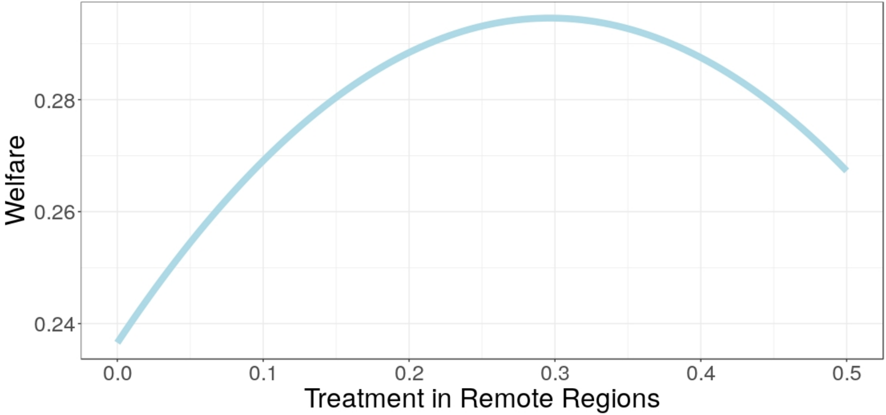

Meta-analysis

"Top-5 econ journals"

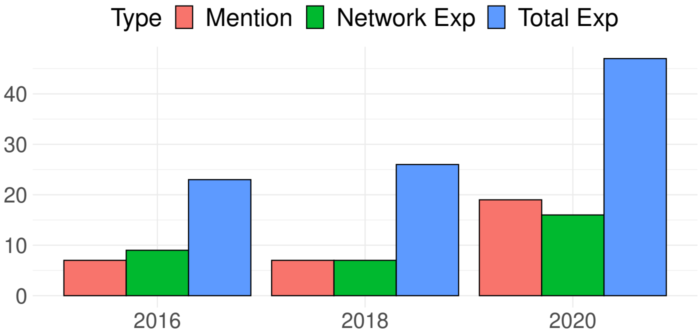

- 40% of experimental **papers** mention spillovers in their analysis;
- **Industry** survey: 22% of companies on online platforms conduct experiments (total: 8 million; Runge et al., 2021).

(back)

Sampling

- Sampling:

$$(X_i^k, U_i^k) \sim_{i.i.d.} F_{U|X} F_X, \quad A_{i,j}^k = f(X_i^k, X_j^k, U_i^k, U_j^k) 1\{i_k \sim j_k\}$$

Sampling

- Sampling:

$$(X_i^k, U_i^k) \sim_{i.i.d.} F_{U|X} F_X, \quad A_{i,j}^k = f(X_i^k, X_j^k, U_i^k, U_j^k) 1\{i_k \sim j_k\}$$

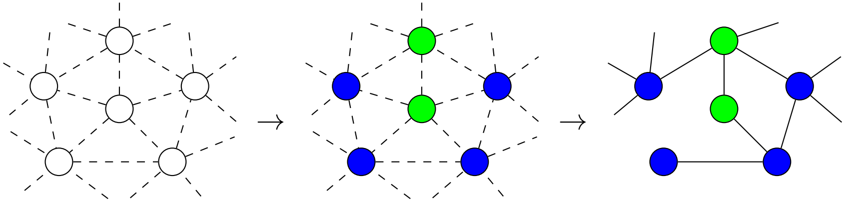

Sampling

- Sampling:

$$(X_i^k, U_i^k) \sim_{i.i.d.} F_{U|X} F_X, \quad A_{i,j}^k = f(X_i^k, X_j^k, U_i^k, U_j^k) 1\{i_k \sim j_k\}$$

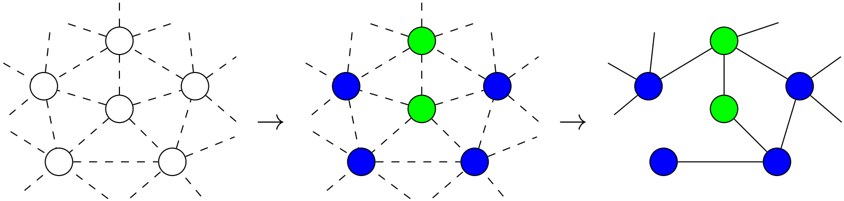

- $Y_{i,t}^k = r\Big(\mathcal{N}_i^k, D_{i,t}^k, D_{\mathcal{N}_i^k,t}, U_i^k, X_i^k, U_{\mathcal{N}_i^k}, X_{\mathcal{N}_i^k}, \nu_{i,t}\Big) + \alpha_t + \tau_k,$
- Ex-ante, individuals can only connect with some units: $|\mathcal{N}_i| \leq \gamma_N$.

(back)

Conditions under cluster heterogeneity

Consider the following model in cluster k

$$Y_{i,1} = r\Big(D_i, D_{\mathcal{N}_i}, U_i, U_{\mathcal{N}_i}, A_{i,\cdot}, \nu_{i,1}, \theta_k\Big) + \tau_k,$$

where

- θ_k is fixed (and observable and captures cluster heterogeneity;
- for each $\theta \in \Theta$, there are two clusters (k, k') such that $\theta_k = \theta_{k'}$.

Conditions under cluster heterogeneity

Consider the following model in cluster k

$$Y_{i,1} = r\Big(D_i, D_{\mathcal{N}_i}, U_i, U_{\mathcal{N}_i}, A_{i,\cdot}, \nu_{i,1}, \theta_k\Big) + \tau_k,$$

where

- θ_k is fixed (and observable and captures cluster heterogeneity;
- for each $\theta \in \Theta$, there are two clusters (k, k') such that $\theta_k = \theta_{k'}$.

Experiments:

- Single wave experiment: clusters with the same θ_k are matched with each other. Test remains valid for

$$H_0 : \beta(\theta) = \beta^*(\theta) \text{ for all } \theta.$$

- Multi-wave experiment: same matching as above + no time dependence: different policies for different groups.

(back, backextensions)

Regularity conditions

Regularities for consistency/high probability bounds

- $\varepsilon_{i,t}$ is sub-gaussian;
- $m(d, \beta)$ is uniformly bounded with bounded derivative.

Additional regularities for asymptotic normality

- $\mathrm{Var}\Big(\sqrt{n}(\bar{Y}_1^k - \bar{Y}_0^k)\Big) > 0$;
- $\gamma_N / N^{1/8} = o(1), n \propto N$;

Regularities for regret bounds (adaptive exp)

- $W(\beta)$ is strongly concave (can be relaxed with strict quasi-concavity).

(back)
(back2)

Rule of thumb

Choose

$$\eta_n = \begin{cases} n^{-1/3} \times \sqrt{\frac{2\hat{\sigma}}{\hat{c}}} \text{ if } n^{-1/3} \times \sqrt{\frac{2\hat{\sigma}}{\hat{c}}} < B \\ B \text{ otherwise,} \end{cases}$$

where

- $\hat{\sigma}$ is the individual outcomes' variance;
- $\hat{c}$ approximates the curvature of $W(\beta)$ (which can be obtained with three clusters and a quadratic model);
- $\hat{c}B$ is the largest bias that the researcher is willing to tolerate (e.g., $\hat{c}B = 0.05$).

(back)

Pivotal test statistics

Theorem: Under regularity conditions (more), for $\eta_n = o(n^{-1/4})$,

$$\frac{\widehat{V}_g(\beta) - V(\beta)}{\sqrt{\mathrm{Var}(\hat{V}_g(\beta))}} \to_d \mathcal{N}(0,1)$$

Corollary: let

$$\mathcal{T}_n = \frac{\sqrt{G}\,\bar{V}}{\sqrt{(G-1)^{-1}\sum_{g=1}^{G}(\widehat{V}_g - \bar{V})^2}}, \quad \bar{V} = \frac{1}{G}\sum_{g=1}^{G}\widehat{V}_g,$$

then (back)

$$\lim_{n\to\infty} P\Big(|\mathcal{T}_n| \le \mathrm{cv}_{K/2-1}(\alpha)|H_0\Big) \ge 1-\alpha.$$

Pairing permits finite cluster asymptotics

$\Rightarrow$ Without pairing concentration at rate $(K\eta_n^2)^{-1/2} \gg \eta_n$.

Pairing permits finite cluster asymptotics

$\Rightarrow$ Without pairing concentration at rate $(K\eta_n^2)^{-1/2} \gg \eta_n$.

Ex. Let $\beta \in [0, 1]$, no fixed effects.

Pairing permits finite cluster asymptotics

$\Rightarrow$ Without pairing concentration at rate $(K\eta_n^2)^{-1/2} \gg \eta_n$.

Ex. Let $\beta \in [0,1]$, no fixed effects.

$$W(\beta + \eta_n) - W(\beta - \eta_n) = \underbrace{W(\beta) - W(\beta)}_{(A)} + \frac{\partial W(\beta)}{\partial \beta}\eta_n - \frac{\partial W(\beta)}{\partial \beta}(-\eta_n) + O(\eta_n^2).$$

Pairing permits finite cluster asymptotics

$\Rightarrow$ Without pairing concentration at rate $(K\eta_n^2)^{-1/2} \gg \eta_n$.

Ex. Let $\beta \in [0,1]$, no fixed effects.

$$W(\beta + \eta_n) - W(\beta - \eta_n) = \underbrace{W(\beta) - W(\beta)}_{(A)} + \frac{\partial W(\beta)}{\partial \beta}\eta_n - \frac{\partial W(\beta)}{\partial \beta}(-\eta_n) + O(\eta_n^2).$$

$\Rightarrow$ (A) equals to zero because of pairing.

$\Rightarrow$ Otherwise, after averaging across all clusters, $(A) = O(1/\sqrt{K})$ with randomization.

(back)

Formula direct and spillovers

Direct effect

$$\hat{\Delta}^k = \frac{1}{n}\sum_{i=1}^{n}\Big[\frac{Y_i^k D_i^k}{\beta_k} - \frac{Y_i^k(1 - D_i^k)}{1 - \beta_k}\Big]$$

Formula direct and spillovers

Direct effect

$$\hat{\Delta}^k = \frac{1}{n}\sum_{i=1}^{n}\Big[\frac{Y_i^k D_i^k}{\beta_k} - \frac{Y_i^k(1-D_i^k)}{1-\beta_k}\Big]$$

Spillover effect (no fixed effects)

$$\hat{S}_0^k = \frac{1}{2n\eta_n}\sum_{i=1}^{n}\Big[\frac{Y_i^k(1-D_i^k)}{1-\beta_k} - \frac{Y_i^{k+1}(1-D_i^{k+1})}{1-\beta_{k+1}}\Big]$$

Formula direct and spillovers

Direct effect

$$\hat{\Delta}^k = \frac{1}{n}\sum_{i=1}^{n}\Big[\frac{Y_i^k D_i^k}{\beta_k} - \frac{Y_i^k(1-D_i^k)}{1-\beta_k}\Big]$$

Spillover effect (no fixed effects)

$$\hat{S}_0^k = \frac{1}{2n\eta_n}\sum_{i=1}^{n}\Big[\frac{Y_i^k(1-D_i^k)}{1-\beta_k} - \frac{Y_i^{k+1}(1-D_i^{k+1})}{1-\beta_{k+1}}\Big]$$

heterogeneous spillover effects (no fixed effects)

$$\hat{S}_0^k(x) = \frac{1}{2n\eta_n}\sum_{i=1}^{n}\Big[\underbrace{\frac{Y_i^k(1-D_i^k)1\{X_i^k = x\}}{(1-\beta_k)P(X_i^k = x)}}_{\text{cluster } k} - \underbrace{\frac{Y_i^{k+1}(1-D_i^{k+1})1\{X_i^{k+1} = x\}}{(1-\beta_{k+1})P(X_i^{k+1} = x)}}_{\text{cluster } k+1}\Big]$$

(back)

Adaptive experiment: details 1

Outcome (no covariates for simplicity)

$$Y_{i,t} = r\Big(D_{i,t}, D_{N_i,t}, U_i, U_{N_i}, A_{i,\cdot}, \nu_{i,t}\Big) + \alpha_t + \tau_k, \quad \nu_{i,t} \sim \mathcal{P}$$

Description

- No carry-over effects (relaxed in the extensions);
- Stationary unobservables $\nu_{i,t}$, but time-varying separable fixed effects allowed.

(back)

Adaptive experiment: details 2

Learning Rate

1. Strong concavity $\alpha_{k,t} \propto 1/t$;
2. Strict quasi concavity $\alpha_{k,t} \propto \frac{1}{||\hat{V}_{k,t}||_2 \sqrt{T}}$.

In practice recommended 2.
(back)

Details on global optimality

Model

$$Y_i = s\Big(\underbrace{\frac{1}{|\mathcal{N}_i^1|} \sum_{j \in \mathcal{N}_i^1} D_j}_{\text{Spill from } j \in \mathcal{N}_i : X_j = 1}, \cdots, \frac{1}{|\mathcal{N}_i^{|\mathcal{X}|}|} \sum_{j \in \mathcal{N}_i^{|\mathcal{X}|}} D_j \Big) + \nu_i, \quad |\mathcal{X}| < \infty$$

- Finitely many types and spillovers are heterogeneous in types;
- No direct effects: effects are the opportunity costs of the intervention.

Details on global optimality cont'd

Why is it sufficient to randomize at the **individual level**?

- Ideally, we would like to find for all units

$$s\Big(\beta_1^*, \cdots, \beta_{|\mathcal{X}|}^*\Big), \quad \beta^* \in \arg\max_{\beta} s(\beta).$$

Details on global optimality cont'd

Why is it sufficient to randomize at the **individual level**?

- Ideally, we would like to find for all units

$$s\Big(\beta_1^*, \cdots, \beta_{|\mathcal{X}|}^*\Big), \quad \beta^* \in \arg\max_{\beta} s(\beta).$$

- This choice (weakly) dominates any assignment $\mathbf{D} \in \{0,1\}^N$;

Details on global optimality cont'd

Why is it sufficient to randomize at the **individual level**?

- Ideally, we would like to find for all units

$$s\Big(\beta_1^*, \cdots, \beta_{|\mathcal{X}|}^*\Big), \quad \beta^* \in \arg\max_{\beta} s(\beta).$$

- This choice (weakly) dominates any assignment $\mathbf{D} \in \{0,1\}^N$;
- If we randomize treatments **independently**, we can show

$$D_i | X_i = x \sim_{i.i.d.} \mathrm{Bern}(\beta_x^*) \Rightarrow \frac{1}{|\mathcal{N}_i^x|} \sum_{j \in \mathcal{N}_i^x} D_j \simeq \beta_x^*!$$

Details on global optimality cont'd

Why is it sufficient to randomize at the **individual level**?

- Ideally, we would like to find for all units

$$s\Big(\beta_1^*, \cdots, \beta_{|\mathcal{X}|}^*\Big), \quad \beta^* \in \arg\max_{\beta} s(\beta).$$

- This choice (weakly) dominates any assignment $\mathbf{D} \in \{0,1\}^N$;
- If we randomize treatments **independently**, we can show

$$D_i | X_i = x \sim_{i.i.d.} \mathrm{Bern}(\beta_x^*) \Rightarrow \frac{1}{|\mathcal{N}_i^x|} \sum_{j \in \mathcal{N}_i^x} D_j \simeq \beta_x^*!$$

$\Rightarrow$ Using bounds on derivatives of $s(\cdot)$, we can obtain the desired result.

(back)

Regret

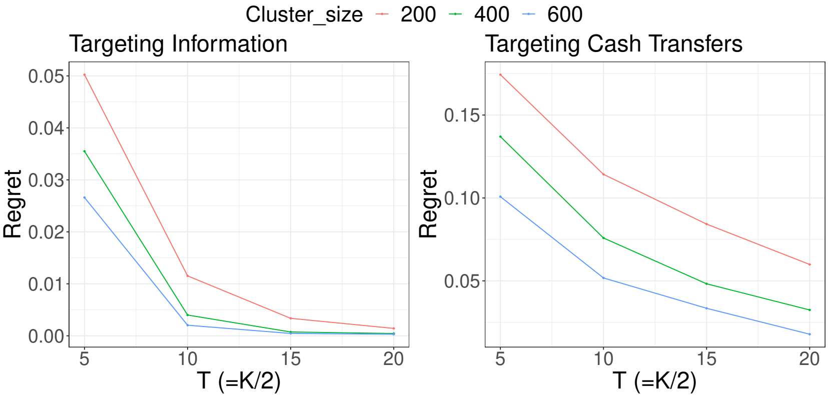
Cluster_size 200 400 600
Targeting Information
Targeting Cash Transfers
Regret
0.05
0.04
0.03
0.02
0.01
0.00
0.15
0.10
0.05
5
10
15
20
T (=K/2)

(back)

Welfare improvement relative to grid-search

	Information				Cash Transfer			
$T =$	5	10	15	20	5	10	15	20
200	-0.02	0.01	0.04	0.03	0.29	0.39	0.52	0.32
400	0.00	0.02	0.02	0.03	0.46	0.44	0.58	0.56
600	0.00	0.03	0.02	0.02	0.48	0.48	0.62	0.64

(back)

Coverage from simulations

	Information				Cash Transfer			
$K =$	10	20	30	40	10	20	30	40
$n = 200$	0.90	0.95	0.90	0.90	0.92	0.94	0.91	0.89
$n = 400$	0.98	0.96	0.90	0.92	0.98	0.96	0.89	0.93
$n = 600$	0.97	0.97	0.95	0.94	0.97	0.99	0.96	0.93

(back)

Vary η_n

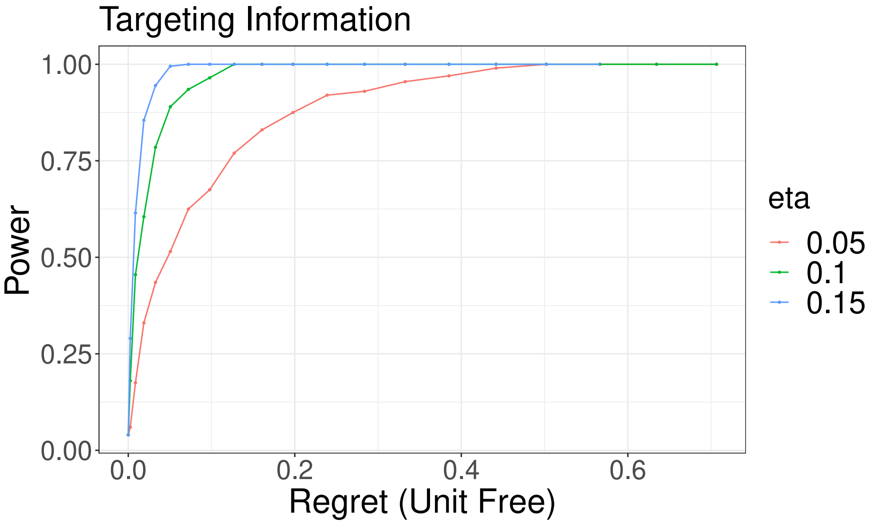

(back)

Global interference: example

Asm Individuals depend on everybody else through some statistics;

Ex. Change in average number of treated friends in the cluster.

Global interference: example

Asm Individuals depend on everybody else through some statistics;

Ex. Change in average number of treated friends in the cluster.

Example: outcomes depend on the average treatment in the cluster

$$Y_i = \underbrace{\frac{1}{N}\sum_{i=1}^{N} D_i}_{Global} + \nu_i$$

Global interference: example

Asm Individuals depend on everybody else through some statistics;

Ex. Change in average number of treated friends in the cluster.

Example: outcomes depend on the average treatment in the cluster

$$Y_i = \underbrace{\frac{1}{N}\sum_{i=1}^{N} D_i}_{Global} + \nu_i = \beta + \nu_i + \underbrace{(\bar{D} - \beta)}_{O_p(n^{-1/2})\text{:can be ignored}}$$

Global interference: example

Asm Individuals depend on everybody else through some statistics;

Ex. Change in average number of treated friends in the cluster.

Example: outcomes depend on the average treatment in the cluster

$$Y_i = \underbrace{\frac{1}{N}\sum_{i=1}^{N} D_i}_{Global} + \nu_i = \beta + \nu_i + \underbrace{(\bar{D} - \beta)}_{O_p(n^{-1/2})\text{:can be ignored}}$$

With high probability:

$$\left|\hat{V} - V(\beta)\right| = \mathcal{O}_p\Big(\underbrace{\sqrt{\frac{1}{n\eta_n^2} + \eta_n}}_{\text{from before without } \gamma_n} + \underbrace{1/\sqrt{n\eta_n^2}}_{\text{additional term}} \Big)$$

(back)

Clusters with different distributions

- Cluster specific distributions: $X_i^k \sim_{i.i.d.} F_X^k$;

Clusters with different distributions

- Cluster specific distributions: $X_i^k \sim_{i.i.d.} F_X^k$;
- Treatments randomized with some probabilities: $D_i^k \sim \mathrm{Bern}(\beta_k)$;

Clusters with different distributions

- Cluster specific distributions: $X_i^k \sim_{i.i.d.} F_X^k$;
- Treatments randomized with some probabilities: $D_i^k \sim \mathrm{Bern}(\beta_k)$;
- Outcomes can be written as $Y_i^k = y(X_i, \beta_k) + \varepsilon_i^k$

Clusters with different distributions

- Cluster specific distributions: $X_i^k \sim_{i.i.d.} F_X^k$;
- Treatments randomized with some probabilities: $D_i^k \sim \mathrm{Bern}(\beta_k)$;
- Outcomes can be written as $Y_i^k = y(X_i, \beta_k) + \varepsilon_i^k$

Target:

$$\omega_k = \mathbb{E}_{\beta_k}[Y_i^k] - \mathbb{E}_{\beta_{k+1}}[Y_i^k]$$

Clusters with different distributions

- Cluster specific distributions: $X_i^k \sim_{i.i.d.} F_X^k$;
- Treatments randomized with some probabilities: $D_i^k \sim \mathrm{Bern}(\beta_k)$;
- Outcomes can be written as $Y_i^k = y(X_i, \beta_k) + \varepsilon_i^k$

Target:

$$\omega_k = \mathbb{E}_{\beta_k}[Y_i^k] - \mathbb{E}_{\beta_{k+1}}[Y_i^k] = \int y(x;\beta_k)dF_X^k - \underbrace{\int y(x;\beta_{k+1})dF_X^k}_{\text{not identified}}$$

Matching clusters with embeddings

Bias characterization:

$$\left|\mathbb{E}\left[\bar{Y}_k - \bar{Y}_{k+1}\right] - \omega_k\right| \leq \sup_{y(\cdot,\beta_{k+1}) \in \mathcal{M}} \int y(x;\beta_{k+1})(dF_X^k - dF_X^{k+1})$$

$$= \underbrace{\mu_{X^k} - \mu_{X^{k+1}}}_{embeddings}$$

Matching clusters with embeddings

Bias characterization:

$$\left|\mathbb{E}\left[\bar{Y}_k - \bar{Y}_{k+1}\right] - \omega_k\right| \leq \sup_{y(\cdot,\beta_{k+1})\in\mathcal{M}} \int y(x;\beta_{k+1})(dF_X^k - dF_X^{k+1})$$

$$= \underbrace{\mu_{X^k} - \mu_{X^{k+1}}}_{embeddings}$$

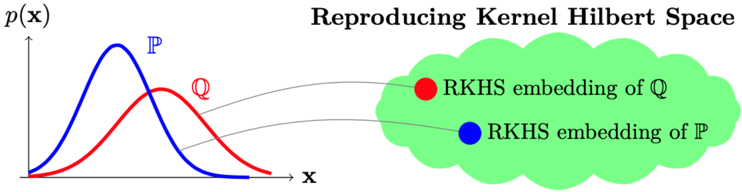

Matching clusters with embeddings

Bias characterization:

$$\left|\mathbb{E}\left[\bar{Y}_k - \bar{Y}_{k+1}\right] - \omega_k\right| \leq \sup_{y(\cdot,\beta_{k+1})\in\mathcal{M}} \int y(x;\beta_{k+1})(dF_X^k - dF_X^{k+1})$$

$$= \underbrace{\mu_{X^k} - \mu_{X^{k+1}}}_{embeddings}$$

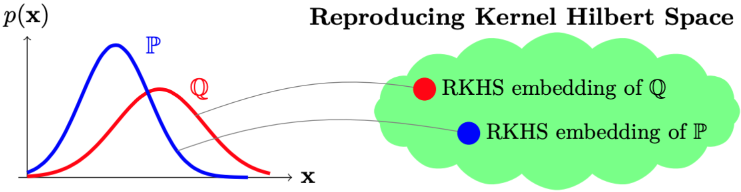

- Consistent estimator of $\hat{\mu}_X$ are available (Smola et al., 2007; Muandet, 2016; ...)

(back)

Dynamic treatments

$\Rightarrow$ Consider a sequence of parameters

$$\{\beta_t\}_{t=1}^{T}, \quad D_{i,t} \sim \mathrm{Bern}(\beta_t)$$

Dynamic treatments

⇒ Consider a sequence of parameters

$$\{\beta_t\}_{t=1}^{T}, \quad D_{i,t} \sim \mathrm{Bern}(\beta_t)$$

- Take assignments drawn with parameters $\{\beta_t\}$. Let

$$Y_{i,t} = \Gamma(\beta_t, \beta_{t-1}) + \varepsilon_{i,t}.$$

Dynamic treatments

$\Rightarrow$ Consider a sequence of parameters

$$\{\beta_t\}_{t=1}^T, \quad D_{i,t} \sim \mathrm{Bern}(\beta_t)$$

- Take assignments drawn with parameters $\{\beta_t\}$. Let

$$Y_{i,t} = \Gamma(\beta_t, \beta_{t-1}) + \varepsilon_{i,t}.$$

- Long-run welfare

$$\sum_{t=1}^T q^t \Gamma(\beta_t, \beta_{t-1}).$$

Dynamic treatments

⇒ Consider a sequence of parameters

$$\{\beta_t\}_{t=1}^{T}, \quad D_{i,t} \sim \mathrm{Bern}(\beta_t)$$

- Take assignments drawn with parameters $\{\beta_t\}$. Let

$$Y_{i,t} = \Gamma(\beta_t, \beta_{t-1}) + \varepsilon_{i,t}.$$

- Long-run welfare

$$\sum_{t=1}^{T} q^t \Gamma(\beta_t, \beta_{t-1}).$$

- Future decisions must depend on past decisions: optimal path

$$\beta_t^* = h_\theta(\beta_{t-1}^*, \beta_{t-2}^*)$$

Dynamic treatments

$\Rightarrow$ Consider a sequence of parameters

$$\{\beta_t\}_{t=1}^T, \quad D_{i,t} \sim \mathrm{Bern}(\beta_t)$$

- Take assignments drawn with parameters $\{\beta_t\}$. Let

$$Y_{i,t} = \Gamma(\beta_t, \beta_{t-1}) + \varepsilon_{i,t}.$$

- Long-run welfare

$$\sum_{t=1}^{T} q^t \Gamma(\beta_t, \beta_{t-1}).$$

- Future decisions must depend on past decisions: optimal path

$$\beta_t^* = h_\theta(\beta_{t-1}^*, \beta_{t-2}^*)$$

$\Rightarrow$ Main idea: marginal effect improves rate of convergence. (back)

Policy gradient: estimation with $t = 2$

$\Rightarrow$ Many K and only $t = 2$.

a. Construct triads of clusters;

Policy gradient: estimation with $t = 2$

$\Rightarrow$ Many K and only $t = 2$.

a. Construct triads of clusters;
b. Randomize pairs (β_1, β_2) to each triad;

Policy gradient: estimation with $t = 2$

$\Rightarrow$ Many K and only $t = 2$.

a. Construct triads of clusters;
b. Randomize pairs (β_1, β_2) to each triad;
c. Estimate the marginal effects from past and present;

Policy gradient: estimation with $t = 2$

$\Rightarrow$ Many K and only $t = 2$.

a. Construct triads of clusters;
b. Randomize pairs (β_1, β_2) to each triad;
c. Estimate the marginal effects from past and present;
d. Estimate the functions $\frac{\partial \Gamma(\beta_1, \beta_2)}{\partial \beta_1}, \frac{\partial \Gamma(\beta_1, \beta_2)}{\partial \beta_2}$ regressing the noisy measurement from c. onto (β_1, β_2)

Policy gradient: estimation with $t = 2$

$\Rightarrow$ Many K and only $t = 2$.

a. Construct triads of clusters;
b. Randomize pairs (β_1, β_2) to each triad;
c. Estimate the marginal effects from past and present;
d. Estimate the functions $\frac{\partial \Gamma(\beta_1,\beta_2)}{\partial \beta_1}, \frac{\partial \Gamma(\beta_1,\beta_2)}{\partial \beta_2}$ regressing the noisy measurement from c. onto (β_1, β_2)

$\Rightarrow$ Estimate

$$\theta^* = \arg\max_{\theta \in \Theta} \widetilde{W}(\{\beta_t^*\}), \quad \beta_t^* = h_\theta(\beta_{t-1}^*, \beta_{t-2}^*).$$

with gradient descent, with gradient estimated using $\widehat{\frac{\partial \Gamma(\beta_1,\beta_2)}{\partial \beta_1}}, \widehat{\frac{\partial \Gamma(\beta_1,\beta_2)}{\partial \beta_2}}$.

(back)

Endogenous network formation

Asm In each cluster k (dropping the superscript k)

$$Y_{i,t} = r\Big(D_{i,t}, D_{\mathcal{N}_{i,t}}, U_i, U_{\mathcal{N}_{i,t}}, A_{i,\cdot,t}, \nu_{i,t}\Big) + \tau_k + \alpha_t.$$

Endogenous network formation

Asm In each cluster k (dropping the superscript k)

$$Y_{i,t} = r\Big(D_{i,t}, D_{\mathcal{N}_{i,t}}, U_i, U_{\mathcal{N}_{i,t}}, A_{i,\cdot,t}, \nu_{i,t}\Big) + \tau_k + \alpha_t.$$

Asm Let

$$A^k_{i,j,t} = f(U^k_i, U^k_j, \underbrace{D^k_{i,t}, D^k_{j,t}}_{Treatments})1\{i \sim j\}$$

Lem We can write (in cluster k), for $\nu_{i,t}$ i.i.d. exogenous unobservables

Endogenous network formation

Asm In each cluster k (dropping the superscript k)

$$Y_{i,t} = r\Big(D_{i,t}, D_{\mathcal{N}_{i,t}}, U_i, U_{\mathcal{N}_{i,t}}, A_{i,\cdot,t}, \nu_{i,t}\Big) + \tau_k + \alpha_t.$$

Asm Let

$$A^k_{i,j,t} = f(U^k_i, U^k_j, \underbrace{D^k_{i,t}, D^k_{j,t}}_{Treatments})1\{i \sim j\}$$

Lem We can write (in cluster k), for $\nu_{i,t}$ i.i.d. exogenous unobservables

$$\mathbb{E}_\beta[Y_{i,t}|D_{i,t}] = m(D_{i,t}, \beta) + \tau_k + \alpha_t, \quad \text{for} \quad D_{i,t} \sim \text{Bern}(\beta);$$

Endogenous network formation

Asm In each cluster k (dropping the superscript k)

$$Y_{i,t} = r\Big(D_{i,t}, D_{\mathcal{N}_{i,t}}, U_i, U_{\mathcal{N}_{i,t}}, A_{i,\cdot,t}, \nu_{i,t}\Big) + \tau_k + \alpha_t.$$

Asm Let

$$A^k_{i,j,t} = f(U^k_i, U^k_j, \underbrace{D^k_{i,t}, D^k_{j,t}}_{Treatments})1\{i \sim j\}$$

Lem We can write (in cluster k), for $\nu_{i,t}$ i.i.d. exogenous unobservables

$$\mathbb{E}_\beta[Y_{i,t}|D_{i,t}] = m(D_{i,t}, \beta) + \tau_k + \alpha_t, \quad \text{for} \quad D_{i,t} \sim \mathrm{Bern}(\beta);$$

(a) Network changes over time and depends on treatment assignments;

(b) The people I can connect ex-ante is invariant + no carry-overs;

(c) Marginal effect also captures the effect on the network. (back beg), (back ext)

Individuals depend on higher-order degree neighbors

Consider

$$Y_{i,1} = r\Big(D_{i,1}, D_{\mathcal{N}_i}, U_i, U_{\mathcal{N}_i}, A_{i,\cdot}, \nu_i\Big) + \tau_k + \underbrace{\text{Rem}}_{\text{effect from remaining units}}$$

For consistency, we require that $\mathrm{Rem} = o(1/\eta_n)$.

$\Rightarrow$ the higher order effect needs to decay at a sufficiently faster rate than the degree. (back ext)

Alternative condition:

$\Rightarrow$ By letting $\mathrm{Var}(\sqrt{n}\hat{V}_{k,k+1}) = \mathcal{O}(\rho_n)$, we have $\rho_n/(n\eta_n^2) = o(1)$ (Theorem 3.1).

References I

[1] Adusumilli, K., F. Geiecke, and C. Schilter (2019). Dynamically optimal treatment allocation using reinforcement learning. *arXiv preprint arXiv:1904.01047*.

[2] Agarwal, A., O. Dekel, and L. Xiao (2010). Optimal algorithms for online convex optimization with multi-point bandit feedback. In *COLT*, pp. 28–40. Citeseer.

[3] Akbarpour, M., S. Malladi, and A. Saberi (2018). Just a few seeds more: value of network information for diffusion. *Available at SSRN 3062830*.

[4] Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016). Network structure and the aggregation of information: Theory and evidence from indonesia. *American Economic Review 106*(7), 1663–1704.

[5] Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, and J. Tobias (2012). Targeting the poor: evidence from a field experiment in indonesia. *American Economic Review 102*(4), 1206–40.

[6] Andrews, I., T. Kitagawa, and A. McCloskey (2019). Inference on winners. Technical report, National Bureau of Economic Research.

[7] Aronow, P. M. and C. Samii (2017). Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics 11*(4), 1912–1947.

References II

[8] Athey, S. and S. Wager (2021). Policy learning with observational data. *Econometrica 89*(1), 133–161.

[9] Bai, Y. (2019). Optimality of matched-pair designs in randomized controlled trials. *Available at SSRN 3483834*.

[10] Baird, S., J. A. Bohren, C. McIntosh, and B. Özler (2018). Optimal design of experiments in the presence of interference. *Review of Economics and Statistics 100*(5), 844–860.

[11] Basse, G. W. and E. M. Airoldi (2018). Model-assisted design of experiments in the presence of network-correlated outcomes. *Biometrika 105*(4), 849–858.

[12] Bhattacharya, D., P. Dupas, and S. Kanaya (2013). Estimating the impact of means-tested subsidies under treatment externalities with application to anti-malarial bednets. Technical report, National Bureau of Economic Research.

[13] Bloch, F., M. O. Jackson, and P. Tebaldi (2019). Centrality measures in networks. *Available at SSRN 2749124*.

[14] Bubeck, S., N. Cesa-Bianchi, et al. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning 5*(1), 1–122.

References III

[15] Cai, J., A. De Janvry, and E. Sadoulet (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics 7*(2), 81–108.

[16] Canay, I. A., J. P. Romano, and A. M. Shaikh (2017). Randomization tests under an approximate symmetry assumption. *Econometrica 85*(3), 1013–1030.

[17] Carneiro, P., J. J. Heckman, and E. Vytlacil (2010). Evaluating marginal policy changes and the average effect of treatment for individuals at the margin. *Econometrica 78*(1), 377–394.

[18] Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ. 1*(1), 451–488.

[19] Duchi, J., F. Ruan, and C. Yun (2018). Minimax bounds on stochastic batched convex optimization. In *Conference On Learning Theory*, pp. 3065–3162. PMLR.

[20] Eckles, D., B. Karrer, and J. Ugander (2017). Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference 5*(1).

[21] Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. W. Walker (2019). General equilibrium effects of cash transfers: experimental evidence from kenya. Technical report, National Bureau of Economic Research.

References IV

[22] Flaxman, A. D., A. T. Kalai, and H. B. McMahan (2004). Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*.

[23] Goldsmith-Pinkham, P. and G. W. Imbens (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics 31*(3), 253–264.

[24] Graham, B. S., G. W. Imbens, and G. Ridder (2010). Measuring the effects of segregation in the presence of social spillovers: A nonparametric approach. Technical report, National Bureau of Economic Research.

[25] Hu, Y., S. Li, and S. Wager (2021). Average treatment effects in the presence of interference. *arXiv preprint arXiv:2104.03802*.

[26] Hudgens, M. G. and M. E. Halloran (2008). Toward causal inference with interference. *Journal of the American Statistical Association 103*(482), 832–842.

[27] Ibragimov, R. and U. K. Müller (2010). t-statistic based correlation and heterogeneity robust inference. *Journal of Business & Economic Statistics 28*(4), 453–468.

[28] Ibragimov, R. and U. K. Müller (2016). Inference with few heterogeneous clusters. *Review of Economics and Statistics 98*(1), 83–96.

[29] Imai, K., G. King, and C. Nall (2009). The essential role of pair matching in cluster-randomized experiments, with application to the mexican universal health insurance evaluation. *Statistical Science 24*(1), 29–53.

[30] Jagadeesan, R., N. S. Pillai, A. Volfovsky, et al. (2020). Designs for estimating the treatment effect in networks with interference. *Annals of Statistics 48*(2), 679–712.

[31] Johari, R., H. Li, I. Liskovich, and G. Weintraub (2020). Experimental design in two-sided platforms: An analysis of bias. *arXiv preprint arXiv:2002.05670*.

[32] Karrer, B., L. Shi, M. Bhole, M. Goldman, T. Palmer, C. Gelman, M. Konutgan, and F. Sun (2021). Network experimentation at scale. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 3106–3116.

[33] Kasy, M. (2016). Partial identification, distributional preferences, and the welfare ranking of policies. *Review of Economics and Statistics 98*(1), 111–131.

[34] Kasy, M. and A. Sautmann (2019). Adaptive treatment assignment in experiments for policy choice. *Econometrica*.

[35] Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica 86*(2), 591–616.

[36] Kitagawa, T. and G. Wang (2021). Who should get vaccinated? individualized allocation of vaccines over sir network. *Journal of Econometrics*.

References VI

[37] Leung, M. P. (2020). Treatment and spillover effects under network interference. *Review of Economics and Statistics 102*(2), 368–380.

[38] Li, S. and S. Wager (2020). Random graph asymptotics for treatment effect estimation under network interference. *arXiv preprint arXiv:2007.13302*.

[39] Manresa, E. (2013). Estimating the structure of social interactions using panel data. *Unpublished Manuscript. CEMFI, Madrid*.

[40] Manski (2004). Statistical treatment rules for heterogeneous populations. *Econometrica 72*(4), 1221–1246.

[41] Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The review of economic studies 60*(3), 531–542.

[42] Rai, Y. (2018). Statistical inference for treatment assignment policies.

[43] Russo, D., B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen (2017). A tutorial on thompson sampling. *arXiv preprint arXiv:1707.02038*.

[44] Saez, E. (2001). Using elasticities to derive optimal income tax rates. *The review of economic studies 68*(1), 205–229.

[45] Sävje, F., P. M. Aronow, and M. G. Hudgens (2021). Average treatment effects in the presence of unknown interference. *The Annals of Statistics 49*(2), 673–701.

References VII

[46] Stoye, J. (2012). Minimax regret treatment choice with covariates or with limited validity of experiments. *Journal of Econometrics 166*(1), 138–156.

[47] Sutton, R. S. and A. G. Barto (2018). *Reinforcement learning: An introduction*. MIT press.

[48] Tabord-Meehan, M. (2018). Stratification trees for adaptive randomization in randomized controlled trials. *arXiv preprint arXiv:1806.05127*.

[49] Ugander, J., B. Karrer, L. Backstrom, and J. Kleinberg (2013). Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 329–337. ACM.

[50] Viviano, D. (2019). Policy targeting under network interference. *arXiv preprint arXiv:1906.10258*.

[51] Viviano, D. (2020a). Experimental design under network interference. *arXiv preprint arXiv:2003.08421*.

[52] Viviano, D. (2020b). Policy design in experiments with unknown interference. *arXiv preprint arXiv:2011.08174*.

[53] Wager, S. and K. Xu (2021). Experimenting in equilibrium. *Management Science*.