Randomization Inference for Spillover Effects

Panagiotis (Panos) Toulis
panos.toulis@chicagobooth.edu

Econometrics and Statistics
University of Chicago, Booth School of Business

WNAR/IMS 2023, Anchorage, Alaska

Introduction

Standard causal inference assumes no interference;
i.e., a unit's treatment cannot affect other units.

This assumes a simple, static world.

In many interesting problems, units interact in a complex way.
—spillovers, peer effects, contagion, equilibrium effects, etc.

Pervasive in most social studies. Can either be a nuisance to be addressed by design, or the quantity of interest.

New methods and tools are needed. Many applications:
e.g., policy making, marketplace algorithms, climate science, healthcare, etc.

Overview

Current approaches tend to be heavily model-based.

In complex domains, this causes problems with inference and even with identification.

Randomization tests (e.g., permutations) are nonparametric procedures that are **model-agnostic** and **finite-sample exact** under certain conditions.

However, they have been limited in scope.

A lot of recent research work in extending the scope of randomization tests to complex domains. I will present such a line of work today.

Motivation: Crime spillovers in Medellin, Colombia

(Collazos, 2019), (Puelz et al, 2021)

Crime spillovers from nearby **treated streets** on **control streets**?

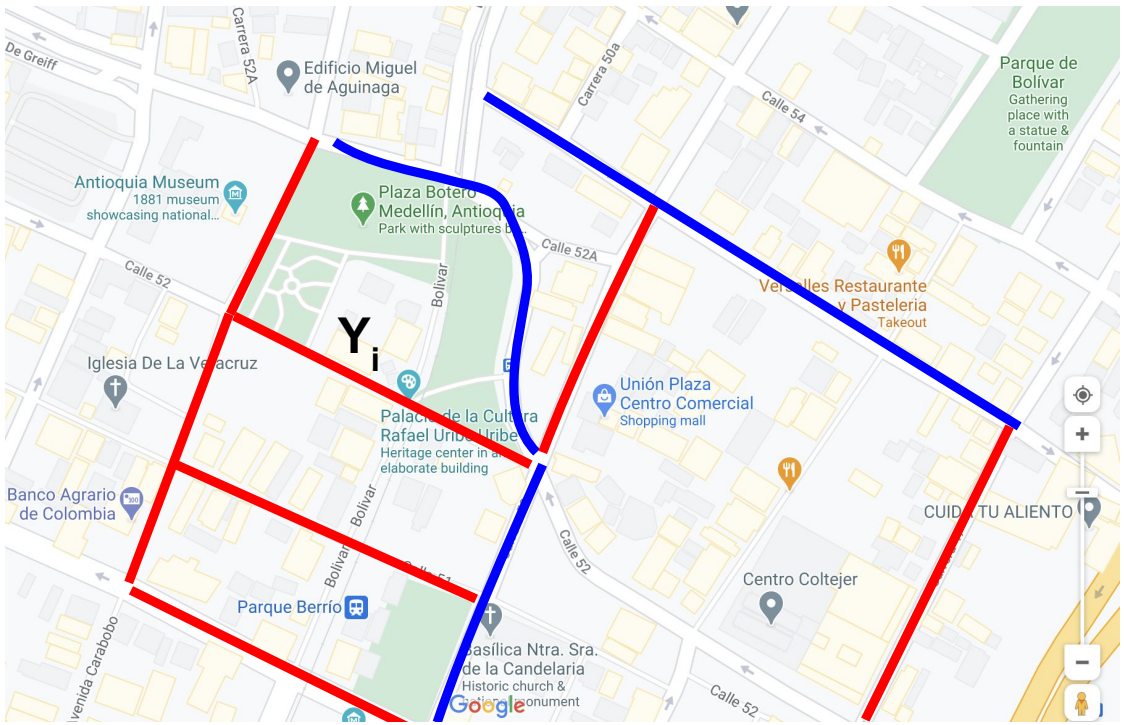

treatment = increased policing; control = baseline policing.

- What is a proper definition of a spillover effect?
- How to estimate it?

Causal Inference

Suppose data $\{(Y_i, Z_i, X_i)\}$, $i = 1, \ldots, N$.

$\mathbf{Y}$ = outcomes, $\mathbf{Z}$ = treatments, $\mathbf{X}$ = covariates (features).

$Y_i(\mathbf{z})$ is the *potential outcome* of unit i under treatment $\mathbf{z} \in \{0,1\}^N$,

Why potential outcomes?

- Used to define causal estimands: e.g., $(1/N)[\sum_i Y_i(\mathbf{1}) - Y_i(\mathbf{0})]$.
  $\hookrightarrow$ Separates the science from the statistical model.
- Makes identification assumptions more transparent and general.

Consistency assumption: $\mathbf{Y} = \mathbf{Y}(\mathbf{Z})$ — Outcomes are only a function of treatment.Variation only comes from treatment assignment ("design-based inference"). See, e.g., (Abadie et al, 2020).

No Interference

In classical causal inference, every unit i has only two potential outcomes, namely "$Y_i(0), Y_i(1)$" for control and treatment, respectively.

This is equivalent to assuming that

$$Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \text{ for all } \mathbf{z}, \mathbf{z}', i \text{ if } z_i = z_i' \; - \text{"SUTVA" (Rubin, 1974).}$$

No Interference

In classical causal inference, every unit i has only two potential outcomes, namely "$Y_i(0), Y_i(1)$" for control and treatment, respectively.

This is equivalent to assuming that

$$Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \text{ for all } \mathbf{z}, \mathbf{z}', i \text{ if } z_i = z_i' \; - \text{"SUTVA" (Rubin, 1974).}$$

However, in many problems there is treatment interference. (spillovers, peer effects, contagion, dynamics etc.)

Under interference, a unit is exposed to "something more" than Z_i, a sum effect from the entire population treatment, $\mathbf{Z}$.

Think of a vaccine trial. A control unit (unvaccinated) is still "protected" by treated units (vaccinated) in proximity.

Effective treatments

Under interference, it is popular to use treatment exposures, $f_i(\mathbf{Z}) \in \mathbb{F}$.

Although not necessary, it is useful to think that the exposure is the "effective treatment" (Manski, 2013).

Assumption

$$Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \quad \textit{for all } \mathbf{z}, \mathbf{z}', i \textit{ if } f_i(\mathbf{z}) = f_i(\mathbf{z}').$$

Examples of treatment exposure:

- $f_i(\mathbf{z}) = z_i$. Standard, no interference setting.
- $f_i(\mathbf{z}) = \left(z_i, \sum_{j:d(i,j)<d_0} z_j\right)$. "Treatments nearby matter".
  We use this in the Medellin application.
- $f_i(\mathbf{z}) = (z_i, \mathbf{z}_{\text{neighborhood}_i})$.

Effective treatments

Under interference, it is popular to use treatment exposures, $f_i(\mathbf{Z}) \in \mathbb{F}$.

Although not necessary, it is useful to think that the exposure is the "effective treatment" (Manski, 2013).

Assumption

$$Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \quad \textit{for all } \mathbf{z}, \mathbf{z}', i \textit{ if } f_i(\mathbf{z}) = f_i(\mathbf{z}').$$

Examples of treatment exposure:

- $f_i(\mathbf{z}) = z_i$. Standard, no interference setting.
- $f_i(\mathbf{z}) = \left(z_i, \sum_{j:d(i,j)<d_0} z_j\right)$. "Treatments nearby matter".
  We use this in the Medellin application.
- $f_i(\mathbf{z}) = (z_i, \mathbf{z}_{\text{neighborhood}_i})$.

Goal: Learn the effect of $f_i(\mathbf{z})$ on outcome Y_i?

Wait, could I just fit a regression?

A popular approach is still to fit:

$$Y_i = \alpha + \beta Z_i + \gamma \underbrace{f_i(\mathbf{Z})}_{\text{exposure}} + \delta' X_i + \epsilon_i.$$

But there are caveats:

- Correct specification is crucial.
- $f_i(\mathbf{Z})$ may have a complex correlation structure with other covariates, and usually an underlying network as well.
- Cannot accurately quantify uncertainty, in general.
- Asymptotics on $\hat{\gamma}$ may well be intractable.

Finally, it is not uncommon to use a model with Ys on the "left and right" of the regression. This is almost never a good idea. (Angrist, 2019)

Wait, could I just fit a regression?

A popular approach is still to fit:

$$Y_i = \alpha + \beta Z_i + \gamma \underbrace{f_i(\mathbf{Z})}_{\text{exposure}} + \delta' X_i + \epsilon_i.$$

But there are caveats:

- Correct specification is crucial.
- $f_i(\mathbf{Z})$ may have a complex correlation structure with other covariates, and usually an underlying network as well.
- Cannot accurately quantify uncertainty, in general.
- Asymptotics on $\hat{\gamma}$ may well be intractable.

Finally, it is not uncommon to use a model with Ys on the "left and right" of the regression. This is almost never a good idea. (Angrist, 2019)

* These problems can be avoided in experimental studies through methods such as randomization tests.

Hypotheses for spillovers

Let's consider a large family of hypotheses about $f_i(\mathbf{z})$:

$$H_0 : Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \text{ for all } i, \mathbf{z}, \mathbf{z}' \text{ st } f_i(\mathbf{z}), f_i(\mathbf{z}') \in \mathbb{F}_0 \subseteq \mathbb{F}.$$

(Manski, 2009), (Aronow, 2012), (T. and Kao, 2013), (Bowers et al., 2013), (Athey et al., 2019), Basse et al, 2019), (Puelz et al, 2021).

This null tests whether certain kinds of exposures are equivalent in their outcomes;
e.g., $\mathbb{F}_0 = \{$"control-spillovers", "pure-control"$\}$ (coming soon).

If $\mathbb{F}_0 = \mathbb{F}$, then the null can be tested **exactly** through the celebrated Fisherian randomization test (Fisher, 1935) (Lehmann and Romano, 2005).

General Idea: Fisher's Randomization Test

When $\mathbb{F}_0 = \mathbb{F}$, then all exposures give identical outcomes under the null. This is **equivalent to** the global null of no effect:

$$H_0 : Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \text{ for all } \mathbf{z}, \mathbf{z}', i.$$

This can be tested through Fisher's randomization test (Fisher, 1935),

1. Calculate test statistic, $T = t(\mathbf{Z}, \mathbf{Y})$; e.g., ANOVA statistic
2. pval $= \mathsf{E}[t(\mathbf{Z}', \mathbf{Y}) > T], \quad \mathbf{Z}' \sim P.$

General Idea: Fisher's Randomization Test

When $\mathbb{F}_0 = \mathbb{F}$, then all exposures give identical outcomes under the null. This is **equivalent to** the global null of no effect:

$$H_0 : Y_i(\mathbf{z}) = Y_i(\mathbf{z}') \text{ for all } \mathbf{z}, \mathbf{z}', i.$$

This can be tested through Fisher's randomization test (Fisher, 1935),

1. Calculate test statistic, $T = t(\mathbf{Z}, \mathbf{Y})$; e.g., ANOVA statistic
2. pval $= \mathsf{E}[t(\mathbf{Z}', \mathbf{Y}) > T], \quad \mathbf{Z}' \sim P.$

♠ Works because $t(\mathbf{Z}', \mathbf{Y}) \overset{H_0}{=} t(\mathbf{Z}', \mathbf{Y}') \overset{d}{=} T.$

An assessment of FRT

Main advantages:

- The test is exact in finite samples. No asymptotics.
- Not necessary to have correct Y-model specification.
- The test is robust. Same answer under transformations of Y.

Some disadvantages:

- Can only test "strong" hypotheses. (Currently, a lot of research activity in this area).
- Cannot generalize to population.

An assessment of FRT

Main advantages:

- The test is exact in finite samples. No asymptotics.
- Not necessary to have correct Y-model specification.
- The test is robust. Same answer under transformations of Y.

Some disadvantages:

- Can only test "strong" hypotheses. (Currently, a lot of research activity in this area).
- Cannot generalize to population.

* But can we use it for the Medellin application?

Medellin example

Crime spillovers from nearby **treated streets** on **control streets**?

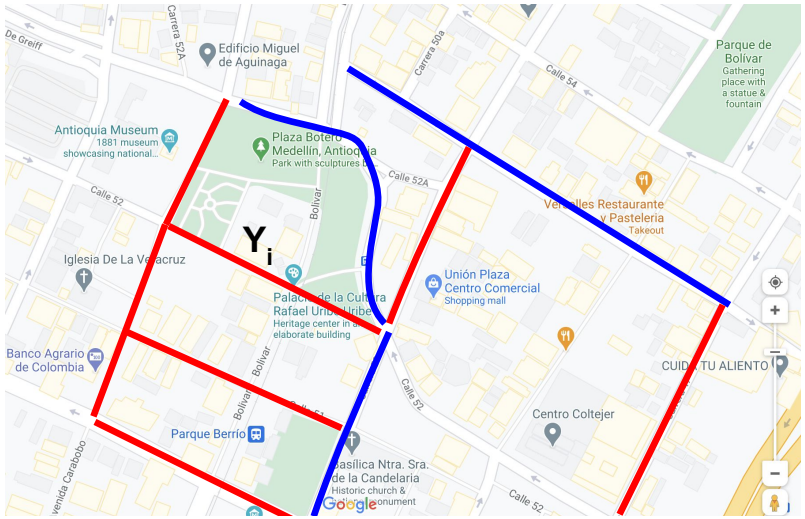

Recall

$$f_i(\mathbf{z}) = \Big(z_i, \sum_{j:d(i,j)<d_0} z_j\Big).$$

Let $\mathbb{F}_0 = \{$"control-spillovers", "pure-control"$\}$ where

- "control-spillovers" if $f_i(\mathbf{z}) = (0, +)$.
- "pure-control" if $f_i(\mathbf{z}) = (0, 0)$.

Thus, $H_0 : Y_i($"control-spill"$) = Y_i($"pure-control"$)$ precludes spillover effects on the treated.

FRT problems under interference

Suppose we resample $\mathbf{z}'$ in the FRT as shown below:

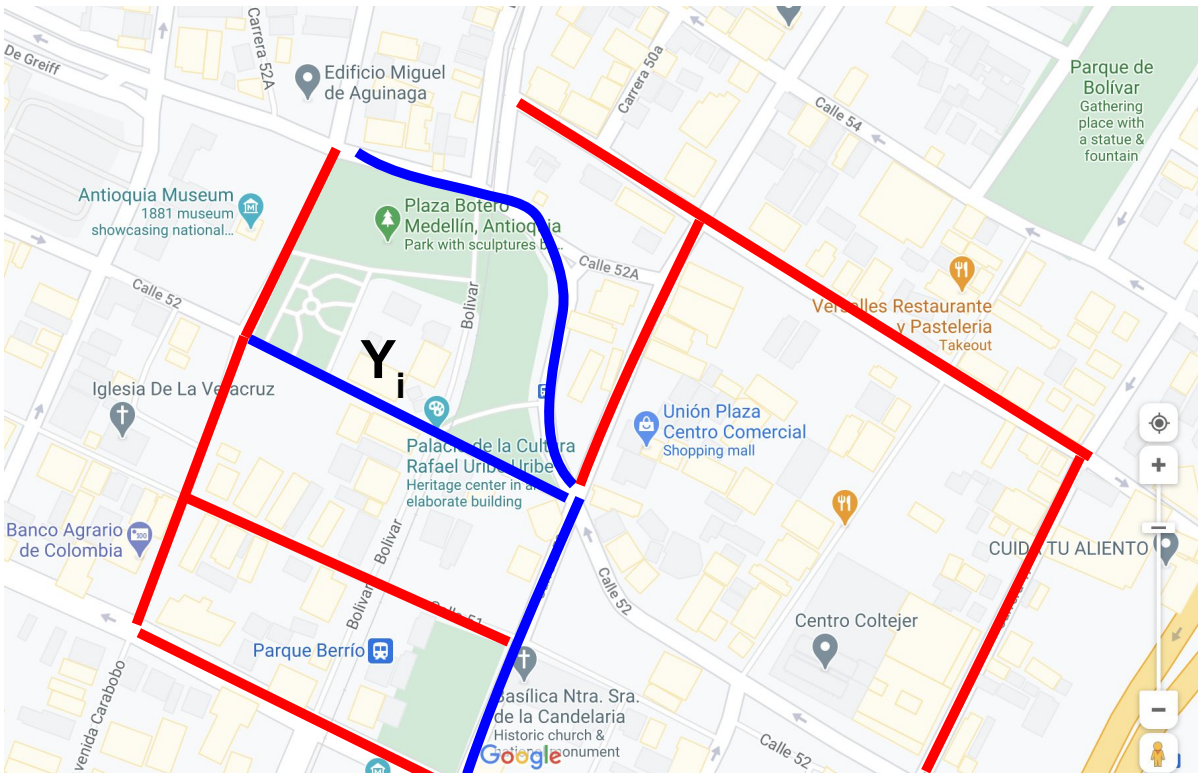

The exposure of i is not in $\mathbb{F}_0$. Thus, $Y_i(\mathbf{z}')$ **cannot be imputed** under H_0.

This means that, under interference, we cannot naively apply the FRT.

Recent developments

Athey et al (2019) recently proposed a general approach to apply FRTs under interference:

1. Randomly select **subset** of units U.
2. Calculate $\mathcal{Z}_U = \{\mathbf{z} : Y_i(\mathbf{z})$ can be imputed under H_0 for all $i \in U\}$.
3. Run FRT focused on $U, \mathcal{Z}_U$ (discard all other data/randomizations). Resample $\mathbf{Z}'$ as $P(\mathbf{Z}')/P(\mathcal{Z}_U)$.

Practical issues:

i. Random selection of "focal units" (U) can be too naive. Does not exploit the problem structure, and can lead to data waste (see Medellin example before)
ii. Step 2 is computationally demanding. Needs to enumerate $\mathcal{Z}_U$.

Recent developments

Basse et al (2019) showed that random selection of focal units is not necessary.

In fact, any selection $P(U \mid \mathbf{Z})$ is ok can as long as in we resample from the correct conditional randomization distribution:

$$P(\mathbf{Z} \mid U) \propto P(U \mid \mathbf{Z}) \cdot P(\mathbf{Z}).$$

- $P(U \mid \mathbf{Z})$ is the "conditioning mechanism". Controlled by the analyst.
- e.g., the approach in (Athey et al, 2019) is to define $P(U \mid \mathbf{Z}) = P(U) \propto 1$, and so $P(\mathbf{Z} \mid U) \propto P(\mathbf{Z})$.

While this result leads to more general tests (and often more power), it does not specify how to choose $P(U \mid \mathbf{Z})$.

Conditioning the FRT for spillovers

Puelz et al. (2021) developed a general method to construct such valid conditioning for FRTs under spillovers.

Main idea: Connect every pair $(i, \mathbf{z})$ iff $f_i(\mathbf{z}) \in \mathbb{F}_0 \Rightarrow$ null exposure graph.

Conditioning the FRT for spillovers

Puelz et al. (2021) developed a general method to construct such valid conditioning for FRTs under spillovers.

Main idea: Connect every pair $(i, \mathbf{z})$ iff $f_i(\mathbf{z}) \in \mathbb{F}_0 \Rightarrow$ null exposure graph.

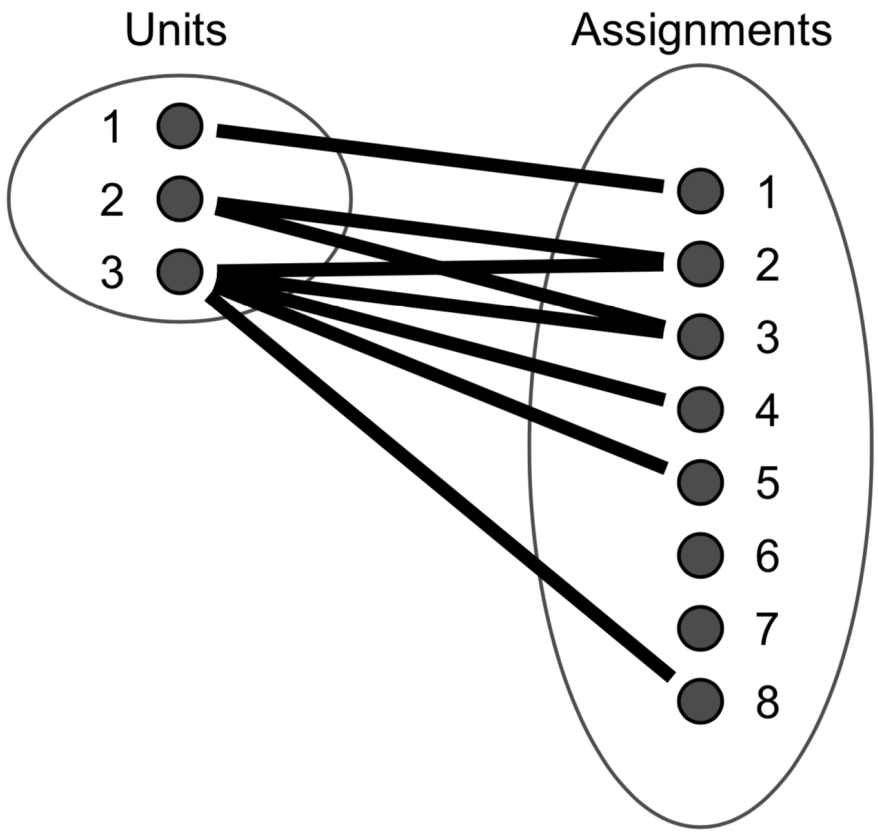

Null-exposure graph

The null exposure graph has some very nice properties.

- It encodes the problem structure (only a function of H_0).
- The density of the graph reveals the "support" for testing H_0. (is H_0 easy or hard to test?)
- An edge in the graph is equivalent to imputability.

Theorem

Under any conditional randomization test that uses focal units U and assignments $\mathcal{Z}_U$, the potential outcomes are imputable if and only if the sets $(U, \mathcal{Z}_U)$ form a biclique in the null exposure graph.

FRT for spillovers

This leads to the following extension of the classical FRT.

1. Calculate NE graph based on H_0.
2. Calculate a *biclique decomposition* of NE.
3. Condition the randomization on the biclique that contains $\mathbf{Z}$. Resample any $\mathbf{z}$ in the biclique proportional to $P(\mathbf{z})$.

FRT for spillovers

This leads to the following extension of the classical FRT.

1. Calculate NE graph based on H_0.
2. Calculate a *biclique decomposition* of NE.
3. Condition the randomization on the biclique that contains $\mathbf{Z}$. Resample any $\mathbf{z}$ in the biclique proportional to $P(\mathbf{z})$.

This algorithm automatically generates a conditioning mechanism that fits the particular problem structure.

FRT for spillovers

This leads to the following extension of the classical FRT.

1. Calculate NE graph based on H_0.
2. Calculate a *biclique decomposition* of NE.
3. Condition the randomization on the biclique that contains $\mathbf{Z}$. Resample any $\mathbf{z}$ in the biclique proportional to $P(\mathbf{z})$.

This algorithm automatically generates a conditioning mechanism that fits the particular problem structure.

Essentially, it defines the following conditioning mechanism

$$P(U \mid \mathbf{Z}) = 1\{U = \text{units}(\text{clique}_{\mathbf{Z}})\}. \tag{1}$$

where "clique$_{\mathbf{z}}$" is the (unique) clique that contains $\mathbf{z}$;
"units(c)" is the unit set of clique c.

Thus, $P(\mathbf{z} \mid U) \propto P(\mathbf{z})$ for any $\mathbf{z}$ in the biclique that contains $\mathbf{Z}$.

Medellin application

- treated units
- "pure-control" units
- "control-spillovers" units
- focal units

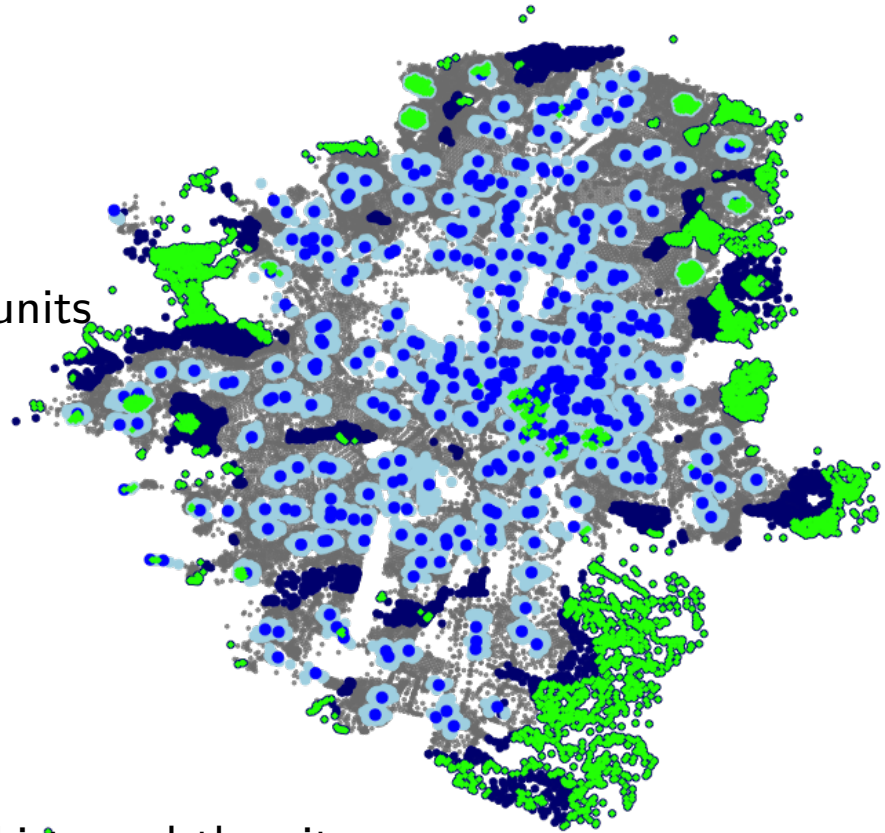

Only units in the outskirts and the city center are pertinent to testing H_0.

The picture reveals a complex conditioning structure for this particular H_0.

Very hard to obtain such conditioning by random sampling, if not impossible.

Results

Spillover effect is marginally significant.

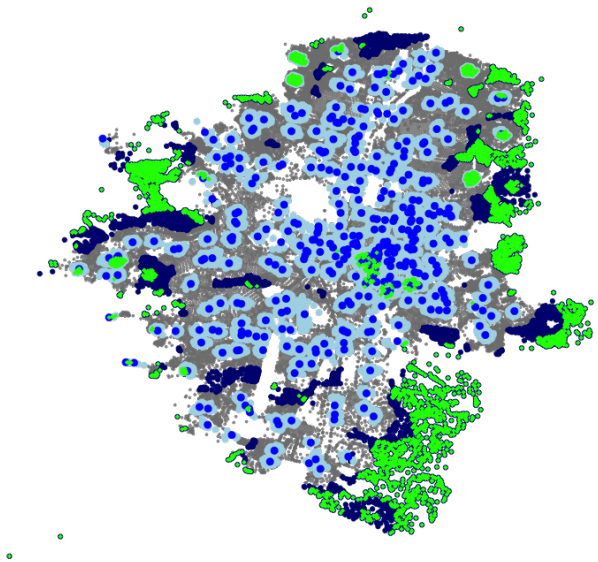

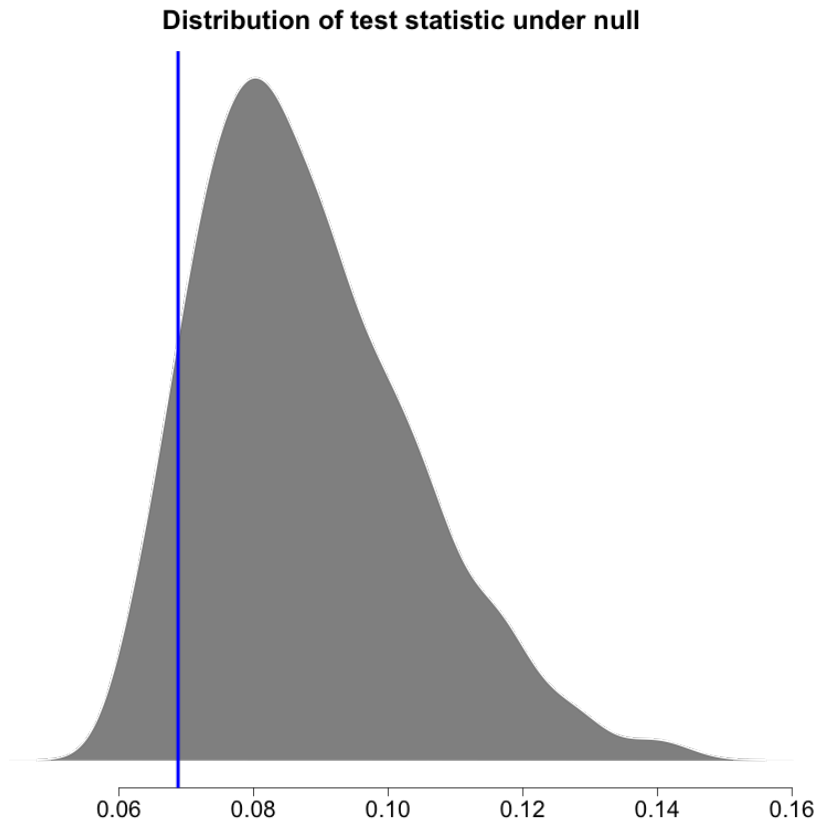

Results

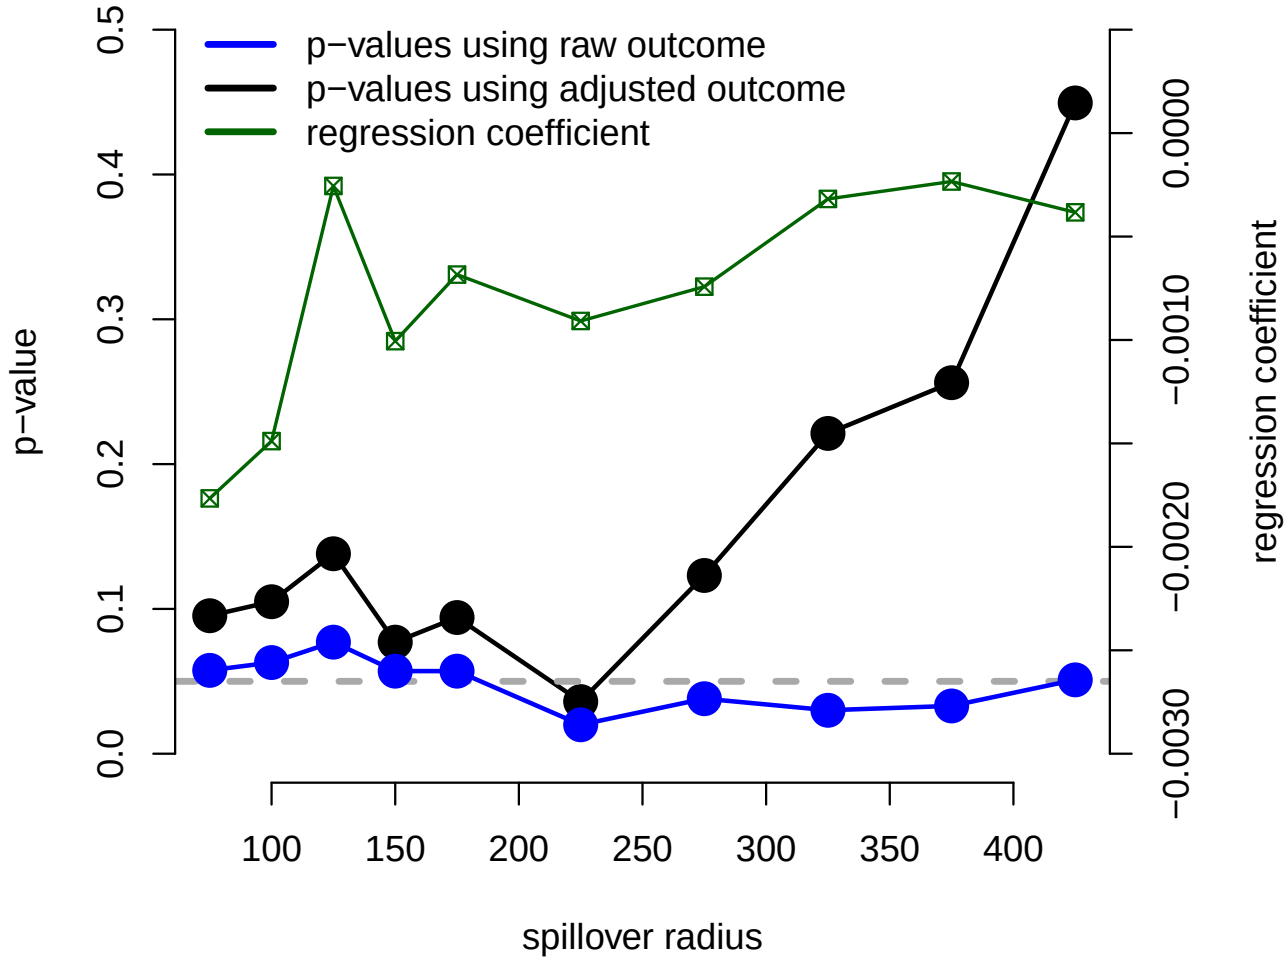
p−values using raw outcome
p−values using adjusted outcome
regression coefficient
p−value
0.0
0.1
0.2
0.3
0.4
0.5
regression coefficient
0.0000
−0.0010
−0.0020
−0.0030
100
150
200
250
300
350
400
spillover radius

Concluding remarks

Randomization tests can be extended to problems with interference.

These are robust, **finite-sample exact** procedures.

Many open problems remain. — Inference, average spillover effects, etc.

More challenges: Marketplace dynamics, game theory etc.

Thank you!

Basse, Ding, Feller, Toulis "Randomization tests for group formation experiments" , (R&R, 2023)

Puelz, Basse, Feller, Toulis "A graph-theoretic approach to randomization tests of causal effects under interference" , (JRSS-B, 2021)

Basse, Feller, Toulis, "Randomization tests of causal effects under interference" (Biometrika, 2019)