CHICAGO BOOTH

Randomization Tests of Causal Effects Under General Interference

David Puelz

joint with Panos Toulis

March 4, 2019

Medellín, Colombia

Medellín, Colombia

Medellín, Colombia

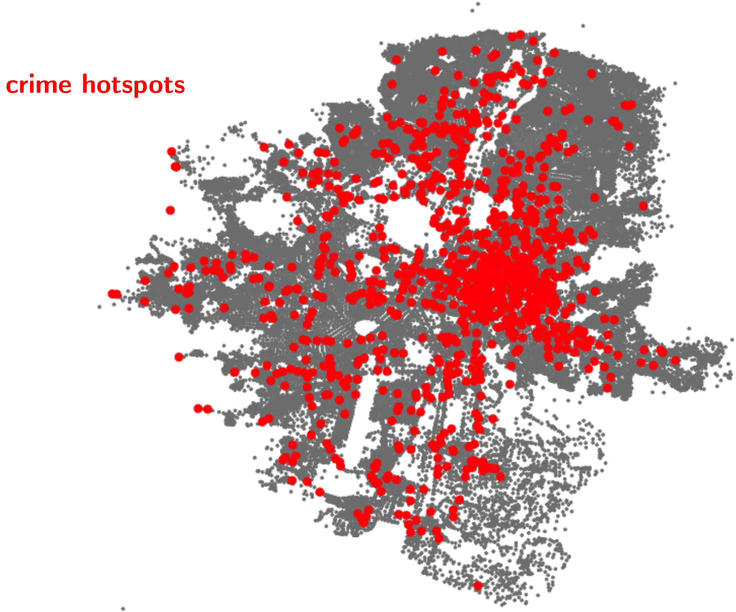

Medellín, Colombia

observed treatment

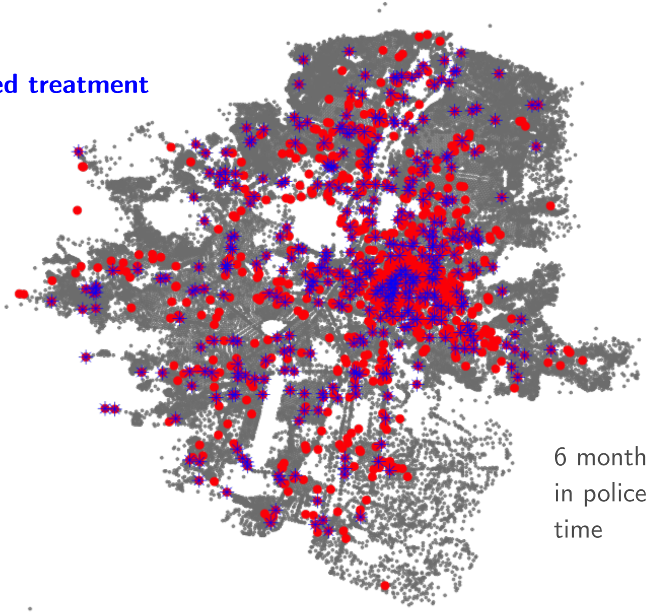

6 month increase in police patrolling time

Experiment and data

Units and treatment assignment

- 37,055 total streets (units)
- 967 streets are identified as crime "hotspots"
- 384 are treated with increased police presence

Access to randomizations based on the design, pr(Z)

Experiment and data

Units and treatment assignment

- 37,055 total streets (units)
- 967 streets are identified as crime "hotspots"
- 384 are treated with increased police presence

Access to randomizations based on the design, pr(Z)

Outcomes and covariates

- Crime counts on all streets (murders, car and motorbike thefts, personal robberies, assaults)
- Survey data on hotspot streets
- Characteristics of hotspots (distance from school, bus stop, rec center, church, neighborhood, ...)

Questions we aim to answer

How does the intervention affect crime?

$\rightarrow$ direct effect?

$\rightarrow$ spillovers to adjacent streets?

Questions we aim to answer

How does the intervention affect crime?

$\rightarrow$ direct effect?

$\rightarrow$ spillovers to adjacent streets?

We will answer these through hypothesis testing.

We would like to be model-free, so we will use the randomization method of inference.

A classical test

Define potential outcome of unit i under assignment Z: $Y_i(Z)$

i.e., number of thefts over measurement interval.

Y = vector of observed outcomes.

Assume: $Y_i(Z)$ depends only on Z_i (no interference)

$$\mathbf{H_0}: \quad Y_i(Z_i = 0) = Y_i(Z_i = 1) \text{ for every } i.$$

We can use a Fisher exact test here!

$\mathbf{H_0}: \quad Y_i(Z_i = 0) = Y_i(Z_i = 1)$ for every i.

The procedure:

Choose test statistic $T = T(y, z)$ (e.g., difference in means).

1. $T_{\text{obs}} = T(Y, Z)$.
2. Sample $Z' \sim \text{pr}(Z')$, store $T_r = T(Y', Z') \stackrel{H_0}{=} T(Y, Z')$.
3. p-value $= \mathbb{E}\left[\mathbb{1}\{T_r \geq T_{\text{obs}}\}\right]$.

Fisher exact test (1935)

$\mathbf{H_0}$: $Y_i(Z_i = 0) = Y_i(Z_i = 1)$ for every i.

The procedure:

Choose test statistic $T = T(y, z)$ (e.g., difference in means).

1. $T_{\text{obs}} = T(Y, Z)$.
2. Sample $Z' \sim \text{pr}(Z')$, store $T_r = T(Y', Z') \overset{H_0}{=} T(Y, Z')$.
3. p-value $= \mathbb{E}[\mathbb{1}\{T_r \geq T_{\text{obs}}\}]$.

Proof of validity:

$$T(Y', Z') \overset{H_0}{=} T(Y, Z') \overset{d}{=} T(Y, Z)$$

"$T_{\text{obs}} \sim T_r$ (under null)"

Why is this great?

- Fisher test is exact.
- No model for Y.
- Valid in finite samples.
- Robustness since it is a rank test (the same cannot be said for regression).

The original assumption ...

Assume: $Y_i(Z)$ depends only on Z_i (no interference)
$\rightarrow$ not very realistic for our application.

In reality, $Y_i(Z)$ is **exposed** to (depends on) multiple parts of Z.

The original assumption ...

Assume: $Y_i(Z)$ depends only on Z_i (no interference)
$\rightarrow$ not very realistic for our application.

In reality, $Y_i(Z)$ is **exposed** to (depends on) multiple parts of Z.

New question that assumes interference: Is there a difference in outcome between short-range and pure control streets?

Answering this question under interference

Let's suppose, for a given Z, unit i's **exposure** lives in the set

$$\{\text{short-range}, \text{pure control}, \text{neither}\} = \{a, b, c\} = \mathcal{E}.$$

Unit i's exposure function, $f_i : \{0,1\}^N \to \mathcal{E}$. Maps Z to exposure.

Now, assume: $Y_i(Z)$ depends only on $f_i(Z)$. We want to test:

$$\mathbf{H_0}: \quad Y_i(a) = Y_i(b) \text{ for every } i.$$

Can we just use a Fisher exact test again?

Not quite ...

Recall, observed $T \sim$ randomized T for things to work:

$$T(Y', Z') \not\overset{H_0}{=} T(Y, Z') \overset{d}{=} T(Y, Z)$$

The null only assumes 2 of the 3 exposures have equal outcomes

$\mathbf{H_0}$: $\quad Y_i(a) = Y_i(b) \overset{?}{=} Y_i(c)$ for every i

In this case, the null is **not sharp.** We cannot impute potential outcomes Y' freely under any Z'.

Existing approaches and our contribution

We need to find units only exposed to *a* or *b* under some set of assignments ... called **focal units.**

→ make $\mathbf{H_0}$ conditionally sharp (so that $Y' \overset{H_0}{=} Y$)

Aronow 2012, Athey et al. 2017 – Sample focals, enumerate Z

- computational challenges

Basse et al. 2018 – Conditioning mechanisms

- conditioning difficult to execute
  easier when interference has structure
  (e.g. two-stage designs).

Existing approaches and our contribution

We need to find units only exposed to *a* or *b* under some set of assignments ... called **focal units.**

→ make $\mathbf{H_0}$ conditionally sharp (so that $Y' \overset{H_0}{=} Y$)

Aronow 2012, Athey et al. 2017 – Sample focals, enumerate Z

- computational challenges

Basse et al. 2018 – Conditioning mechanisms

- conditioning difficult to execute
  easier when interference has structure
  (e.g. two-stage designs).

Our contribution: A constructive, general approach to find focal units and assignments to make the null sharp.

Revisiting the null based on exposure functions

$$\mathbf{H_0}: \; Y_i(Z) = Y_i(Z') \text{ for every } i, Z, Z',$$

$$\text{such that } f_i(Z), f_i(Z') \in \{a, b\}.$$

$Y_i(Z)$ – potential outcome for street i.

$Z, Z^{'}$ – assignment vectors $\in \{0,1\}^N$.

f_i – deterministic exposure function (takes in Z, outputs exposure).

$\{a, b\}$ – set of possible exposures for units ($\subseteq \text{range}(f_i) = \mathcal{E}$).

Testing $Y_i(a) = Y_i(b) \, \forall i$

Given a null hypothesis and assignment from pr(Z), we know which units are exposed to either a or b using $f_i(\cdot)$.

This is a binary relationship!
How can we visualize?

Our main contribution: The null exposure graph

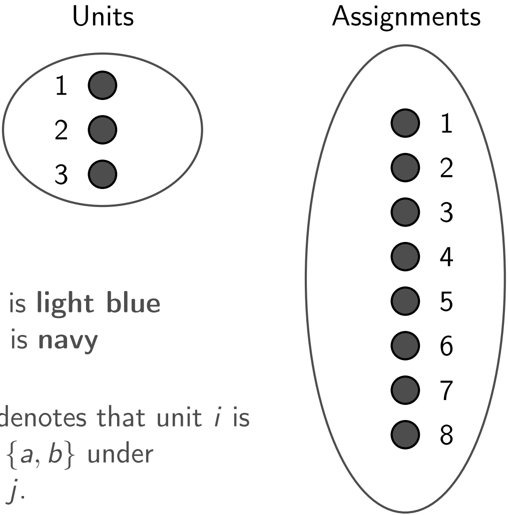

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

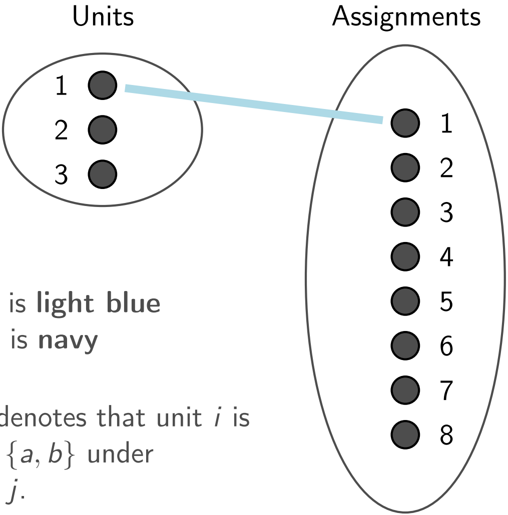

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

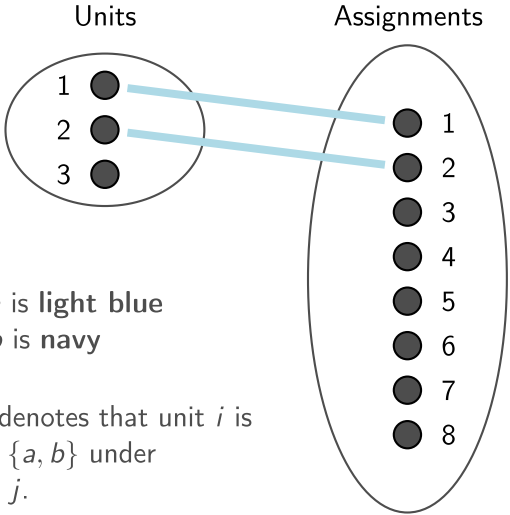

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

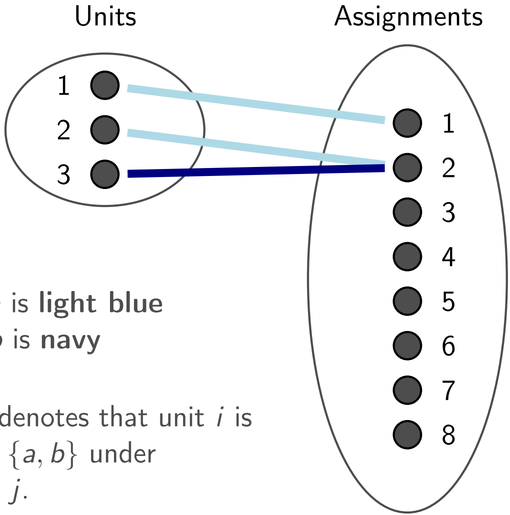

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

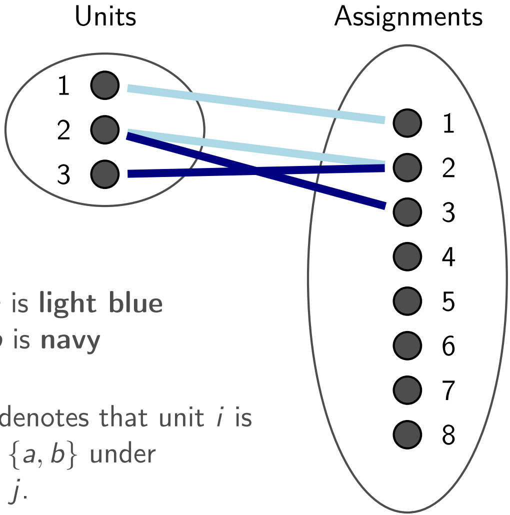

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

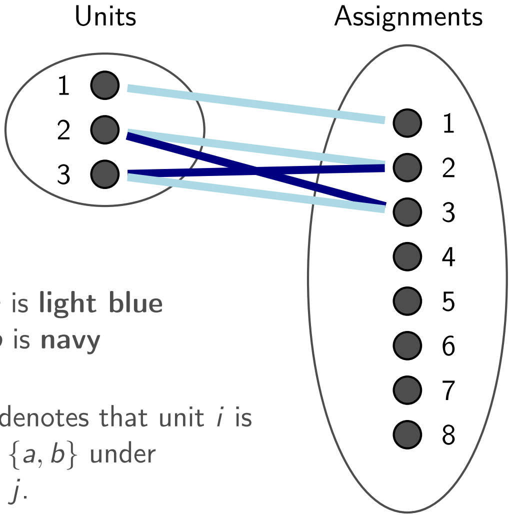

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

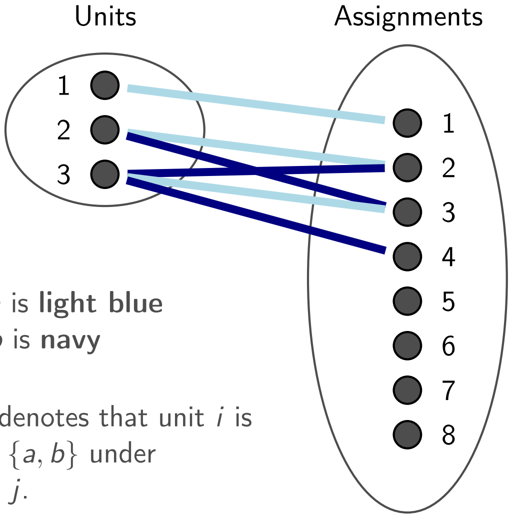

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Our main contribution: The null exposure graph

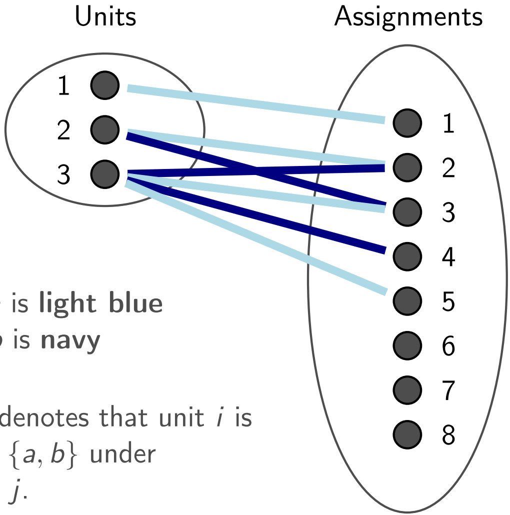

Exposure a is **light blue**
Exposure b is **navy**

edge (i, j) denotes that unit i is exposed to $\{a, b\}$ under assignment j.

Introducing the null exposure graph

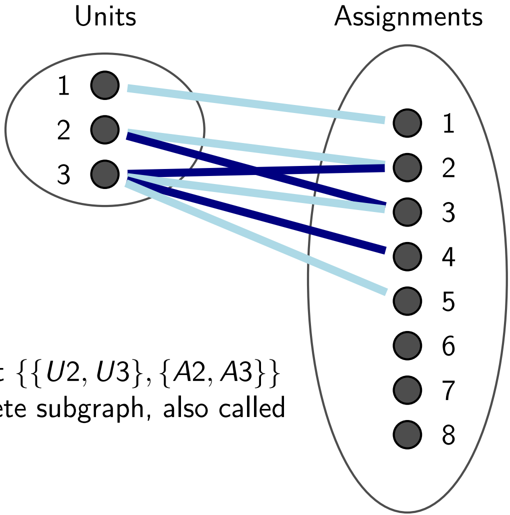

Notice that $\{\{U2, U3\}, \{A2, A3\}\}$ is a complete subgraph, also called a *biclique*!

Why are these bicliques useful?

Within a biclique, every unit is exposed to $\{a, b\}$ under any assignment.

i.e.: If Z_{obs} is in biclique, we can impute potential outcomes, and $H_{\mathbf{0}}$ is **sharp** in the biclique.

Let's outline the method ...

Conditional biclique method

$\rightarrow$ A null exposure graph uniquely defined given $\mathbf{H_0}$.

$\rightarrow$ A test statistic $T = T(y, z)$.

1. **Decompose:** Compute biclique decomposition of null exposure graph. Pick out biclique with Z_{obs}, call it C.

2. **Condition:** Compute test statistic values with units and assignments only in C.

3. **Summarize:** p-value $= \mathbb{E}_{Z_C}\left[\mathbb{1}\{T_C \geq T_{\text{obs}}\}\right]$.

   Here, $P(Z_C) \propto pr(Z_C)\mathbb{1}\{Z_C \in C\}$

Why is this a valid method?

Clique test statistics: $T_C = T(Y_C, Z_C)$

*T is defined only in C by **condition** step in method

For every Z, Z', we need to show $T(Y', Z') \stackrel{d}{=} T(Y, Z) \mid C$

Proof:

$$T(Y', Z') \stackrel{*}{=} T(Y'_C, Z'_C) \stackrel{H_0}{=} T(Y_C, Z'_C) \stackrel{d}{=} T(Y_C, Z_C) \stackrel{*}{=} T(Y, Z)$$

Considerations

- Finding bicliques is hard, actually, **NP-hard**[1]

- The method is **constructive**, still needs to be optimized

  i.e., different biclique decompositions will have different power properties, but all are **valid**!

[1] We use Binary Inclusion-Maximal Biclustering Algorithm, which uses a divide and conquer method to find bicliques.

Example: Is there a short-range spillover effect?

$$\mathbf{H_0}: \; Y_i(Z) = Y_i(Z') \text{ for every } i, Z, Z',$$

$$\text{such that } f_i(Z), f_i(Z') \in \{a, b\}.$$

$$f_i(Z) := \begin{cases} \text{short-range} & Z_i = 0, \text{dist}_i < 125\text{m} \\ \text{control} & Z_i = 0, \text{dist}_i > 500\text{m} \\ \text{neither} & \text{else} \end{cases}$$

$$\{a, b\} := \{\text{short-range}, \; \text{control}\}$$

$$\text{dist}_i := \text{distance to closest treated street.}$$

Returning to the map

The observed assignment

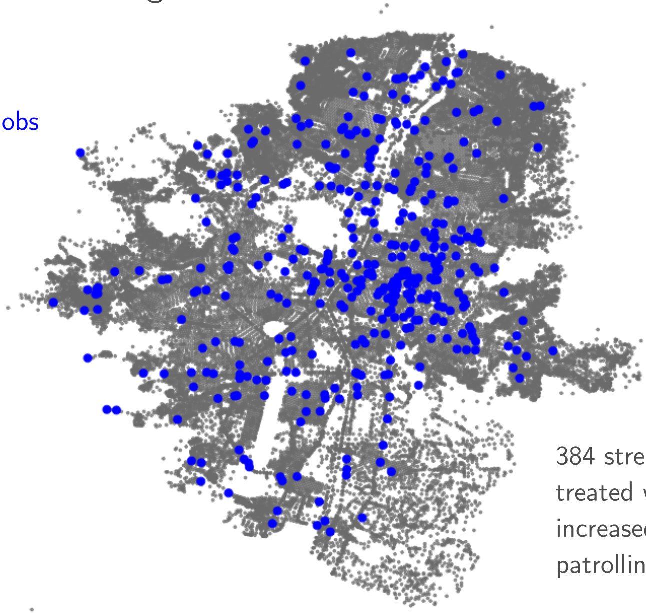

384 streets are treated with increased police patrolling

The observed assignment

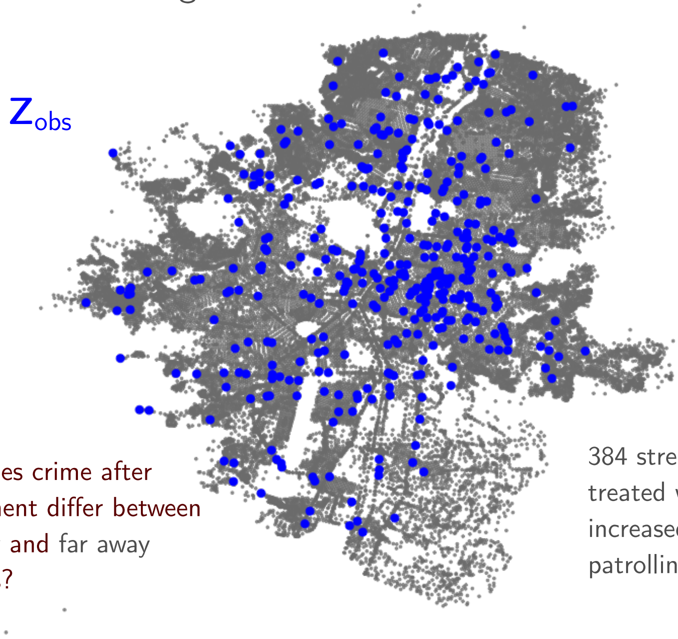

Q: Does crime after treatment differ between nearby and far away streets?

384 streets are treated with increased police patrolling

Short-range spillover units (exposure “a”)

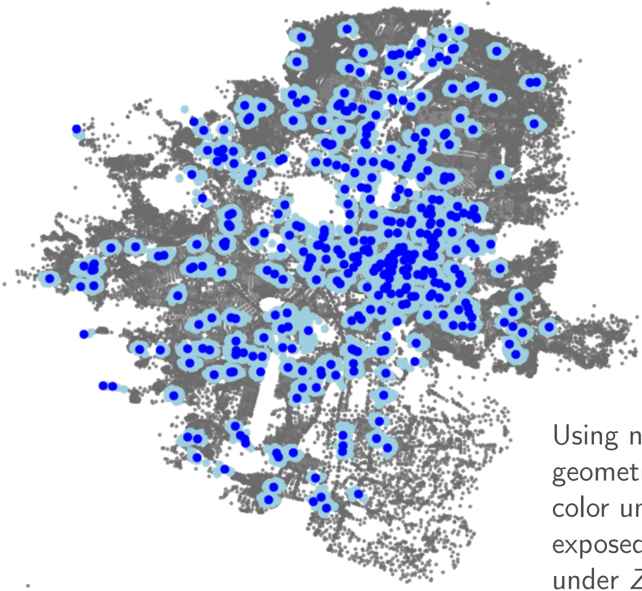

Using network geometry, color units exposed to “a” under Z_{obs}

Pure control units (exposure "b")

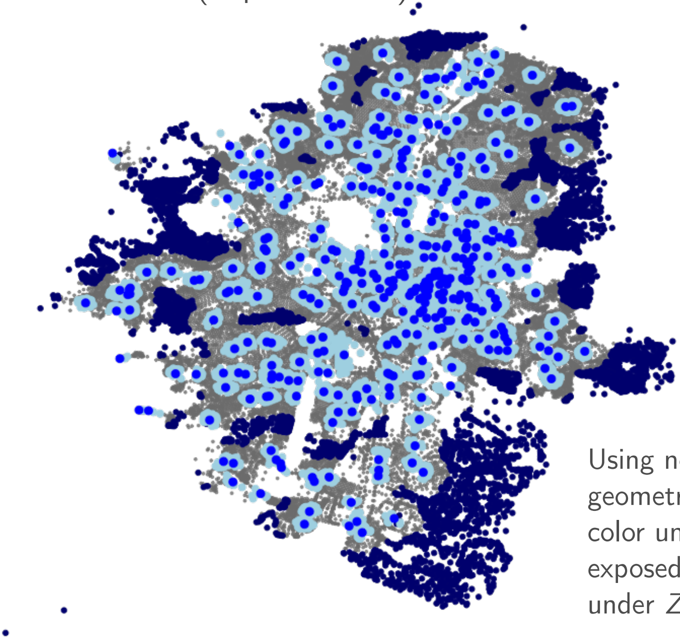

Using network geometry, color units exposed to "b" under Z_{obs}

We can remake these pictures for every assignment Z drawn from pr(Z) ...

We can remake these pictures for every assignment Z drawn from pr(Z) ...

→ The output is our null exposure graph!

Null exposure graph

navy, light blue,
and white

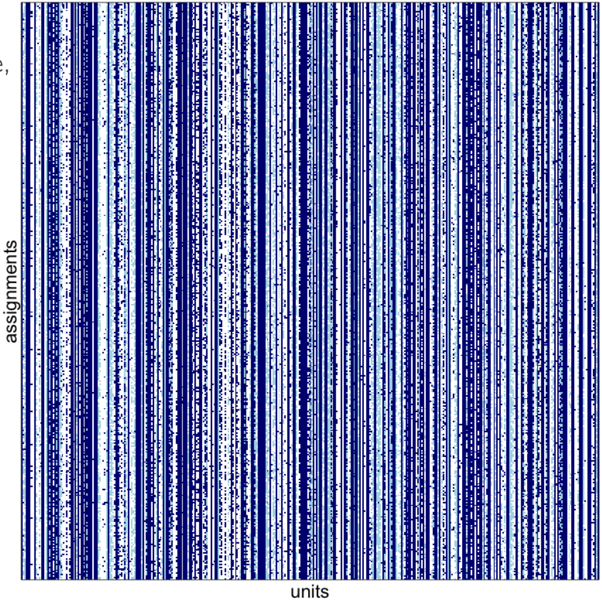

Biclique containing the observed assignment

only navy and light blue!

Where are the focal units?

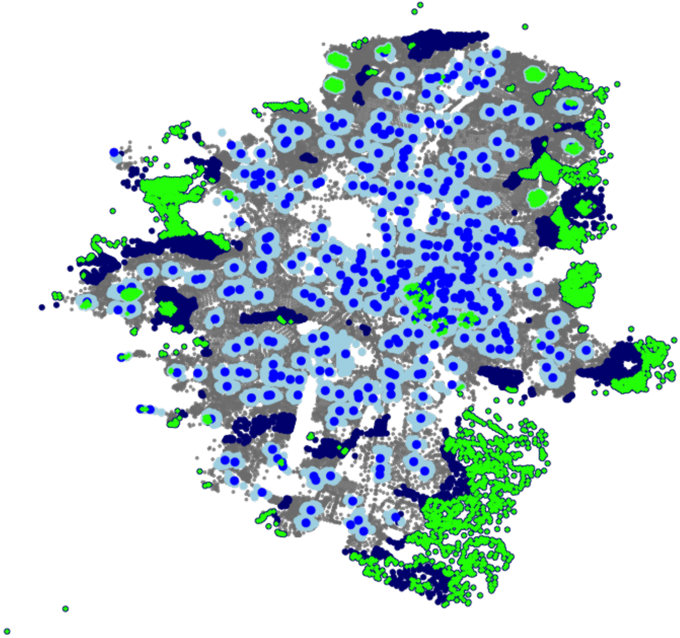

A test of the null

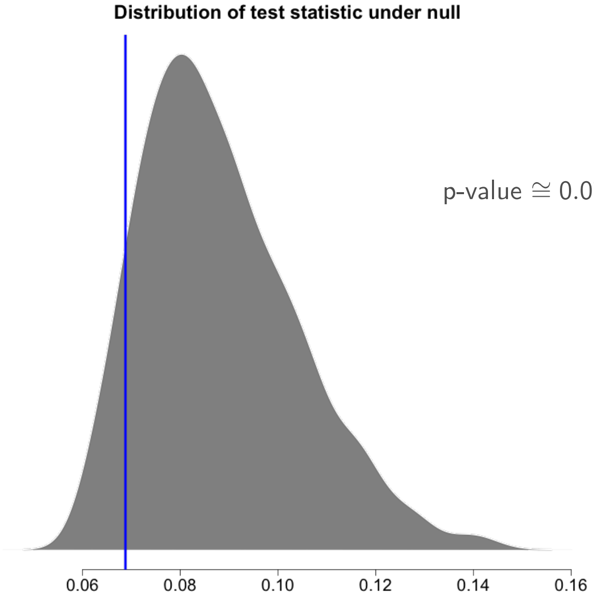

p-value $\cong 0.07$

Concluding thoughts

- New method is presented for testing causal effects under general interference using null exposure graphs and bicliques.

- Structure is placed on null hypothesis through exposure functions.

- More interesting work to be done to improve the method and test interesting hypotheses!