

Lecture 7: Interference and Experiment Design

DATA 8020: Advanced Causal Inference


Jin-Hong Du

Institute of Data Science, The University of Hong Kong

October 23, 2026

Today: what does an experiment measure when people interact?

Lecture question: what does randomization measure, and how can the design measure the effect we want?



By the end, you should be able to:

- explain what an A/B test measures when participants affect each other;
- define the causal effect we want and choose how to randomize; and
- recognize what remains open when connections, exposures, or market prices are uncertain.

One person's treatment can change another's outcome: **interference**.


Before the break: interaction, effects, then design

- Part 1 — interference and spillover effects: what effect do we want?
- Part 2 — exposure mappings and causal effects: what summary of others' treatments matters?
- Part 3 — network experiment design: what do we randomize?

Interference and spillover effects


From earlier lectures: what changes today?

Part 1 — interference: what changes when one person's treatment affects another?




- Earlier lectures assumed no interference between patients: each patient's outcome depends on their own treatment or treatment history.
- Lectures 3–6 *observed* the registry; today we **randomize** an adherence programme on the same patients, now linked by clinics and a peer network.

Units interact in three ways



network



market



time

- **Network:** infection spreads along edges; higher vaccine coverage protected unvaccinated neighbours.
- **Market:** units compete for a limited resource; job counselling displaced untreated job seekers.
- **Time:** when one market is switched between arms, each period is a unit; yesterday's treatment still acts today, so periods interfere with each other.

Perez-Heydrich et al. (2014); Crépon et al. (2013); Blake and Coey (2014)

With two units, each has four potential outcomes

	unit 2: 0	unit 2: 1
unit 1: 0	$Y_1(0,0)$	$Y_1(0,1)$
unit 1: 1	$Y_1(1,0)$	$Y_1(1,1)$


$Y_i(\mathbf{a})$, $\mathbf{a} \in \{0,1\}^n$: 2^n potential outcomes

When the target is unit 1's treatment effect:

- Across a column, only unit 1's own treatment changes: the **direct effect**, $Y_1(1,0) - Y_1(0,0)$.
- Across a row, only unit 2's treatment changes: the **spillover effect**, $Y_1(0,1) - Y_1(0,0)$.
- An experiment reveals one cell per unit; "treat everyone", $\mathbf{a} = \mathbf{1}$, is usually not observed.

No interference, $Y_1(a_1, a_2) = Y_1(a_1)$, drops the second argument (Cox, 1958); with no hidden versions of treatment, it forms SUTVA (Rubin, 1980).

In this example, the A/B test misses the spillover benefit



- Own enrolment adds 6 (**direct**); enrolled peers add up to 8, even to the unenrolled (**spillover**).
- **Global effect:** everyone enrolled vs. no one, $6 + 8 = 14$.
- Both arms have similar peer enrolment, so the contrast estimates only 6.
- Tripling the sample narrows the interval, not the gap to 14.

More patients improve precision; they do not change which effect this design measures.

The bias has no fixed sign



control patients also gain,
so the contrast **shrinks**

peers: diluted



enrolled patients book slots first,
so the contrast **grows**

shared capacity: inflated

- Peers: an enrolled patient encourages their peers, so control patients also improve and the A/B contrast understates the launch effect.
- Shared capacity: the programme helps enrolled patients book a clinic's limited follow-up slots first, so control patients get fewer and the contrast overstates it.
- Knowing that interference exists does not tell you which way the A/B test is wrong.


The same competition on online platforms: Johari et al. (2022b); Li et al. (2022).

Exposure mappings and causal effects

An exposure mapping summarizes who received treatment

Part 2 — causal effects: how do we define and estimate an effect when patients affect each other?


Roadmap: **Q2.1 define exposure** → Q2.2 define effects → Q2.3 estimate, given the design → Q2.4 check the mapping.



An exposure mapping summarizes who received treatment

Part 2 — causal effects: how do we define and estimate an effect when patients affect each other?

Roadmap: **Q2.1 define exposure** → Q2.2 define effects → Q2.3 estimate, given the design → Q2.4 check the mapping.




$$Y_i(\mathbf{a}) = Y_i(\mathbf{a}') \text{ whenever } d_i(\mathbf{a}) = d_i(\mathbf{a}'), \quad d_i \text{ takes a few values}$$

- $d_i(\cdot)$ is the **exposure mapping**, or **effective treatment**; a value d is an **exposure**.
- No interference: $d_i = a_i$.
- **One-hop** (only neighbours matter): $d_i = (a_i, \text{any neighbour treated})$.
- **Partial interference** (only the group matters, next slides): $d_i = (a_i, \text{group coverage})$.


Manski (2013); Aronow and Samii (2017)

Partial interference: units nested in groups

Q2.2 — effects: with exposure = (own enrolment, clinic coverage), which effects can we define?




clinic 1: many connections




clinic 2: few connections

Partial interference: units nested in groups

Q2.2 — **effects:** with exposure = (own enrolment, clinic coverage), which effects can we define?



clinic 1: many connections




clinic 2: few connections

- Two levels: **groups** (clinics) contain **units** (patients); each patient belongs to exactly one clinic.
- **Partial interference:** a patient's outcome depends only on treatments in their own clinic.
- Same structure: members in households, pupils in schools, job seekers in labour markets.
- **Coverage:** the proportion of a group's units treated; the exposure is $d_i = (a_i, \text{coverage of } i\text{'s group})$.

The term is from Sobel (2006); the clinic network is constructed.

Four effects from one design



- **Direct:** $\bar{Y}(1; \alpha') - \bar{Y}(0; \alpha')$, own enrolment, coverage fixed.
- **Indirect:** $\bar{Y}(0; \alpha') - \bar{Y}(0; \alpha)$, spillover on the unenrolled.
- **Total:** $\bar{Y}(1; \alpha') - \bar{Y}(0; \alpha) = \text{direct} + \text{indirect}$.
- **Overall:** $\bar{Y}(\alpha') - \bar{Y}(\alpha)$, where the group average is $\bar{Y}(\alpha) = \alpha \bar{Y}(1; \alpha) + (1 - \alpha) \bar{Y}(0; \alpha)$.

Total $\bar{Y}(1; \alpha') - \bar{Y}(0; \alpha)$ is the **global** effect only if $\alpha' = 1, \alpha = 0$, as in whole-clinic randomization, which loses direct and indirect.

Hudgens and Halloran (2008).

Partial interference: randomize twice

one stage: clinics



coverage 1



coverage 0

whole clinic in or out:
global effect only

one stage: patients



coverage 0.5



coverage 0.5

coverage never varies:
direct effect only

two stages



coverage 0.7



coverage 0.3

both vary:
all four effects

Partial interference: randomize twice

one stage: clinics



coverage 1



coverage 0

whole clinic in or out:

global effect only

one stage: patients



coverage 0.5



coverage 0.5

coverage never varies:

direct effect only

two stages



coverage 0.7



coverage 0.3


both vary:

all four effects

- **Stage 1:** randomly assign each clinic a coverage, $\alpha = 0.3$ or $\alpha' = 0.7$; comparing clinic averages gives the **overall** effect.
- **Stage 2:** within each clinic, randomly enrol that fraction; the spillover cancels within a clinic: the **direct** effect.
- **Both:** **indirect** and **total** compare a subgroup across coverages, so they need random coverage and random subgroups.

Design drawn from Hudgens and Halloran (2008). One connected network has no separate groups, so Part 3 uses graph designs.

Real data: cholera vaccination in Bangladesh



- **Groups:** no natural groups; the authors clustered household locations into 700 neighbourhoods and assumed no transmission between them (similar results with 400 or 1,100).
- **Coverage** was not randomized: people chose whether to join the trial, so exposure probabilities come from a participation model, unlike a two-stage design.
- (A) the direct effect falls as coverage rises; (B)–(D) indirect, total, and overall effects over pairs of coverage levels.

From Perez-Heydrich et al. (2014, Fig. 3); weighting estimators: Tchetgen Tchetgen and VanderWeele (2012).

Technical core 1: weight each unit by its exposure probability

Q2.3 — estimation: how do we estimate an effect defined through exposure, for a given design?

Estimand: $\bar{Y}(d) = \frac{1}{n} \sum_{i=1}^n Y_i(d)$

Estimator: $\hat{\bar{Y}}(d) = \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{D_i = d\} Y_i}{\pi_i(d)}, \quad \pi_i(d) = \mathbb{P}\{D_i = d\}$

- $\bar{Y}(d)$: mean outcome if every patient had exposure d ; an effect is $\tau = \bar{Y}(d) - \bar{Y}(d')$.

Aronow and Samii (2017).

Technical core 1: weight each unit by its exposure probability

Q2.3 — estimation: how do we estimate an effect defined through exposure, for a given design?

Estimand: $\bar{Y}(d) = \frac{1}{n} \sum_{i=1}^n Y_i(d)$

Estimator: $\hat{\bar{Y}}(d) = \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{D_i = d\} Y_i}{\pi_i(d)}, \quad \pi_i(d) = \mathbb{P}\{D_i = d\}$

- $\bar{Y}(d)$: mean outcome if every patient had exposure d ; an effect is $\tau = \bar{Y}(d) - \bar{Y}(d')$.
- $D_i = d_i(\mathbf{A})$: the exposure the random assignment produces; $\pi_i(d)$ is fixed by the design.

Aronow and Samii (2017).

Technical core 1: weight each unit by its exposure probability

Q2.3 — estimation: how do we estimate an effect defined through exposure, for a given design?

Estimand: $\bar{Y}(d) = \frac{1}{n} \sum_{i=1}^n Y_i(d)$

Estimator: $\hat{\bar{Y}}(d) = \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{D_i = d\} Y_i}{\pi_i(d)}, \quad \pi_i(d) = \mathbb{P}\{D_i = d\}$

- $\bar{Y}(d)$: mean outcome if every patient had exposure d ; an effect is $\tau = \bar{Y}(d) - \bar{Y}(d')$.
- $D_i = d_i(\mathbf{A})$: the exposure the random assignment produces; $\pi_i(d)$ is fixed by the design.
- On $D_i = d$, the observed Y_i equals $Y_i(d)$, and $\mathbb{E}[\mathbf{1}\{D_i = d\}] = \pi_i(d)$, so $\mathbb{E}\{\hat{\bar{Y}}(d)\} = \bar{Y}(d)$.

Aronow and Samii (2017).

Technical core 1: weight each unit by its exposure probability

Q2.3 — estimation: how do we estimate an effect defined through exposure, for a given design?


Estimand: $\bar{Y}(d) = \frac{1}{n} \sum_{i=1}^n Y_i(d)$

Estimator: $\hat{\bar{Y}}(d) = \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{1}\{D_i = d\} Y_i}{\pi_i(d)}, \quad \pi_i(d) = \mathbb{P}\{D_i = d\}$

- $\bar{Y}(d)$: mean outcome if every patient had exposure d ; an effect is $\tau = \bar{Y}(d) - \bar{Y}(d')$.
- $D_i = d_i(\mathbf{A})$: the exposure the random assignment produces; $\pi_i(d)$ is fixed by the design.
- On $D_i = d$, the observed Y_i equals $Y_i(d)$, and $\mathbb{E}[\mathbf{1}\{D_i = d\}] = \pi_i(d)$, so $\mathbb{E}\{\hat{\bar{Y}}(d)\} = \bar{Y}(d)$.
- This is the **Horvitz–Thompson** estimator; the normalized (**Hájek**) version divides by the sum of weights and can improve stability.

Aronow and Samii (2017).

Technical core 2: some variance terms cannot be estimated



- The variance of $\widehat{Y}(d)$ needs joint exposure probabilities for every pair; two tied units can never be in these two exposures at once, so that term has only a conservative bound.
- Small $\pi_i(d)$ gives large weights and unstable estimates, even when the estimator is unbiased.

Aronow and Samii (2017); design fixes this, Part 3.

Unknown interference: the estimated effect depends on the design


Q2.4 — **wrong mapping**: what if the exposure mapping is wrong or unknown?

- Standard estimators still converge, to the **expected average treatment effect** (EATE): change one unit's treatment, averaging over others' assignments drawn from the design, so a different design targets a different effect.

Unknown interference: the estimated effect depends on the design

Q2.4 — wrong mapping: what if the exposure mapping is wrong or unknown?

- Standard estimators still converge, to the **expected average treatment effect** (EATE): change one unit's treatment, averaging over others' assignments drawn from the design, so a different design targets a different effect.
- Error falls at the root- n rate if interference is bounded, more slowly if it grows, not at all if proportional.



Amount of interference per unit, d_{avg} :

- | | |
|------------------------------|----------------|
| — 1, none | — 25, bounded |
| - - - $8n^{1/4}$ | — $2.5n^{1/2}$ |
| - - - $0.25n$, proportional | |

Answer: still an average effect over the design, but not the effect of treating everyone.

From Sävje et al. (2021, supplement, Fig. B2, panel A of three designs): error of the normalized estimator against sample size.

Network experiment design

What should we randomize: patients or groups?

Part 3 — design: Part 2 took the design as given; now choose it so the target exposure is common.

Roadmap: **Q3.1 choose the design** → Q3.2 test for spillover → Q3.3 peer influence.

Why it matters: online platforms such as Facebook and LinkedIn run experiments on observed networks at scale; the designs began in computer science (KDD), then statistics (Eckles et al., 2017).

Target: the **global effect**, everyone treated vs. no one.

What should we randomize: patients or groups?


Part 3 — design: Part 2 took the design as given; now choose it so the target exposure is common.

Roadmap: **Q3.1 choose the design** → Q3.2 test for spillover → Q3.3 peer influence.


Why it matters: online platforms such as Facebook and LinkedIn run experiments on observed networks at scale; the designs began in computer science (KDD), then statistics (Eckles et al., 2017).

Target: the **global effect**, everyone treated vs. no one.

- **Rare exposures:** for a patient, “everyone treated” is “you and 4 neighbours treated”: $1/32$.
- **Diluted contrast:** both groups have half their neighbours treated: only the direct effect.
- **Clusters fix both:** neighbours treated together make full exposure common; fewer assignments.



Graph cluster randomization



- Ideal: Universe A, everyone gets A, versus Universe B, everyone gets B.
- (a) Individual randomization: an **ego** (the person we measure) has half A, half B neighbours.
- (b) **Cluster assignment**: one coin per loosely linked cluster; neighbours share the ego's version.
- For unit i , neighbours in clusters of size k_i : full exposure has probability 2^{-k_i} .

Unlike Part 2, clusters are cut from one network, not assumed sealed: edges between them leak, so bias falls but stays.

From Ugander et al. (2013, Fig. 1); registry figures are from constructed data.

Graph cluster randomization: estimand and estimator

Q3.1 — choose the design: what does a cluster design estimate, and with which estimator?

Estimand: $\tau_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \{Y_i(\mathbf{1}) - Y_i(\mathbf{0})\}$

Estimator: $\hat{\tau}_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\mathbf{1}\{E_i^1\} Y_i}{p^{k_i}} - \frac{\mathbf{1}\{E_i^0\} Y_i}{(1-p)^{k_i}} \right\}$

- E_i^a : i and all its neighbours get a (**full exposure**); k_i : clusters touching i 's neighbourhood.

Ugander et al. (2013); one coin per cluster with probability p ; $p = 1/2$ gives the 2^{-k_i} of the previous slide.

Graph cluster randomization: estimand and estimator

Q3.1 — **choose the design**: what does a cluster design estimate, and with which estimator?

Estimand: $\tau_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \{Y_i(\mathbf{1}) - Y_i(\mathbf{0})\}$

Estimator: $\hat{\tau}_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\mathbf{1}\{E_i^1\} Y_i}{p^{k_i}} - \frac{\mathbf{1}\{E_i^0\} Y_i}{(1-p)^{k_i}} \right\}$

- E_i^a : i and all its neighbours get a (**full exposure**); k_i : clusters touching i 's neighbourhood.
- Under **one-hop** interference (only own and neighbours' treatments matter), $\hat{\tau}_{\text{GTE}}$ is unbiased: Part 2's Horvitz–Thompson, $d = \text{full exposure}$.

Ugander et al. (2013); one coin per cluster with probability p ; $p = 1/2$ gives the 2^{-k_i} of the previous slide.

Graph cluster randomization: estimand and estimator

Q3.1 — **choose the design**: what does a cluster design estimate, and with which estimator?

Estimand: $\tau_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \{Y_i(\mathbf{1}) - Y_i(\mathbf{0})\}$


Estimator: $\hat{\tau}_{\text{GTE}} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\mathbf{1}\{E_i^1\} Y_i}{p^{k_i}} - \frac{\mathbf{1}\{E_i^0\} Y_i}{(1-p)^{k_i}} \right\}$

- E_i^a : i and all its neighbours get a (**full exposure**); k_i : clusters touching i 's neighbourhood.
- Under **one-hop** interference (only own and neighbours' treatments matter), $\hat{\tau}_{\text{GTE}}$ is unbiased: Part 2's Horvitz–Thompson, $d = \text{full exposure}$.
- Fewer clusters per neighbourhood (small k_i): full exposure is common, weights won't be small.

Ugander et al. (2013); one coin per cluster with probability p ; $p = 1/2$ gives the 2^{-k_i} of the previous slide.

Exact tests without a model: hold selected patients' treatments fixed

Q3.2 — test: in a **randomized** experiment, is there any spillover at all?



focal f : treatment fixed,
outcome used

others $?$: re-randomized
by the design

- Null of no spillover: $Y_i(\mathbf{a}) = Y_i(a_i)$ for every patient.
- **Core idea:** fix the **focal** units' treatments, re-randomize the rest; under the null, focal outcomes cannot change.
- **Statistic:** focal outcomes vs. treated non-focal neighbours; p -value as in Lecture 2's Fisher test.

Drawn from the construction in Athey et al. (2018); generalizations: Basse et al. (2019), Puelz et al. (2022).

Exact tests: the procedure step by step

Q3.2 — test: $H_0: Y_i(\mathbf{a}) = Y_i(a_i)$ for every patient; the design $\mathbb{P}(\mathbf{A})$ is known.


1. **Choose focal units** F before seeing outcomes, e.g. a random half of the patients.
2. **Statistic:** $T = \text{corr}_{i \in F}(Y_i - \bar{Y}_{\text{arm}(i)}, s_i(\mathbf{A}))$, where $s_i(\mathbf{A})$ is the proportion of i 's non-focal neighbours being treated.
3. **Re-draw:** for $b = 1, \dots, B$, draw $\mathbf{A}^{(b)}$ from the design with $\mathbf{A}_F$ held fixed; recompute $T^{(b)}$ with the same focal outcomes.
4. **p -value:** $p = (1 + \#\{b : |T^{(b)}| \geq |T_{\text{obs}}|\}) / (B + 1)$.

Why it is exact: under H_0 a focal outcome depends only on its own treatment, which is fixed, so every re-draw leaves Y_F unchanged and the null distribution of T is known, with no model. Unlike Lecture 2's sharp null, own treatment may matter, so only focal outcomes can be filled in.

Peer effects: contagion or homophily?

Q3.3 — **peer influence**: does peers' behaviour change a patient's outcome?


New target: the effect of peers' *behaviour* W (e.g. adherence), not of their treatment A .



- **Challenge in observational data**: W is not randomized; similar people connect (**homophily**), so shared U mimics a peer effect.
- Ogburn et al. (2024): Framingham's obesity "contagion" vanishes once network structure is modelled.

Forastiere et al. (2021); Ogburn et al. (2024)

Randomize upstream of behaviour: mechanism and encouragement designs



- Target: the effect of peers' **behaviour** W (e.g. adherence) on a unit's Y . Red = 1, grey = 0; squares are randomized.
- (A) **Mechanism design**: randomly switch on or off the channel by which a unit sees peers' behaviour.
- (B) **Peer encouragement**: randomly nudge peers (Z); Z shifts W and reaches Y only through W , so it is an instrument for W .

From Eckles et al. (2016, Fig. 1); ego clusters at LinkedIn: Saint-Jacques et al. (2019); at scale: Karrer et al. (2021).

Short break

After the break: markets, then research questions

- **Part 4 — marketplace and switchback experiments:** who competes for a shared resource, and how long do treatment effects last?
- **Part 5 — research under interference:** what can we do when connections, exposure summaries, or market responses are uncertain?

Marketplace and switchback experiments

Recall the target: the global treatment effect

Part 4 — markets and time: when customers share supply, or one market is tracked over time, what does an A/B test estimate?

Roadmap: **Q4.1 the target** → Q4.2 which side to randomize → Q4.3 measure the bias → Q4.4 switch over time.

Same target as Parts 1 and 3: **global treatment effect** (GTE), everyone treated vs. no one. A one-sided A/B test instead treats a fraction p of one market and compares its two arms, $\tau_{AB}(p)$.


$$\tau_{\text{GTE}} = \mathbb{E} \bar{Y}(\text{all}) - \mathbb{E} \bar{Y}(\text{none}) \quad \text{vs.} \quad \tau_{AB}(p) = \mathbb{E}(\bar{Y}_T - \bar{Y}_C)$$

Under interference, a one-sided A/B test answers “treated vs. control in a mixed market”, not “everyone vs. no one”; more data does not close the gap, only design can.

GTE: Johari et al. (2022b); toy market constructed, 2,000 replications.

A two-sided market: interference runs through scarce supply

The platform shows some customers a **lower booking fee**. Customers book **listings**: the items on the supply side, such as homes on Airbnb.




- A treated customer who books a home takes it away from a control customer.
- So the control arm looks worse than in an untreated market: the two arms are not independent.
- Randomizing listings instead reverses the roles: treated listings take customers from control listings.

Setting and names: Johari et al. (2022b).

Randomize one side: the scarce side decides which test is biased

Q4.2 — which side: customers or listings?

randomize	who takes from whom	overstates when
customers, customer-side (CR)	treated customers take listings from control customers	listings are scarce
listings, listing-side (LR)	treated listings take customers from control listings	customers are scarce



Johari et al. (2022b); eBay buyer/seller tests: Blake and Coey (2014); Li et al. (2022)

Randomize both sides: treat a pair only if both sides are treated

	listing T	listing C
customer T	lower fee	usual fee
customer C	usual fee	usual fee

only the green cell receives the intervention

- **Two-sided randomization** (TSR): randomize customers *and* listings; a booking gets the lower fee only if both its customer and its listing are treated.
- Treating all listings gives CR back; treating all customers gives LR. TSR sits between them and is robust whichever side is scarce.
- Similar to Part 2's two-stage (nested) design: the four cells separate direct from spillover. Here, the two randomizations are on two sides, customers and listings, and the treated unit is the pair (a **multiple randomization design**).

Two-sided randomization: Johari et al. (2022b); multiple randomization designs: Bajari et al. (2023). Estimators: next slide (optional).

Technical: the market estimators (optional)

Notation: customer i treated with probability p_C , listing j with probability p_L ; Q_{ab} is the booking rate of customers in arm a with listings in arm b (1 = treated, 0 = control).


customer-side (CR), $p_L = 1$	listing-side (LR), $p_C = 1$	two-sided (TSR)
$\hat{\tau}^{\text{CR}} = \frac{Q_{11}}{p_C} - \frac{Q_{01}}{1 - p_C}$	$\hat{\tau}^{\text{LR}} = \frac{Q_{11}}{p_L} - \frac{Q_{10}}{1 - p_L}$	$\hat{\tau}^{\text{TSR}} = \frac{Q_{11}}{p_C p_L} - \frac{Q_{01} + Q_{10} + Q_{00}}{1 - p_C p_L}$

- In Johari et al.'s model, customers arrive at rate λ and booked listings free up at rate μ ; λ/μ is the “customers per available listing” axis of the earlier slide.
- CR is unbiased as $\lambda/\mu \rightarrow 0$ and LR as $\lambda/\mu \rightarrow \infty$; TSR with p_C, p_L chosen from λ/μ is unbiased in both limits.
- Bias corrections for TSR use the three grey cells: less bias, more variance.

Johari et al. (2022b), Thms 3–4, Prop. 4, eq. 21 and Corollary 1.

Measure the bias by design: a meta-experiment on Airbnb

Q4.3 — measure the bias: run two designs at once and compare



$$\widehat{\text{bias}} = \hat{\tau}_{\text{listing}} - \hat{\tau}_{\text{cluster}} = -0.068, \text{ about } 20\% \text{ of } -0.345$$

- A guest-fee change, March 2019, 2.6 million listings and grouped into clusters: 25% randomized listing by listing, 75% cluster by cluster.
- Both designs ran in the same weeks, so the gap is interference bias: a lower bound, since guests also switch between clusters.
- Part 3's cluster randomization, used to measure the bias without a model of the market.

Holtz et al. (2025), Tables 2–3 (published version; bars show 95% intervals). The 2020 arXiv version analyses different data.

Switchbacks: the exposure is the recent assignment path

Q4.4 — over time: one market, switched between arms in blocks of periods




$$\tau_m = \frac{1}{T-m} \sum_{t>m} \{Y_t(\mathbf{1}_{m+1}) - Y_t(\mathbf{0}_{m+1})\}$$
$$\hat{\tau}_m = \frac{1}{T-m} \sum_{t>m} Y_t \left\{ \frac{\mathbf{1}(W_{t-m:t} = \mathbf{1})}{\mathbb{P}(W_{t-m:t} = \mathbf{1})} - \frac{\mathbf{1}(W_{t-m:t} = \mathbf{0})}{\mathbb{P}(W_{t-m:t} = \mathbf{0})} \right\}$$

- **No anticipation** and **m -carryover**: Y_t depends only on $W_{t-m:t}$. Fair coin per block: $\mathbb{P}(W_{t-m:t} = \mathbf{1}) = 2^{-J}$, J = blocks touched; $\hat{\tau}_m$ is unbiased.
- Part 2's HT estimator with exposure $d_t = W_{t-m:t}$; carryover is Lecture 6's treatment history.

Bojinov et al. (2023), Assumptions 1–2, eq. (1), eq. (4), Proposition 1.

Where to switch: the optimal design for a known carryover (optional)



$T = nm, n \geq 4$: switch at $\{1, 2m + 1, 3m + 1, \dots, (n - 2)m + 1\}$, fair coins

- Minimax over bounded outcomes: blocks of m , but $2m$ at the ends, so end periods can be fully exposed.
- Unknown m : assuming too large an m keeps $\hat{\tau}$ unbiased; too small biases it.
- Carryover that fades but never ends: drop a **burn-in** at each block start; error $\approx (\log T / T)^{1/2}$, not $T^{-1/3}$.

Bojinov et al. (2023), Theorems 1–2 and Sections 4.3–4.4; fading carryover: Hu and Wager (2022) (preprint, version 5).

Research questions under interference

What can go wrong, and what researchers try

- ▶ **Q5.1 Incomplete exposure summary:** units outside our exposure mapping may also interfere.
- **Q5.2 Unknown connections:** we do not observe who is connected to whom.
- **Q5.3 Which uncertainty:** our intervals treat the network as fixed.
- **Q5.4 Whom to treat:** Lecture 6's treatment rules assume no spillovers.
- **Q5.5 Changing prices:** treating some units can move prices for everyone.

Q5.1: what if the exposure summary is incomplete?

Q5.1 — exposure summary: what if units outside our exposure mapping also interfere?



influence decays with network distance

- **Core idea:** replace “only these units interfere” with “far units interfere little”.
- The faster influence decays, the more accurate the estimates; a normal approximation gives valid intervals (Leung, 2022).
- Limit: the decay rate cannot be checked from one network, and the target effect still depends on the design.

Leung (2022).


What can go wrong, and what researchers try

- Q5.1 Incomplete exposure summary: units outside our exposure mapping may also interfere.
→ let far units interfere a little, with influence decaying by network distance.
- ▶ Q5.2 **Unknown connections**: we do not observe who is connected to whom.
- Q5.3 Which uncertainty: our intervals treat the network as fixed.
- Q5.4 Whom to treat: Lecture 6's treatment rules assume no spillovers.
- Q5.5 Changing prices: treating some units can move prices for everyone.

Q5.2a: without the network, the contrast gives the direct effect

Q5.2 — unknown connections: without the network, what does treated versus control estimate?


Example: $Y = b + 1 \cdot \text{own treatment} + 0.5 \cdot \# \text{treated neighbours}$, four neighbours each.



treated unit

$$Y = b + 1 + 0.5 \cdot 2$$


—



control unit

$$Y = b + 0 + 0.5 \cdot 2$$

contrast = 1
global = 3



everyone treated: never observed

$$Y = b + 1 + 0.5 \cdot 4$$

- **Why:** treated and control units have the same treated neighbours on average, so the spillover cancels, leaving the **direct** effect, 1.
- **Core idea:** no network needed, only limited interference (next slide).
- Limit: the **global** effect (GTE, Part 4) needs “everyone treated”, which is never observed.

Sävje et al. (2021).

Q5.2a: which estimators still work, and what is “small”?

Q5.2 — unknown connections: which estimators survive, and how much interference is too much?

Estimator	Without the network
difference in means	works: targets the EATE
Horvitz–Thompson/Hájek, $1/\mathbb{P}(A_i = a)$: own coin only	works: targets the EATE
Horvitz–Thompson/Hájek, $1/\mathbb{P}(D_i = d)$, Part 2: needs the neighbours	fails: not computable


- **EATE** (Part 2): $\frac{1}{n} \sum_i \mathbb{E}\{Y_i(1, \mathbf{A}_{-i}) - Y_i(0, \mathbf{A}_{-i})\}$, $\mathbf{A}_{-i}$ from the design; here 1, not the GTE 3.
- **How much interference:** d_i = number of units whose treatment can change Y_i (itself and neighbours, here $1 + 4 = 5$); d_{avg} is its average. It counts *who* interferes, not *how strongly*.
- **“Small”** means $d_{\text{avg}}/n \rightarrow 0$; the error shrinks like $\sqrt{d_{\text{avg}}/n}$, though the usual standard errors need not be valid (Sävje et al., 2021):

d_{avg} fixed	sparse network, few neighbours each	error $1/\sqrt{n}$, the usual rate
$d_{\text{avg}} \propto \sqrt{n}$	neighbourhoods widen with n	error $n^{-1/4}$, slower
$d_{\text{avg}} \propto n$	everyone affects everyone, as in a market	error does not shrink

Q5.2b: recover the global effect by varying the treatment probability

Q5.2 — unknown connections: can we reach “everyone treated” without seeing the connections?

Same example as Q5.2a, now treating each unit with probability p .



- $4p$ treated neighbours on average, so $\mathbb{E}[Y \mid A = 1] = b + p + 0.5 \cdot 4p = b + 3p$: **no network needed**.
- Fit a line at two p ; the rise from 0 to 1 is the **GTE**, 3.
- **Core idea:** a degree- β response gives a degree- β polynomial in p , fixed by $\beta + 1$ values of p .
- Limit: the model replaces the network, and the extrapolation relies on it.

Yu et al. (2022), who call the GTE the total treatment effect; staggered rollouts: Cortez et al. (2022).

Q5.2c: two answers to an unknown network

Q5.2 — unknown connections: which answer should we choose?

	Settle for less: Sävje et al. (2021)	Buy a model: Yu et al. (2022)
Target	direct effect, averaged over the design (1)	GTE, everyone versus no one (3)
Assumption	limited interference, $d_{\text{avg}}/n \rightarrow 0$	polynomial response to treated neighbours
Design	any single randomization	several p , or a staggered rollout
Estimator	difference in means; weighting	fit a polynomial in p ; extrapolate to 0 and 1


Both drop the network: one weakens the target, the other adds a model.

What can go wrong, and what researchers try

- **Q5.1** Incomplete exposure summary: units outside our exposure mapping may also interfere.
→ let far units interfere a little, with influence decaying by network distance.
- **Q5.2** Unknown connections: we do not observe who is connected to whom.
→ settle for the direct effect, or add a response model and vary the treatment probability.
- ▶ **Q5.3 Which uncertainty:** our intervals treat the network as fixed.
- **Q5.4** Whom to treat: Lecture 6's treatment rules assume no spillovers.
- **Q5.5** Changing prices: treating some units can move prices for everyone.

Q5.3: what randomness should our uncertainty describe?

Q5.3 — which uncertainty: should our interval treat the network as fixed, or as one random draw?



- **Core idea:** model the network as a random draw; the uncertainty then also depends on the network model.
- Average direct and indirect effects can be defined so that they stay meaningful under general interference.
- It depends on: do we want uncertainty for this network, or for similar networks?


Li and Wager (2022); Hu et al. (2022)

What can go wrong, and what researchers try

- **Q5.1** Incomplete exposure summary: units outside our exposure mapping may also interfere.
→ let far units interfere a little, with influence decaying by network distance.
- **Q5.2** Unknown connections: we do not observe who is connected to whom.
→ settle for the direct effect, or add a response model and vary the treatment probability.
- **Q5.3** Which uncertainty: our intervals treat the network as fixed.
→ model the network as a random draw; say which population the interval is for.
- ▶ **Q5.4 Whom to treat:** Lecture 6's treatment rules assume no spillovers.
- **Q5.5** Changing prices: treating some units can move prices for everyone.

Q5.4: whom should we treat when benefits spread?

Q5.4 — whom to treat: how do we choose whom to treat when the benefit also reaches neighbours?



1. **Experiment** (left): randomize with known probabilities; record own and neighbours' treatments.
2. **Learn the rule:** estimate each rule's welfare $W(\pi) = \frac{1}{n} \sum_i \mathbb{E}[Y_i(\pi(X_i), \pi(X_{N_i}))]$, the mean outcome if everyone follows π , where $\pi(X_{N_i})$ are the neighbours' treatments; pick the best π , as in Lecture 6.
3. **Apply it** (right) to the whole network, observed at the time of treatment.


From Viviano (2025, Fig. 1).

What can go wrong, and what researchers try

- **Q5.1** Incomplete exposure summary: units outside our exposure mapping may also interfere.
→ let far units interfere a little, with influence decaying by network distance.
- **Q5.2** Unknown connections: we do not observe who is connected to whom.
→ settle for the direct effect, or add a response model and vary the treatment probability.
- **Q5.3** Which uncertainty: our intervals treat the network as fixed.
→ model the network as a random draw; say which population the interval is for.
- **Q5.4** Whom to treat: Lecture 6's treatment rules assume no spillovers.
→ choose whom to treat counting the benefits that reach neighbours.
- ▶ **Q5.5 Changing prices:** treating some units can move prices for everyone.

Q5.5: what if treatment changes market prices?

Q5.5 — changing prices: what if treating some units changes the price for everyone?




- Unit i : a household; A_i : cash transfer, randomized; $Y_i(a, p)$: eggs bought at village price p .
- In equilibrium, $\sum_i Y_i(A_i, p) = S(p)$ (eggs bought = eggs supplied), so $P = p^*(\mathbf{A})$ depends on *everyone's* treatment.
- **Exposure** = (own treatment, price); compare no one vs. everyone treated, each at its clearing price.
- Every unit affects every other, so $d_{\text{avg}} \propto n$: Q5.2's result and cluster designs do not help.

Not a sparse network but one shared price: market structure replaces the graph.


Munro et al. (2025).

Q5.5: correct the direct effect for the price change

Q5.5 — changing prices: how do we get the GTE when treating everyone changes the price?



(a) Estimated expected demand and supply under treatment and control. Prices are per egg in local currency.



1. **Experiment:** both arms pay the same price, so the contrast is the **direct** effect τ_{direct} at the experiment's price; too large (b).
2. **Slopes:** small random price changes reveal the demand slope $\beta_D < 0$ and the supply slope $\beta_S > 0$.
3. **Combine (a):** treating everyone raises demand by τ_{direct} , so the clearing price is higher by $\Delta p^* = \tau_{\text{direct}} / (\beta_S - \beta_D)$:

$$\tau_{\text{GTE}} = \tau_{\text{direct}} + \beta_D \Delta p^* = \tau_{\text{direct}} \cdot \frac{\beta_S}{\beta_S - \beta_D}.$$

Limit: linear approximation; poor if the two prices are far apart.


From Munro et al. (2025, Fig. 1) (their total effect = GTE); Wager and Xu (2021); Bright et al. (2025).

What can go wrong, and what researchers try

- **Q5.1 Incomplete exposure summary:** units outside our exposure mapping may also interfere.
→ let far units interfere a little, with influence decaying by network distance.
- **Q5.2 Unknown connections:** we do not observe who is connected to whom.
→ settle for the direct effect, or add a response model and vary the treatment probability.
- **Q5.3 Which uncertainty:** our intervals treat the network as fixed.
→ model the network as a random draw; say which population the interval is for.
- **Q5.4 Whom to treat:** Lecture 6's treatment rules assume no spillovers.
→ choose whom to treat counting the benefits that reach neighbours.
- **Q5.5 Changing prices:** treating some units can move prices for everyone.
→ correct the direct effect for the price change, using supply and demand slopes.

Summary

Interference is a design problem first



Define the effect, summarize exposure, and choose assignments that produce it often enough.

- Lecture 8, October 30: unmeasured confounding; similarity between connected people can create confounding.
- Reading: Aronow et al. (2021) for Half I; Larsen et al. (2024) for Half II.

References and further reading i

- Aronow, P. M., Eckles, D., Samii, C., and Zonszein, S. (2021). Spillover effects in experimental data. In Druckman, J. N. and Green, D. P., editors, *Advances in Experimental Political Science*, pages 289–319. Cambridge University Press, Cambridge.
- Aronow, P. M. and Samii, C. (2017). Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4):1912–1947.
- Athey, S., Chetty, R., Imbens, G. W., and Kang, H. (2026). The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. *The Review of Economic Studies*, 93(4):2284–2312.
- Athey, S., Eckles, D., and Imbens, G. W. (2018). Exact p-values for network interference. *Journal of the American Statistical Association*, 113(521):230–240.
- Bajari, P., Burdick, B., Imbens, G. W., Masoero, L., McQueen, J., Richardson, T. S., and Rosen, I. M. (2023). Experimental design in marketplaces. *Statistical Science*, 38(3):458–476.
- Basse, G. W., Feller, A., and Toulis, P. (2019). Randomization tests of causal effects under interference. *Biometrika*, 106(2):487–494.
- Blake, T. and Coey, D. (2014). Why marketplace experimentation is harder than it seems: The role of test-control interference. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 567–582.
- Bojinov, I., Rambachan, A., and Shephard, N. (2021). Panel experiments and dynamic causal effects: A finite population perspective. *Quantitative Economics*, 12(4):1171–1196.
- Bojinov, I., Simchi-Levi, D., and Zhao, J. (2023). Design and analysis of switchback experiments. *Management Science*, 69(7):3759–3777.
- Bright, I., Delarue, A., and Lobel, I. (2025). Reducing marketplace interference bias via shadow prices. *Management Science*, 71(8):7094–7112.
- Cortez, M., Eichhorn, M., and Yu, C. L. (2022). Staggered rollout designs enable causal inference under interference without network knowledge. In *Advances in Neural Information Processing Systems*, volume 35, pages 7437–7449.
- Cox, D. R. (1958). *Planning of Experiments*. Wiley, New York.
- Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. *The Quarterly Journal of Economics*, 128(2):531–580.
- Deng, A., Xu, Y., Kohavi, R., and Walker, T. (2013). Improving the sensitivity of online controlled experiments by utilizing pre-experiment data. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, pages 123–132.
- Eckles, D., Karrer, B., and Ugander, J. (2017). Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1):20150021.

References and further reading ii

- Eckles, D., Kizilcec, R. F., and Bakshy, E. (2016). Estimating peer effects in networks with peer encouragement designs. *Proceedings of the National Academy of Sciences*, 113(27):7316–7322.
- Fabijan, A., Gupchup, J., Gupta, S., Omhover, J., Qin, W., Vermeer, L., and Dmitriev, P. (2019). Diagnosing sample ratio mismatch in online controlled experiments: A taxonomy and rules of thumb for practitioners. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2156–2164.
- Forastiere, L., Airoidi, E. M., and Mealli, F. (2021). Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534):901–918.
- Holtz, D., Lobel, F., Lobel, R., Liskovich, I., and Aral, S. (2025). Reducing interference bias in online marketplace experiments using cluster randomization: Evidence from a pricing meta-experiment on Airbnb. *Management Science*, 71(1):390–406.
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080.
- Hu, Y., Li, S., and Wager, S. (2022). Average direct and indirect causal effects under interference. *Biometrika*, 109(4):1165–1172.
- Hu, Y. and Wager, S. (2022). Switchback experiments under geometric mixing. Preprint arXiv:2209.00197, version 5 (September 2026).
- Hudgens, M. G. and Halloran, M. E. (2008). Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842.
- Johari, R., Koomen, P., Pekelis, L., and Walsh, D. (2022a). Always valid inference: Continuous monitoring of A/B tests. *Operations Research*, 70(3):1806–1821.
- Johari, R., Li, H., Liskovich, I., and Weintraub, G. Y. (2022b). Experimental design in two-sided platforms: An analysis of bias. *Management Science*, 68(10):7069–7089.
- Karrer, B., Shi, L., Bhole, M., Goldman, M., Palmer, T., Gelman, C., Konutgan, M., and Sun, F. (2021). Network experimentation at scale. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 3106–3116.
- Kohavi, R., Longbotham, R., Sommerfield, D., and Henne, R. M. (2009). Controlled experiments on the web: Survey and practical guide. *Data Mining and Knowledge Discovery*, 18(1):140–181.
- Kohavi, R., Tang, D., and Xu, Y. (2020). *Trustworthy Online Controlled Experiments: A Practical Guide to A/B Testing*. Cambridge University Press, Cambridge.
- Larsen, N., Stallrich, J., Sengupta, S., Deng, A., Kohavi, R., and Stevens, N. T. (2024). Statistical challenges in online controlled experiments: A review of A/B testing methodology. *The American Statistician*, 78(2):135–149.
- Leung, M. P. (2022). Causal inference under approximate neighborhood interference. *Econometrica*, 90(1):267–293.
- Li, H., Zhao, G., Johari, R., and Weintraub, G. Y. (2022). Interference, bias, and variance in two-sided marketplace experimentation: Guidance for platforms. In *Proceedings of the ACM Web Conference 2022*, pages 182–192.

References and further reading iii

- Li, S. and Wager, S. (2022). Random graph asymptotics for treatment effect estimation under network interference. *The Annals of Statistics*, 50(4):2334–2358.
- Lin, W. (2013). Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics*, 7(1):295–318.
- Manski, C. F. (2013). Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23.
- Munro, E., Kuang, X., and Wager, S. (2025). Treatment effects in market equilibrium. *American Economic Review*, 115(10):3273–3321.
- Ogburn, E. L., Sofrygin, O., Díaz, I., and van der Laan, M. J. (2024). Causal inference for social network data. *Journal of the American Statistical Association*, 119(545):597–611.
- Paluck, E. L., Shepherd, H., and Aronow, P. M. (2016). Changing climates of conflict: A social network experiment in 56 schools. *Proceedings of the National Academy of Sciences*, 113(3):566–571.
- Perez-Heydrich, C., Hudgens, M. G., Halloran, M. E., Clemens, J. D., Ali, M., and Emch, M. E. (2014). Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70(3):731–741.
- Puelz, D., Basse, G., Feller, A., and Toulis, P. (2022). A graph-theoretic approach to randomization tests of causal effects under general interference. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):174–204.
- Rubin, D. B. (1980). Randomization analysis of experimental data: The Fisher randomization test—comment. *Journal of the American Statistical Association*, 75(371):591–593.
- Saint-Jacques, G., Varshney, M., Simpson, J., and Xu, Y. (2019). Using ego-clusters to measure network effects at LinkedIn.
- Sävje, F., Aronow, P. M., and Hudgens, M. G. (2021). Average treatment effects in the presence of unknown interference. *The Annals of Statistics*, 49(2):673–701.
- Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476):1398–1407.
- Tchetgen Tchetgen, E. J. and VanderWeele, T. J. (2012). On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75.
- Ugander, J., Karrer, B., Backstrom, L., and Kleinberg, J. (2013). Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 329–337.
- Viviano, D. (2025). Policy targeting under network interference. *The Review of Economic Studies*, 92(2):1257–1292.
- Wager, S. and Xu, K. (2021). Experimenting in equilibrium. *Management Science*, 67(11):6694–6715.
- Waudby-Smith, I., Wu, L., Ramdas, A., Karampatziakis, N., and Mineiro, P. (2024). Anytime-valid off-policy inference for contextual bandits. *ACM / IMS Journal of Data Science*, 1(3):1–42.
- Yu, C. L., Airolidi, E. M., Borgs, C., and Chayes, J. T. (2022). Estimating the total treatment effect in randomized experiments with unknown network structure. *Proceedings of the National Academy of Sciences*, 119(44):e2208975119.

Questions?