

Experimentation on Endogenous Graphs

Wenshuo Wang*, Edvard Bakhitov, Dominic Coey

Central Applied Science, Meta
wenshuowang@meta.com, edbakhitov@meta.com, coey@meta.com

Abstract

We study experimentation under endogenous network interference. Interference patterns are mediated by an endogenous graph, where edges can be formed or eliminated as a result of treatment. We show that conventional estimators are biased in these circumstances, and present a class of unbiased, consistent and asymptotically normal estimators of total treatment effects in the presence of such interference. We show via simulation that our estimator outperforms existing estimators in the literature. Our results apply both to bipartite experimentation, in which the units of analysis and measurement differ, and the standard network experimentation case, in which they are the same.

Extended version — <https://arxiv.org/abs/2410.09267>

1 Introduction

In many settings, the fundamental no-interference assumption common in experimental analysis (Cox 1958; Rubin 1980) is implausible, and there are strong grounds for believing that the treatment assigned to one unit may affect outcomes of other units. This is especially true when outcomes are in part determined by peoples’ interactions with each other, as is often the case in public health, education, voting, social media, and other social domains (Calvó-Armengol, Patacchini, and Zenou 2009; DiMaggio and Garip 2012; Halloran and Hudgens 2016; Taylor and Eckles 2018). A large literature has arisen developing the theory of estimation and inference for various causal quantities of interest under these conditions.

Typically, theoretical and empirical analysis proceeds assuming that the interference structure is known—that is, the researcher has prior knowledge of the set of units which might be affected by the treatment of any given unit. This is usually expressed in terms of a known “exposure mapping” or “effective treatment” function, which determines the level of treatment exposure a unit receives as a function of the treatment assignment vector to all units, or in terms of a “partial interference” or “neighborhood interference” assumption, according to which interference takes place only

within groups, where the groups are known ex-ante (Hudgens and Halloran 2008; Tchetgen and VanderWeele 2012; Manski 2013; Liu and Hudgens 2014; Aronow and Samii 2017; Forastiere et al. 2022). This is a natural starting point that greatly facilitates the development of treatment effect estimators and the characterization of their behavior.

In practice, the interference structure is rarely known (Egami 2021; Sävje 2024). Analysts often rely on past user interactions to construct proxy interference graphs (Aral, Muchnik, and Sundararajan 2009; Bakshy et al. 2012; Bond et al. 2012; Coppock, Guess, and Ternovski 2016; Harshaw et al. 2021; Karrer et al. 2021). These graphs—such as friendship networks in social media or engagement-based links in marketplaces—are typically crude approximations of the true, unobserved interference graph.¹ Complicating matters further, they may themselves be causally affected by the treatment of interest. The analyst is presented with an uncomfortable choice—either use a graph constructed only on the basis of pre-treatment data, which may omit relevant information about how treatments affect outcomes, or incorporate post-treatment data too, which can lead to complex biases in treatment effect estimators.

Our Contributions. Motivated by this gap between theory and practice, this paper takes a different approach, explicitly allowing for the observed graph to be determined by treatment assignment. We work in a general bipartite setting, in which units of analysis and randomization are allowed to differ. We propose the notion of an *endogenous* bipartite graph which may be treatment-dependent, and which contains information on interference from randomization to analysis units. Generalizing the outcome model of Harshaw et al. (2021) to this setting, we give an unbiased, consistent and asymptotically normal estimator of the total treatment effect, as well as a consistent testing procedure for the sharp null of no treatment effect. Our identification and estimation approach relies on a treatment-invariant “anchor” sub-graph, from which instrumental variable estimates of unit-level treatment responses can be obtained.

Our results apply directly to bipartite experiments, which are of growing interest in their own right. Units of analy-

*Corresponding author.
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹For example, social network users frequently interact with “unconnected nodes” by being recommended to join new groups or follow new users.

sis and randomization may differ because it is not practical or impossible to measure outcomes of interest at the level at which treatments are being assigned (Zigler and Papadogeorgou 2021). Alternatively, there may be no conceptual difficulty in defining outcomes at the randomization unit-level, but accounting for the bipartite structure is a natural way to allow for cross-unit interference. Causal inference problems with this structure are common in social networks, recommender systems, digital advertising, and multi-sided platforms broadly (Chawla, Hartline, and Nekipelov (2016); Gilotte et al. (2018); Pouget-Abadie et al. (2019); Harshaw et al. (2021); Nandy et al. (2021); Johari et al. (2022); Bajari et al. (2023); Shi et al. (2024)).

Our framework extends naturally to the standard unipartite setting, where the units of analysis and randomization are the same. There, we combine unbiased inverse-propensity weighted estimators of the direct effect, while estimating the indirect effect using the previously developed logic for bipartite estimators. Unbiasedness, consistency and asymptotic normality follow as corollaries to our results for bipartite graphs.

Related Work. Two subareas of network science are especially related to this paper: models of network formation and estimation of peer effects (see An (2011); Chandrasekhar (2016); Bramoullé, Djebbari, and Fortin (2020) for overviews). Particularly relevant is an emerging body of work that questions standard assumptions about interference structure and allows for misspecification of the interference graph (Aronow and Samii 2017; Eckles, Karrer, and Ugander 2017; Wang et al. 2020; Leung 2022; Sävje 2024).

Related work also tackles causal inference on unknown graphs (Basse and Airolidi 2018; Chin 2018; Egami 2021; Sävje, Aronow, and Hudgens 2021; Cortez, Eichhorn, and Yu 2022; Yu et al. 2022; Shirani and Bayati 2023; Halloran and Hudgens 2016). Like these papers, we are concerned with the implausibility of correctly specifying the interference structure. However, we focus on the case of *endogenous* misspecification: treatment effects propagate along observed edges, but those edges may themselves depend on treatment assignment realizations. Notable recent exceptions that address this setting include Comola and Prina (2021); Gao (2024); Ryu (2024). Unlike Comola and Prina (2021), we estimate the effect of treating everyone vs no one, the policy-relevant estimand. Ryu (2024) takes a non-instrumental variable-based identification approach in the unipartite setting, using conditional means for outcome and edge formation with parameters assumed to be shared across units. Most closely related is Gao (2024), which also develops an instrumental variable strategy based on a pre-treatment graph, proceeding from quite different modeling assumptions (a unipartite graph with undirected, unweighted edges, and an outcome model with parameters shared across units).

To the best of our knowledge, in addition to providing novel insights into identification and estimation of treatment effects on endogenous unipartite graphs, these are also the first available results on treatment effect estimation for endogenous bipartite graphs. Several important questions re-

main open, including consistent variance estimation for the total treatment effect estimator.

Paper Outline. The rest of the paper is organized as follows. Section 2 introduces the basic setup for endogenous bipartite interference graphs and discusses the edge endogeneity bias. Section 3 proposes several models of endogenous edge formation. Section 4 includes our main estimator, and for which we demonstrate unbiasedness, consistency and asymptotic normality under certain assumptions. We also propose a statistical test for the sharp null with asymptotic power equal to 1. Our theory is supported by simulations in Section 5, where we examine the biases of our estimator and two exposure reweighted linear estimators (Harshaw et al. 2021), and demonstrate the confidence interval coverage properties of our estimator based on a variance proxy. Section 6 extends the results to unipartite graphs, and Section 7 concludes. All proofs and additional details are included in the Supplementary material.

2 Setup

There are n_a analysis units $a \in \mathcal{A}$ and n_r randomization units $r \in \mathcal{R}$. The treatment assignment random vector is denoted by $\mathbf{T} = (T_1, T_2, \dots, T_{n_r})$. For each analysis unit a , the potential outcome function $Y_a : \{0, 1\}^{n_r} \rightarrow \mathbb{R}$ maps treatment assignment vectors to outcomes. Throughout, the object of interest is the total treatment effect (TTE), $\mu = \frac{1}{n_a} \sum_a [Y_a(\mathbf{1}) - Y_a(\mathbf{0})]$. All randomness comes from treatment assignment. In the usual bipartite experimentation setting, the graph is defined so that an edge exists between an analysis unit and a randomization unit if the potential outcomes of the analysis unit depend on the treatment of the randomization unit. The next definition formalizes this.

Definition 1. A bipartite interference graph is a triple $(\mathcal{A}, \mathcal{R}, E)$ of analysis units \mathcal{A} , randomization units \mathcal{R} , and an adjacency matrix E with representative element $e_{ar} \in \{0, 1\}$. For all $a \in \mathcal{A}$, potential outcomes satisfy $Y_a(\mathbf{T}) = Y_a(\mathbf{T}')$ for all \mathbf{T}, \mathbf{T}' with $T_r = T'_r$ for all $r \in \mathcal{R}$ such that $e_{ar} = 1$.

In this definition, the adjacency matrix E restricts how each analysis unit’s outcomes can vary in response to treatment assignments, and furthermore is invariant to \mathbf{T} . Beyond the trivial example of a fully connected graph, it is rarely plausible to assume that this knowledge is available in applications. We take a different approach, motivated by practical settings in which the analyst observes and wishes to make use of a bipartite graph which contains some important information about how treatments affect outcomes, but cannot plausibly satisfy the definition above. In particular, we allow for the bipartite graph to be determined by the treatment assignment.

Formally, potential outcomes for edges are a function of treatment assignments. Each pair $(a, r) \in \mathcal{A} \times \mathcal{R}$ has a corresponding edge potential outcome function $E_{ar} : \{0, 1\}^{n_r} \rightarrow \{0, 1\}$, i.e. specifying whether the corresponding edge exists as a function of the treatment assignment vector. Let $\mathcal{R}_a(\mathbf{T}) = \{r \in \mathcal{R} : E_{ar}(\mathbf{T}) = 1\}$ denote the randomization units a is connected to under the treatment \mathbf{T} . We assume that as long as a treatment assignment does not change

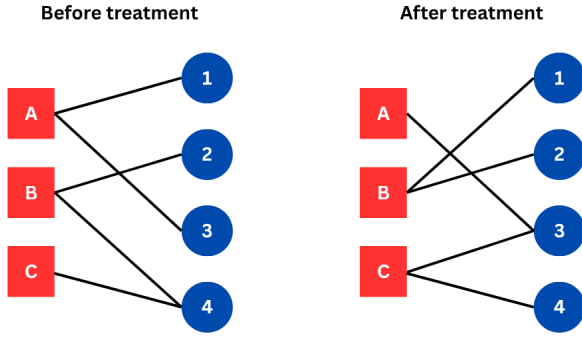


Figure 1: Endogenous graph example. Note that after the treatment is applied, edges (A-1) and (B-4) disappear, while new edges (B-1) and (C-3) emerge.

the randomization units that analysis unit a is connected to and does not change the treatment assignments to those randomization units, it cannot change a 's outcome. This gives the definition of an endogenous bipartite interference graph.

Definition 2. An endogenous bipartite interference graph is a triple $(\mathcal{A}, \mathcal{R}, E)$ of analysis units \mathcal{A} , randomization units \mathcal{R} , and an adjacency matrix E with representative element $E_{ar} : \{0, 1\}^{n_r} \rightarrow \{0, 1\}$. For all $a \in \mathcal{A}$, potential outcomes satisfy $Y_a(\mathbf{T}) = Y_a(\mathbf{T}')$ for all \mathbf{T}, \mathbf{T}' with $\mathcal{R}_a(\mathbf{T}) = \mathcal{R}_a(\mathbf{T}')$ and $T_r = T'_r$ for $r \in \mathcal{R}_a(\mathbf{T})$.

We assume throughout that for all $(a, r) \in \mathcal{A} \times \mathcal{R}$ the functions $E_{ar}(\cdot)$ are unknown, but their realizations given the treatment assignment $E_{ar}(\mathbf{T})$ are observed (see Figure 1 for illustration). Some observations on the relation between bipartite vs. unipartite and endogenous vs. exogenous interference graphs follow.

Remark 1. Definition 2 generalizes Definition 1, with the latter corresponding to the special case of the former where $\mathcal{R}_a(\mathbf{T})$ is constant.

Remark 2. Definition 2 also encompasses the case of endogenous unipartite experimentation, which is the special case of $\mathcal{A} = \mathcal{R}$. In this case, we typically assume $E_{aa} = 1$ for all a , so that each unit's treatment assignment is allowed to affect itself. The interference graph is allowed to be directed, in the sense that $E_{aa'}$ need not necessarily equal $E_{a'a}$ for $a \neq a'$.

Remark 3. Given any endogenous bipartite interference graph $(\mathcal{A}, \mathcal{R}, E)$, it is always possible to find *some* (exogenous) bipartite interference graph $(\mathcal{A}, \mathcal{R}, E')$ consistent with $(\mathcal{A}, \mathcal{R}, E)$. Trivially, we can take E' to be the fully connected bipartite graph $\mathcal{A} \times \mathcal{R}$. Somewhat less trivially, we can define E' as the set of edges (a, r) such that there exists a treatment realization \mathbf{T} satisfying $E_{ar}(\mathbf{T}) = 1$. In principle, conventional methods for causal inference on networks could be applied on the graph $(\mathcal{A}, \mathcal{R}, E')$. This observation is of limited value to practitioners, if E' is unobserved and they only observe a particular, treatment-dependent realization $E(\mathbf{T})$ of endogenous edges.

Graphs built from behavioral signals may not satisfy Definition 1 if they are likely to be causally affected by treat-

ments, but they may still satisfy Definition 2. For instance, a user-item graph in an online marketplace where edges represent shown items, may be causally affected by item-level treatments. Yet we might still assume that user outcomes remain unchanged if neither the set of shown items nor their treatments change. Similarly, in a Facebook Groups experiment, user outcomes may plausibly be unaffected if a user's group memberships and the treatments of those groups remain the same.

In what follows, we assume we have access to a pre-treatment graph, constructed from data before the experiment starts. It is denoted $G^{pre} = (\mathcal{A}, \mathcal{R}, E^{pre})$, with representative element $e_{ar}^{pre} \in \{0, 1\}$.² For conciseness, unless we wish to make their dependence on \mathbf{T} explicit, we write the realized metric outcomes $Y_a(\mathbf{T})$ as y_a and realized edge outcomes $E_{ar}(\mathbf{T})$ as e_{ar} . We now give examples of how naively applying existing bipartite estimators without accounting for the endogenous nature of the graph can bias conventional inference-aware estimators.

Examples of Edge Endogeneity Bias. Treating the realized graph $E(\mathbf{T})$ observed in the setting of Definition 2 as fixed can bias otherwise unbiased estimators, including the Horvitz-Thompson (HT) inverse propensity weighted estimator (Horvitz and Thompson (1952)). We assume that treatment assignments are independent Bernoulli random variables with probability p . The HT estimator of the TTE is

$$\hat{\mu}_{\text{HT}} = \frac{1}{n_a} \sum_{a=1}^{n_a} y_a \left[\frac{\prod_{r \in \mathcal{R}_a(\mathbf{T})} T_r}{p^{|\mathcal{R}_a(\mathbf{T})|}} - \frac{\prod_{r \in \mathcal{R}_a(\mathbf{T})} (1 - T_r)}{(1 - p)^{|\mathcal{R}_a(\mathbf{T})|}} \right],$$

where the term in the square brackets is defined to be zero if $\mathcal{R}_a(\mathbf{T}) = \emptyset$. The first term simply sums up outcomes y_a for the "fully treated" analysis units, i.e. those units a such that all their corresponding randomization units are treated ($\prod_{r \in \mathcal{R}_a(\mathbf{T})} T_r = 1$), and rescales by the factor $p^{-|\mathcal{R}_a(\mathbf{T})|}$ to account for the probability of that event. This gives the estimate of the average outcomes under $\mathbf{T} = \mathbf{1}$, and analogously the second term gives the estimate of average outcomes under $\mathbf{T} = \mathbf{0}$. By standard arguments (Aronow and Samii 2017), this is an unbiased estimate of the TTE under graph exogeneity, where $\mathcal{R}_a(\mathbf{T})$ does not vary with \mathbf{T} .

Under graph endogeneity, this unbiasedness property no longer holds. We illustrate mechanisms by which bias can arise in the following minimal examples. There is a single randomization unit in each, implying the HT estimator and the exposure reweighted linear (ERL) estimator of Harshaw et al. (2021) are identical, so these examples also demonstrate bias in the ERL estimator.

²Even though no-one is treated pre-experiment, we need not have $E^{pre} = E(\mathbf{0})$. The potential outcomes for the edges and the metric of interest, $(E(\mathbf{T}), Y_a(\mathbf{T}))$ should be interpreted as their values at some fixed time post-experiment start. So if edges form and disappear over time regardless of whether treatments are being applied, the pre-experiment edges E^{pre} and the counterfactual edges that would exist without any treatment as measured at that future time, $E(\mathbf{0})$, may differ.

Example 1. $\mathcal{R} = \{r_1\}$ and $\mathcal{A} = \{a_1\}$. Unit r_1 forms an edge with a_1 if and only if $T_1 = 1$. Outcomes for a_1 are $y \neq 0$ regardless of the treatment assignment, and $T_1 \sim \text{Bernoulli}(1/2)$. The HT estimator is $\frac{T_1 y}{1/2}$. This is y in expectation, which is biased for the TTE of zero. In this example the number of edges varies with the treatment assignment, but that is not necessary for bias to exist, as the next example shows.

Example 2. $\mathcal{R} = \{r_1\}$ and $\mathcal{A} = \{a_1, a_2\}$. Unit r_1 forms an edge with a_1 if $T_1 = 1$, and with a_2 otherwise. Outcomes for a_1 and a_2 are y_1 and y_2 regardless of the treatment assignment, and $T_1 \sim \text{Bernoulli}(1/2)$. The HT estimator is $\frac{T_1 y_1}{1/2} - \frac{(1-T_1)y_2}{1/2}$ which has expectation $y_1 - y_2$, and again is biased for the TTE of zero. In this example, edge formation induces correlation between treatment intensity and analysis unit potential outcomes, generating bias.

The preceding examples distinguish between randomization and analysis units and so relate to bipartite experimentation. Graph endogeneity can also bias treatment effect estimators in the case of unipartite network interference where this distinction does not exist, as the next example shows.

Example 3. $\mathcal{R} = \mathcal{A} = \{a_1, a_2\}$. Outcomes for a_1 and a_2 are y_1 and y_2 regardless of treatment assignment, and T_1, T_2 are independent $\text{Bernoulli}(1/2)$ draws. A directed edge runs from a_1 to a_2 if and only if $T_1 = 1$. The HT estimator is

$$\hat{\mu}_{\text{HT}} = \begin{cases} -y_1/2 - y_2/2 & \text{if } T_1 = T_2 = 0 \\ -y_1/2 + y_2/2 & \text{if } T_1 = 0, T_2 = 1 \\ y_1/2 & \text{if } T_1 = 1, T_2 = 0 \\ y_1/2 + y_2/4 & \text{if } T_1 = T_2 = 1, \end{cases}$$

with expectation $y_2/16$, and is biased for the TTE of zero.

3 Models of Edge Formation

Just as estimating total treatment effects under exogenous interference requires imposing some additional structure on the interference mechanism (Basse and Airolidi 2018), we will need to rule out arbitrary dependence of edges on treatment assignments. Let S_{ar} be a subset of \mathcal{R} satisfying $E_{ar}(\mathbf{T}) = E_{ar}(\mathbf{T}')$ for all \mathbf{T}, \mathbf{T}' such that $T_s = T'_s$ for all $s \in S_{ar}$. That is, whether the edge (a, r) exists depends only on the treatment assignments to the randomization units in S_{ar} . Examples of various edge formation assumptions follow.

- *Unrestricted edges.* $S_{ar} = \mathcal{R}$. Edges can depend arbitrarily on treatments.
- *Exogenous edges.* $S_{ar} = \emptyset$. Edges are unaffected by treatments, as per Definition 1.
- *r-driven edges.* $S_{ar} = \{r\}$. Edges depend only on the randomization unit's treatment status. This can be interpreted as a SUTVA (Stable Unit Treatment Value Assumption) assumption on edge formation: edge outcomes depend only on the treatment assignment to the randomization unit in question, not by treatment assignments to other units.

- *S_a -driven edges.* $S_{ar} = S_a \subset \mathcal{R}$. Analysis unit a 's edge formation depends only on treatment assignments to a subset of randomization units.
- *S_r -driven edges.* $S_{ar} = S_r \subset \mathcal{R}$. Randomization unit r 's edge formation depends only on treatment assignments to a subset of randomization units.

The following statistical tests can provide evidence on how appropriate the assumptions of exogenous or r -driven edges may be.

Detecting general departures from edge exogeneity. Define the null hypothesis as $H_0 : S_{ar} = \emptyset$ for all a, r , and consider the following test statistic based on the binomial likelihood for each analysis unit's number of treated edges:

$$W(\mathbf{T}, E(\mathbf{T})) = \sum_a \log \left[\left(\frac{\sum_r e_{ar}}{\sum_r e_{ar} T_r} \right) p^{\sum_r e_{ar} T_r} (1-p)^{\sum_r e_{ar} (1-T_r)} \right].$$

Under H_0 , $E(\mathbf{T})$ is independent of \mathbf{T} , and the conditional distribution of W given E can be simulated by resampling \mathbf{T} according to the experiment design. Following the usual randomization inference logic, we obtain a valid statistical test by rejecting the null if the value of W is larger than the upper α^{th} quantile of this null conditional distribution.

This test has power against a variety of departures from exogenous edges, including but not limited to r -driven edges. If treated randomization units disproportionately form new edges, this will lead to analysis units having more treated units than would be expected compared to the binomial benchmark. And if analysis units are more likely to form additional edges based on the treatment intensity of their existing edges, this will similarly lead to a departure from the binomial benchmark.³ It is however not powered to detect all conceivable departures from exogenous edges, and there are somewhat contrived examples of endogenous edges which will not be detected by this test.⁴

Detecting r -driven edges. If the alternative hypothesis is that edges only depend on the randomization unit's treatment status, a straightforward and valid procedure is to test whether treatments are generating more or fewer edges in aggregate with a simple difference-in-means t -test. Let $E_r = \sum_a (e_{ar} - e_{ar}^{\text{pre}})$ be the net number of edges added relative to the pre-experiment graph. We can then test the null hypothesis that $\mathbb{E}[E_r | T_r = 1] = \mathbb{E}[E_r | T_r = 0]$.

³To build intuition for the former case, imagine units form edges if and only if they are treated. Then for all analysis units that have any edges, they are all maximally treated. For the latter case, imagine analysis units will form an edge with a single pre-determined randomization unit, and then form another edge with another randomization unit if and only if the previous unit was treated. Then there will be no analysis units connected to a single treated randomization unit.

⁴Consider the case where there is a single analysis unit and two randomization units, and $e_{11} = \mathbf{1}(T_2 = 1)$. The treatment intensity for the analysis unit conditional on an edge existing is indeed binomial, despite the fact that whether an edge exists in the first place depends on the other randomization unit.

For TTE estimation, we focus exclusively on the case of r -driven edges. This strikes a balance between meaningfully enriching the standard exogenous bipartite graph model, and still imposing enough structure on the problem to be able to derive useful statements on treatment effect magnitudes. This choice is also motivated by practicality—while setting S_{ar} to be some strict superset of $\{r\}$ is more general, in applications there may be no clear principles available for constructing such sets. Note that we are relaxing the common assumption of the interference graph being fixed, and in this respect generalize, rather than specialize, existing work—the r -driven edges assumption is trivially satisfied for a fixed graph. Some restrictions on edges and outcomes are indeed necessary to make meaningful claims about treatment effects in this more general setting.

We opt for the r -driven model of edge endogeneity as the simplest non-trivial generalization of exogenous edges. Robustness under the violation of this assumption and relaxing the assumption to S_r -driven edges, where edge formation can depend on peers' treatments, are interesting future directions to explore. We wish to note that, while it is attractive to dispense with this assumption entirely, we see it as being analogous to the required restrictions on the interference structure in the exogenous graph case—without some structure, causal inference is impossible (see, e.g., Basse and Airolidi (2018)).

4 A Class of Unbiased TTE Estimators

Generalized Linear Exposure-Response Model. We specify the outcome model $y_a = \alpha_a + \beta_a x_a$, where $x_a = \sum_r T_r e_{ar} w_{ar}$ is the weighted number of treated edges a is connected to in the realized, post-treatment graph, and w_{ar} 's are known weights. The w_{ar} 's are a modeling choice, and we recommend the default choice of $w_{ar} = 1/(\# \text{ pre-treatment edges connected to } a)$. The researcher might choose to incorporate domain knowledge into these weights based on how they expect a node to be affected by others (e.g., a user may be more affected by friends with whom they message frequently, or a buyer may be more affected by a seller they have frequently bought from in the past). These weights play the same role as the weights in the outcome model of Harshaw et al. (2021). This outcome model satisfies the restrictions on potential outcomes stipulated in Definition 2, and is a generalization of the linear exposure-response model in Harshaw et al. (2021) to endogenous bipartite graphs. Outcomes are linear in the number of treated randomization units the analysis unit is exposed to, but the linear response is allowed to vary across analysis units. Treatments only affect outcomes through the observed edges, but the set of observed edges is potentially treatment dependent. By definition, the TTE is

$$\mu = \frac{1}{n_a} \sum_a [Y_a(\mathbf{1}) - Y_a(\mathbf{0})] = \frac{1}{n_a} \sum_a \mathcal{W}_a(\mathbf{1}) \beta_a, \quad (1)$$

where $\mathcal{W}_a(\mathbf{1}) = \sum_r w_{ar} E_{ar}(\mathbf{1})$.

Estimation Strategy. We now give a high-level overview of our approach to estimating the TTE. The two components of (1) are β_a , the incremental effect of treating an additional

randomization unit for analysis unit a , and $\mathcal{W}_a(\mathbf{1})$, the total number of edges a would have under a full treatment rollout. To construct an estimator of the TTE, we propose a class of unbiased estimators for each of these two components. We prove that the estimators for these two components are in addition uncorrelated with each other, so their product is an unbiased estimator for the summand of (1).

Considering β_a first, Harshaw et al. (2021) propose $y_a(x_a - \mathbb{E}x_a)/\text{Var}(x_a)$ as an unbiased estimator in the exogenous bipartite graph case. This estimator is no longer feasible because the expectation and variance of x_a are unknown, due to the endogenous nature of the graph edges. Moreover, as the examples in Section 2 show, naively applying this estimator assuming the realized graph is exogenous will lead to bias. To remedy this issue, we borrow from the literature on instrumental variables. Let $z_a^u = \sum_r T_r u_{ar}$ be a weighted exposure variable, where u_{ar} 's are fixed weights chosen by the researcher. If $\text{Cov}(x_a, z_a^u)$ is known and non-zero, we have a class of unbiased estimators for β_a ,

$$\left\{ \hat{\beta}_a^u = y_a(z_a^u - \mathbb{E}z_a^u)/\text{Cov}(x_a, z_a^u) \mid u : \mathcal{A} \times \mathcal{R} \rightarrow \mathbb{R} \right\}. \quad (2)$$

For $\mathcal{W}_a(\mathbf{1})$, let

$$\widehat{\mathcal{W}}_{ar}(\mathbf{1})^c = \frac{T_r w_{ar} (e_{ar} - c_{ar})}{p} + w_{ar} c_{ar}.$$

Under r -driven edges it follows that $\mathbb{E}[\widehat{\mathcal{W}}_{ar}(\mathbf{1})^c] = w_{ar} \mathbb{E}[e_{ar} \mid T_r = 1] = w_{ar} E_{ar}(\mathbf{1})$, and so a class of unbiased augmented inverse probability weighted estimators is

$$\left\{ \widehat{\mathcal{W}}_a(\mathbf{1})^c = \sum_r \widehat{\mathcal{W}}_{ar}(\mathbf{1})^c \mid c : \mathcal{A} \times \mathcal{R} \rightarrow \mathbb{R} \right\}. \quad (3)$$

Our estimation strategy involves giving conditions ensuring that $\text{Cov}(x_a, z_a^u)$ is known and non-zero, and that our estimator for β_a is uncorrelated with our estimator for $\mathcal{W}_a(\mathbf{1})$. The key assumption is the existence of an observed “anchor” subgraph, the edges of which would exist under any treatment. This suffices for unbiasedness of our TTE estimator. Consistency and asymptotic normality follow under additional assumptions on graph sparsity, which limit dependence across analysis units and allow the application of convergence theorems for dependent data.

Using Anchor Instruments. In general, though the experimental design and hence the joint distribution of \mathbf{T} is known, the covariance between the treatment intensity $x_a = \sum_r T_r e_{ar} w_{ar}$ and the instrument $z_a^u = \sum_r T_r u_{ar}$ is unknown. This is because it is a function of the unobserved values that the random variables e_{ar} would take on under counterfactual treatment assignments. This suggests a path forward for a particular class of instruments for which $\text{Cov}(x_a, z_a^u)$ can be calculated. If there were a known set of pre-experiment edges that would continue to exist regardless of treatment assignments, this would allow us to circumvent the difficulties posed by not observing e_{ar} under counterfactual treatments. To start formalizing these observations, we present the definition of an anchor subgraph.

Definition 3 (Anchor Subgraph). In Definition 2, we say $G \subset \mathcal{A} \times \mathcal{R}$ is an anchor subgraph, if $E_{ar}(\mathbf{T}) = 1$ for any $\mathbf{T} \in \{0, 1\}^{n_r}$, $(a, r) \in G$.

An anchor subgraph is a set of edges that would exist under any treatment.⁵ A subgraph in the pre-experiment data may plausibly satisfy this property, in which case we can use it to construct an instrument for the endogenous x_a 's.

As a practical example, consider Facebook groups, where the bipartite graph links users (analysis units) to groups (randomization units). User outcomes depend on the treatment of connected groups, and those connections may themselves be affected by the treatment. To construct an anchor subgraph, we can identify groups a user is unlikely to leave—e.g., using a predictive model trained on pre-treatment data. Similarly, in unipartite settings, persistent friendships may serve as anchor edges. These assumptions can be validated empirically—for any candidate anchor subgraph, we can estimate what fraction of edges disappear under treatment. To continue the Facebook groups example, less than 0.5% of all group memberships disappear over two weeks. By contrast, experimental treatments may cause a larger number of new memberships to form over the same period, with a recent test generating a rate of edge formation of 0.75%.

The following theorem gives conditions for the TTE estimator we propose to be unbiased. Assumption (a) of r -driven edges restricts how treatment assignment affects edge creation or deletion and is used throughout the proof. Assumption (b) describes the experimental design. Assumption (c) requires that the variation in the instrument z_a^u is entirely driven by treatment assignments to the known anchor subgraph G . Assumption (d) ensures $\text{Cov}(\hat{\beta}_a^u, \widehat{\mathcal{W}}_a(\mathbf{1})^c) = 0$. Assumptions (e) and (f) guarantee that the instrument z_a^u has some non-zero correlation with the treatment intensity x_a .

Theorem 1. *In the setting of Definition 2, suppose G is a known anchor subgraph. If (a) edges are r -driven; (b) the treatment assignments are independent Bernoulli random variables with probability p ; (c) $\{(a, r) \mid u_{ar} \neq 0\} \subset G$; (d) $c_{ar} = \mathbb{I}((a, r) \in G)$; (e) $w_{ar} \neq 0$ for all $(a, r) \in G$; and (f) $|\{r \mid u_{ar} \neq 0\}| > 0$ for all a , then*

$$\hat{\mu}^{u,c} = \frac{1}{n_a} \sum_a \hat{\beta}_a^u \cdot \widehat{\mathcal{W}}_a(\mathbf{1})^c \quad (4)$$

is an unbiased estimator of the TTE in (1).

Let $R_a = \{r \in \mathcal{R} \mid E_{ar}(\mathbf{1}) = 1\}$ be the set of randomization units connected to a under treatment, and $A_r = \{a \in \mathcal{A} \mid E_{ar}(\mathbf{1}) = 1\}$ the set of analysis units connected to r under treatment. Let $V_a = \{r \in \mathcal{R} \mid u_{ar} \neq 0\}$ be the set of randomization units connected to a in the anchor

⁵For our Theorem 1 and Theorem 2, we can relax this assumption to requiring only $E_{ar}(\mathbf{1}) = 1$ under r -driven edges. Thus the existence of a known set of edges in the pre-experiment graph which satisfy the anchor subgraph property can be seen as a milder version of a “compliance”-type assumption familiar from the instrumental variables literature (Angrist, Imbens, and Rubin 1996), which in this setting would amount to requiring that all pre-treatment edges exist after the treatment has been applied.

subgraph. To see the time complexity of computing (4), we note that

$$\hat{\beta}_a^u = \frac{y_a \sum_r u_{ar}(T_r - p)}{p(1-p) \sum_r w_{ar} u_{ar}}$$

requires $\Theta(|V_a|)$ calculations, and $\widehat{\mathcal{W}}_a(\mathbf{1})^c$ requires $\Theta(|R_a|)$ calculations. Thus, the total time complexity is $\Theta(\sum_a |R_a|)$.

To characterize the limiting behavior of the estimator proposed in Theorem 1, we consider an asymptotic regime in which the bipartite graph increases in size. If the graph is sufficiently sparse in this limit—i.e. the maximum degree of the analysis units and randomization units grows sufficiently slowly relative to the number of analysis units—then we can apply convergence results for dependent data, which allow us to derive consistency and asymptotic normality. Asymptotic normality relies on an application of Stein’s method (Stein 1986; Ross 2011), which has proven to be broadly applicable in demonstrating asymptotic normality in network settings (Chin 2018; Harshaw et al. 2021). Theorem 2 below summarizes these results. Assumptions (a), (b) and (c) ensure the instrument is positively correlated with the treatment exposure level x_a . Assumption (c) limits the total weight of randomization units connected to the same user and the discrepancy between those weights, and assumption (d) bounds outcomes.

Theorem 2. *We maintain the assumptions in Theorem 1, and further assume (a) $|V_a| > 0$; (b) $u_{ar} \geq 0$; (c) $W_l/|R_a| \leq |w_{ar}| \leq W_h/|R_a|$ for $r \in R_a$; and (d) $|y_a| \leq M$, where $p, 1-p, W_l, W_h, M$ are constants bounded away from 0 and ∞ . Then, letting $d_A = \max_{a \in \mathcal{A}} |R_a|$ and $d_{\mathcal{R}} = \max_{r \in \mathcal{R}} |A_r|$,*

1. $\text{Var}(\hat{\mu}^{u,c}) = O(d_A^3 d_{\mathcal{R}}/n_a)$, and thus $\hat{\mu}^{u,c}$ is consistent if $d_A^3 d_{\mathcal{R}}/n_a \rightarrow 0$;
2. if $n_a \text{Var}(\hat{\mu}^{u,c})$ is bounded away from zero and $d_A^{10} d_{\mathcal{R}}^4/n_a \rightarrow 0$, then $(\hat{\mu}^{u,c} - \mu)/\sqrt{\text{Var}(\hat{\mu}^{u,c})}$ converges in distribution to $\mathcal{N}(0, 1)$.

With asymptotic normality and a variance estimator, we can construct confidence intervals. Section 4 provides a modified estimator with tractable variance that can be used as a proxy for uncertainty quantification for the original estimator. Section 5 shows this proxy variance yields confidence intervals with close-to-nominal coverage in simulations.

Statistical Testing under the Sharp Null. Based on the estimators of (2), we can construct a test for the sharp null $H_0^{\text{sharp}}: \beta_a = 0$ for all a . Under H_0^{sharp} , $y_a = \alpha_a$ is independent of \mathbf{T} . Define a modified variant of (4),

$$\begin{aligned} \tilde{\mu}^{u,c} &= \frac{1}{n_a} \sum_a \hat{\beta}_a^u \sum_r (w_{ar} c_{ar}) \\ &= \frac{1}{n_a} \sum_a \frac{y_a (z_a^u(\mathbf{T}) - \mathbb{E} z_a^u)}{\text{Cov}(x_a, z_a^u)} \sum_r (w_{ar} c_{ar}), \end{aligned} \quad (5)$$

where the dependence of z_a^u on \mathbf{T} is explicit. Under H_0^{sharp} and the assumptions of Theorem 1, all randomness in $\tilde{\mu}^{u,c}$ is from $z_a^u(\mathbf{T})$, the distribution of which is known.

Theorem 3. We maintain the assumptions in Theorem 2. Define the sharp null $H_0^{\text{sharp}}: \beta_a = 0$ for all a . The H_0^{sharp} -null distribution of $|\tilde{\mu}^{u,c}|$ in (5) conditional on $\{y_a : a \in \mathcal{A}\}$ is known. Denote c_α as its α -upper quantile, $\alpha \in (0, 1)$. Rejecting if $|\tilde{\mu}^{u,c}| \geq c_\alpha$ is a valid size- α test for H_0^{sharp} , and the power goes to 1 if $d_{\mathcal{A}}^3 d_{\mathcal{R}}/n_a \rightarrow 0$ and

$$\liminf \left| \frac{1}{n_a} \sum_a \beta_a \sum_r (w_{ar} c_{ar}) \right| > 0.$$

The time complexity to calculate (5) is $\Theta(\sum_a |V_a|)$. The exact c_α is expensive to compute, but noting that $E[\tilde{\mu}^{u,c} | \{y_a : a \in \mathcal{A}\}, H_0^{\text{sharp}}] = 0$, we can obtain a more conservative cutoff based on $\text{Var}(\tilde{\mu}^{u,c} | \{y_a : a \in \mathcal{A}\}, H_0^{\text{sharp}})$, the time complexity to calculate which is also $\Theta(\sum_a |V_a|)$, because we have

$$\begin{aligned} \text{Var}(\tilde{\mu}^{u,c} | \{y_a : a \in \mathcal{A}\}, H_0^{\text{sharp}}) \\ = \frac{1}{p(1-p)} \sum_r \left(\sum_a \frac{y_a C_a u_{ar}}{n_a U_a} \right)^2, \end{aligned}$$

where $C_a = \sum_r w_{ar} c_{ar}$ and $U_a = \sum_r w_{ar} u_{ar}$.

5 Simulations

We let $n_a = 2000$ and $n_r = 400$. We generate pre-treatment edges e_{ar}^{pre} s as independent Bern(0.1) random variables. Our algorithm takes the anchor edges to be the same as the pre-treatment edges. The edge outcome function is r -driven and defined as $E_{ar}(T_r) = e_{ar}^{\text{pre}} + T_r(1 - 2e_{ar}^{\text{pre}})\text{Bern}(p_{\text{anchor}})$ if (a, r) is anchor, and $e_{ar}^{\text{pre}} + T_r(1 - 2e_{ar}^{\text{pre}})\text{Bern}(p_{\text{non-anchor}})$ otherwise. That is, the edge outcome is the same as the pre-treatment edge if no treatment is applied to the randomization unit; otherwise, it will flip to the opposite of the pre-treatment edge with probability p_{anchor} if the pre-treatment edge exists, and $p_{\text{non-anchor}}$ if it does not.

The α_a 's are sampled from independent Unif([0, 1]) variables, and the true treatment effect is set to be $\beta_a = \beta$ for all $a \in \mathcal{A}$. We run 1000 Monte Carlo simulations for various values of β , p_{anchor} and $p_{\text{non-anchor}}$ and compare our anchor-based estimator (Anchor) with the ERL estimator (Harshaw et al. 2021) based on (a) pre-treatment graph (ERL_{pre}) and (b) post-treatment graph (ERL_{post}). The ERL estimators are

$$\hat{\mu}^{\text{ERL-pre}} = \frac{1}{n_a} \sum_a \frac{y_a(x_a^{\text{pre}}(\mathbf{T}) - p)}{p(1-p)/(\sum_{r \in \mathcal{R}} e_{ar}^{\text{pre}})}$$

and

$$\hat{\mu}^{\text{ERL-post}} = \frac{1}{n_a} \sum_a \frac{y_a(x_a^{\text{post}}(\mathbf{T}) - p)}{p(1-p)/(\sum_{r \in \mathcal{R}} e_{ar}^{\text{post}})},$$

where $x_a^{\text{pre}}(\mathbf{T}) = \frac{\sum_{r \in \mathcal{R}} e_{ar}^{\text{pre}} T_r}{\sum_{r \in \mathcal{R}} e_{ar}^{\text{pre}}}$ and $x_a^{\text{post}}(\mathbf{T}) = \frac{\sum_{r \in \mathcal{R}} e_{ar}^{\text{post}} T_r}{\sum_{r \in \mathcal{R}} e_{ar}^{\text{post}}}$.

Here, e_{ar}^{pre} and e_{ar}^{post} are the pre-treatment and post-treatment edges between a and r , respectively. We set $\beta = 2$, $p_{\text{anchor}} = 2\%$ and $p_{\text{non-anchor}} = 0.2\%$ when they are not being varied.

Bias Comparison. Figure 2 shows the biases of the estimators under different parameter values. The anchor estimator results are consistent with unbiasedness with respect to different values of treatment effect (β) and edge formation intensity ($p_{\text{non-anchor}}$), as expected given our theoretical results. Using the post-treatment graph (ERL_{post}) naively can lead to severe overestimation due to the positive correlation between the marginal effect of an edge and the number of treated edges. Estimates based on the pre-treatment graph are still biased, but in the opposite direction, and the magnitude of the bias is much smaller as it stems solely from graph misspecification. Unbiasedness of the anchor estimator is achieved at the cost of higher variance, and the variance grows with $p_{\text{non-anchor}}$; with more edges being formed in response to treatment, we expect larger metric fluctuations between test and control contributed by $\widehat{\mathcal{W}}_{ar}(\mathbf{1})^c$, and hence, higher variance of the anchor estimator.

The anchor estimator is quite robust to violations of the anchor subgraph assumption, as we increase p_{anchor} . In Figure 2, we can see that its bias increases with p_{anchor} , but stays relatively small compared to the magnitude of the treatment effect, which is 2. ERL_{pre} is not directly affected by p_{anchor} and is stable. On the other hand, the bias of ERL_{post} changes its sign with the number of anchor edges disappearing. This is most likely due to the correlation between the marginal effect (which is positive) and the number of edges (that starts to diminish) becoming negative leading to underestimation of the true effect.

Uncertainty Quantification. In order to provide a confidence interval (CI) for the anchor estimator, we appeal to asymptotic normality and use $\text{Var}(\tilde{\mu}^{u,c})$ under the sharp null as a proxy for the anchor estimator $\text{Var}(\hat{\mu}^{u,c})$, where $\tilde{\mu}^{u,c}$ is the estimator (5) defined in Section 4. We observe in Figure 3 that this approach delivers close to the nominal level coverage under different values of treatment effect (β) and anchor subgraph violation (p_{anchor}), with only modest under-coverage. By contrast, coverage is low when the edge formation density ($p_{\text{non-anchor}}$) is high, which is not surprising since the variance proxy is for an estimator that ignores the direct effect of newly formed edges.

6 Endogenous Unipartite Graphs

An endogenous unipartite interference graph is the special case of Definition 2, where the analysis and randomization units are the same: $\mathcal{A} = \mathcal{R}$. This definition allows for *directed* interference, in the sense that E_{ar} may not equal E_{ra} . We specify the outcome model for all a as $y_a = \alpha_a + \beta_a x_a + \gamma_a T_a$, where $x_a = \sum_r T_r e_{ar} w_{ar}$ and we impose $w_{aa} = 0$. Thus, $\beta_a x_a$ is the indirect effect on a 's outcome coming from the treatment of others, and $\gamma_a T_a$ is the direct effect coming from a itself being treated. The TTE is $\frac{1}{n_a} \sum_a [\mathcal{W}_a(\mathbf{1})\beta_a + \gamma_a]$. Let $\hat{\gamma}_a = T_a y_a / p - (1 - T_a) y_a / (1 - p)$. The previous result on unbiasedness in the bipartite setting can be directly applied to construct a class of unbiased estimators in the unipartite setting. We employ the previously developed bipartite estimator to estimate the indirect effect and use $\hat{\gamma}_a$ as an unbiased estimator of the direct effect. Summing the two gives an unbiased estimator of

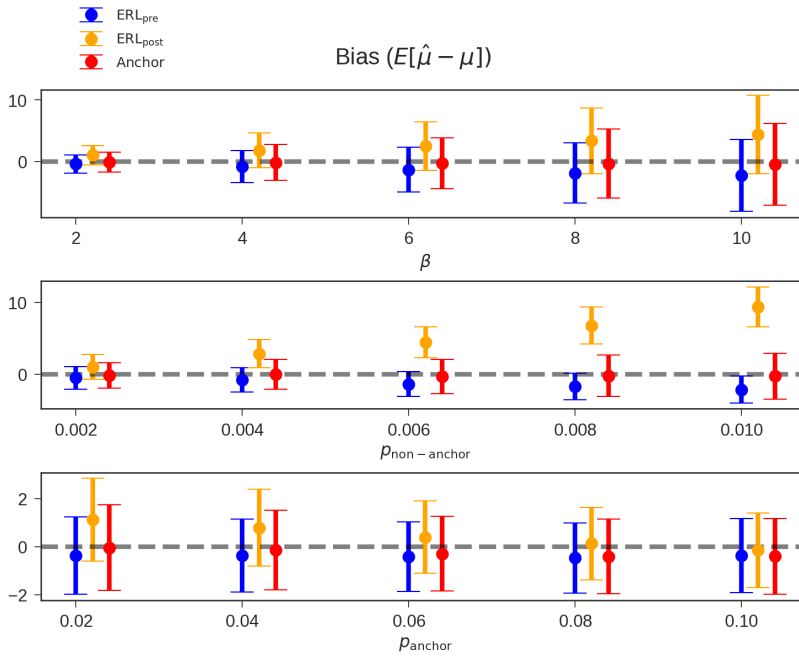


Figure 2: Bias with respect to different parameter values. Error bars are empirical standard deviations of the estimators over 1000 Monte Carlo runs.




Figure 3: 95%-CI coverage from 1000 independent Monte Carlo runs with respect to different parameter values. Error bars represent the 95%-CI for the true coverage probability.

the TTE, $\hat{\mu}_{uni}^{u,c}$, defined as

$$\hat{\mu}_{uni}^{u,c} = \frac{1}{n_a} \sum_a \hat{\mu}_{a,uni}^{u,c} = \frac{1}{n_a} \sum_a \left[\hat{\beta}_a^u \cdot \widehat{\mathcal{W}}_a(\mathbf{1})^c + \hat{\gamma}_a \right]. \quad (6)$$

Corollary 4. Under the assumptions of Theorem 1, with the outcome model $y_a = \alpha_a + \beta_a x_a + \gamma_a T_a$, if $w_{aa} = u_{aa} = 0$ for all a , an unbiased estimator of the TTE is $\hat{\mu}_{uni}^{u,c}$.

The restriction $w_{aa} = u_{aa} = 0$ is innocuous given the presence of a separate term $\gamma_a T_a$ in the outcome model for the direct effect. The summands in (6) follow the same dependence structure as those in (4). Consistency and asymptotic normality consequently follow directly from the same arguments as the bipartite case, under the same conditions on the growth of the maximum degrees, d_A and d_R .

Corollary 5. Under the assumptions of Theorem 2, with the outcome model $y_a = \alpha_a + \beta_a x_a + \gamma_a T_a$, if $w_{aa} = u_{aa} = 0$ for all a , then

1. $\text{Var}(\hat{\mu}_{uni}^{u,c}) = O(d_A^3 d_R / n_a)$, and thus $\hat{\mu}_{uni}^{u,c}$ is consistent if $d_A^3 d_R / n_a \rightarrow 0$;

2. if $n_a \text{Var}(\hat{\mu}_{uni}^{u,c})$ is bounded away from zero and $d_A^{10} d_R^4 / n_a \rightarrow 0$, then $(\hat{\mu}_{uni}^{u,c} - \mu) / \sqrt{\text{Var}(\hat{\mu}_{uni}^{u,c})}$ converges in distribution to $\mathcal{N}(0, 1)$.

7 Conclusion

We establish a foundation for estimating and testing treatment effects on endogenous graphs, but significant extensions and open questions remain. Future research could characterize the variance-minimizing choice of weights u_{ar} used to construct instruments z_a^u . A consistent and scalable variance estimator could be used in conjunction with Theorem 2 to construct confidence intervals around the TTE estimator, complementing the testing results of Section 4. And while we generalize the linear outcome model from the conventional exogenous edge setting, the outcome model may be implausible in some applications. We believe the research direction of causal inference on endogenous graphs holds promise for advancing our understanding of treatment effects in complex, real-world networks.

References

- An, W. 2011. Models and methods to identify peer effects. *The Sage handbook of social network analysis*, 514–532.
- Angrist, J. D.; Imbens, G. W.; and Rubin, D. B. 1996. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434): 444–455.
- Aral, S.; Muchnik, L.; and Sundararajan, A. 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51): 21544–21549.
- 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*, 1912–1947.
- 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.
- Bakshy, E.; Eckles, D.; Yan, R.; and Rosenn, I. 2012. Social influence in social advertising: evidence from field experiments. In *Proceedings of the 13th ACM conference on electronic commerce*, 146–161.
- Basse, G. W.; and Airolidi, E. M. 2018. Limitations of design-based causal inference and A/B testing under arbitrary and network interference. *Sociological Methodology*, 48(1): 136–151.
- Bond, R. M.; Fariss, C. J.; Jones, J. J.; Kramer, A. D.; Marlow, C.; Settle, J. E.; and Fowler, J. H. 2012. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415): 295–298.
- Bramoullé, Y.; Djebbari, H.; and Fortin, B. 2020. Peer effects in networks: A survey. *Annual Review of Economics*, 12(1): 603–629.
- Calvó-Armengol, A.; Patacchini, E.; and Zenou, Y. 2009. Peer effects and social networks in education. *The review of economic studies*, 76(4): 1239–1267.
- Chandrasekhar, A. 2016. Econometrics of network formation. *The Oxford handbook of the economics of networks*, 303–357.
- Chawla, S.; Hartline, J.; and Nekipelov, D. 2016. A/B testing of auctions. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 19–20.
- Chin, A. 2018. Central limit theorems via Stein’s method for randomized experiments under interference. *arXiv preprint arXiv:1804.03105*.
- Comola, M.; and Prina, S. 2021. Treatment effect accounting for network changes. *Review of Economics and Statistics*, 103(3): 597–604.
- Coppock, A.; Guess, A.; and Ternovski, J. 2016. When treatments are tweets: A network mobilization experiment over Twitter. *Political Behavior*, 38: 105–128.
- Cortez, M.; Eichhorn, M.; and Yu, C. 2022. Staggered rollout designs enable causal inference under interference without network knowledge. *Advances in Neural Information Processing Systems*, 35: 7437–7449.
- Cox, D. R. 1958. *Planning of experiments*. Wiley.
- DiMaggio, P.; and Garip, F. 2012. Network effects and social inequality. *Annual review of sociology*, 38(1): 93–118.
- 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.
- Egami, N. 2021. Spillover effects in the presence of unobserved networks. *Political Analysis*, 29(3): 287–316.
- Forastiere, L.; Mealli, F.; Wu, A.; and Airolidi, E. M. 2022. Estimating causal effects under network interference with bayesian generalized propensity scores. *Journal of Machine Learning Research*, 23(289): 1–61.
- Gao, M. 2024. Endogenous Interference in Randomized Experiments. *arXiv preprint arXiv:2412.02183*.
- Gilotte, A.; Calauzènes, C.; Nedelec, T.; Abraham, A.; and Dollé, S. 2018. Offline a/b testing for recommender systems. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 198–206.
- Halloran, M. E.; and Hudgens, M. G. 2016. Dependent happenings: a recent methodological review. *Current epidemiology reports*, 3: 297–305.
- Harshaw, C.; Sävje, F.; Eisenstat, D.; Mirrokni, V.; and Pouget-Abadie, J. 2021. Design and Analysis of Bipartite Experiments under a Linear Exposure-Response Model. *arXiv:2103.06392*.
- Horvitz, D. G.; and Thompson, D. J. 1952. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260): 663–685.
- 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.; Li, H.; Liskovich, I.; and Weintraub, G. Y. 2022. 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*, 3106–3116.
- Leung, M. P. 2022. Causal inference under approximate neighborhood interference. *Econometrica*, 90(1): 267–293.
- Liu, L.; and Hudgens, M. G. 2014. Large sample randomization inference of causal effects in the presence of interference. *Journal of the American statistical association*, 109(505): 288–301.
- Manski, C. F. 2013. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1): S1–S23.
- Nandy, P.; Venugopalan, D.; Lo, C.; and Chatterjee, S. 2021. A/b testing for recommender systems in a two-sided marketplace. *Advances in Neural Information Processing Systems*, 34: 6466–6477.

- Pouget-Abadie, J.; Aydin, K.; Schudy, W.; Brodersen, K.; and Mirrokni, V. 2019. Variance reduction in bipartite experiments through correlation clustering. *Advances in Neural Information Processing Systems*, 32.
- Ross, N. 2011. Fundamentals of Stein's method. *Probability Surveys*, 8: 210–293.
- Rubin, D. B. 1980. Randomization analysis of experimental data: The Fisher randomization test comment. *Journal of the American Statistical Association*, 75(371): 591–593.
- Ryu, S. 2024. Causal Effects of Treatments with Network Changes.
- Sävje, F. 2024. Causal inference with misspecified exposure mappings: separating definitions and assumptions. *Biometrika*, 111(1): 1–15.
- Sävje, F.; Aronow, P.; and Hudgens, M. 2021. Average treatment effects in the presence of unknown interference. *Annals of statistics*, 49(2): 673.
- Shi, L.; Bakhitov, E.; Hung, K.; Karrer, B.; Walker, C.; Bhole, M.; and Schrijvers, O. 2024. Scalable Analysis of Bipartite Experiments. *arXiv preprint arXiv:2402.11070*.
- Shirani, S.; and Bayati, M. 2023. Causal Message Passing: A Method for Experiments with Unknown and General Network Interference. *arXiv preprint arXiv:2311.08340*.
- Stein, C. 1986. Approximate Computation of Expectations. *Lecture Notes-Monograph Series*, 7: i–164.
- Taylor, S. J.; and Eckles, D. 2018. Randomized experiments to detect and estimate social influence in networks. *Complex spreading phenomena in social systems: Influence and contagion in real-world social networks*, 289–322.
- Tchetgen, E. J. T.; and VanderWeele, T. J. 2012. On causal inference in the presence of interference. *Statistical methods in medical research*, 21(1): 55–75.
- Wang, Y.; Samii, C.; Chang, H.; and Aronow, P. M. 2020. Design-based inference for spatial experiments under unknown interference. *arXiv preprint arXiv:2010.13599*.
- Yu, C. L.; Airoidi, 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.
- Zigler, C. M.; and Papadogeorgou, G. 2021. Bipartite causal inference with interference. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 36(1): 109.