

Causal Inference Under Approximate Neighborhood Interference*

Michael P. Leung[†]

April 14, 2021

ABSTRACT. This paper studies causal inference in randomized experiments under network interference. Commonly used models of interference posit that treatments assigned to alters beyond a certain network distance from the ego have no effect on the ego’s response. However, this assumption is violated in common models of social interactions. We propose a substantially weaker model of “approximate neighborhood interference” (ANI) under which treatments assigned to alters further from the ego have smaller, but potentially nonzero, effects on the ego’s response. We formally verify that ANI holds for well-known models of social interactions. Under ANI, restrictions on the network topology, and asymptotics under which the network size increases, we prove that standard inverse-probability weighting estimators consistently estimate useful exposure effects and are approximately normal. For inference, we consider a network HAC variance estimator. Under a finite population model, we show that the estimator is biased but that the bias can be interpreted as the variance of unit-level exposure effects. This generalizes Neyman’s well-known result on conservative variance estimation to settings with interference.

JEL CODES: C22, C31, C57

KEYWORDS: causal inference, network interference, social networks

*This research is supported by National Science Foundation grant SES-1755100. I thank Guido Imbens for suggestions that led to a simpler formulation of Assumption 4. I also thank Dean Eckles, Fredrik Sävje, and seminar audiences at Harvard, MIT, Northwestern, Rutgers, Stanford, UCL, and the 2019 SAMSI Causal Inference Workshop for useful feedback.

[†]Department of Economics, University of Southern California. E-mail: leungm@usc.edu.

1 Introduction

Randomized experiments in settings with network interference have seen increasing use in economics and the social sciences.¹ This paper develops methods for causal inference in experiments under weak restrictions on interference. We consider a finite population of n units connected through a network \mathbf{A} . Let $Y_i(\mathbf{d})$ denote the potential outcome of unit i under the counterfactual that the network is assigned treatment vector $\mathbf{d} = (d_i)_{i=1}^n \in \{0, 1\}^n$. The dependence of $Y_i(\mathbf{d})$ on the entire vector of assignments allows for “interference” or “spillovers” and distinguishes this setup from the standard potential outcomes model. Interest centers on “exposure effects,” defined below, which summarize how outcomes change in response to manipulations of \mathbf{d} .

The main inferential challenge is that, with a single network, the econometrician observes only one realization of the treatment assignment vector $\mathbf{D} = (D_i)_{i=1}^n$, where $D_i \in \{0, 1\}$ denotes unit i ’s realized assignment. Identification of exposure effects is therefore impossible without restrictions on the manner in which $Y_i(\cdot)$ varies with \mathbf{d} . The predominant approach in the literature is to assume interference operates through a low-dimensional vector of sufficient statistics.² That is, $Y_i(\mathbf{D})$ is only a function of \mathbf{D} through a vector-valued *exposure mapping*

$$T_i \equiv T(i, \mathbf{D}, \mathbf{A})$$

whose dimension is fixed with respect to n , unlike that of \mathbf{D} .

For example, [Cai et al. \(2015\)](#) run an experiment to study the effect of providing information on the benefits of weather insurance on farmers’ take-up of insurance. The authors are interested in spillover effects since a farmer who obtains the information may disseminate it to her social contacts. They estimate linear versions of the model

$$Y_i(\mathbf{D}) = \tilde{Y}_i(T_i), \quad \text{where} \quad T_i = \left(D_i, \frac{\sum_j A_{ij} D_j}{\sum_j A_{ij}} \right) \quad (1)$$

and A_{ij} is an indicator for whether farmers i and j are friends. Here the exposure mapping T_i is two-dimensional. Variation in its first component identifies the direct

¹E.g. [Bandiera et al. \(2009\)](#), [Bond et al. \(2012\)](#), [Bursztyrn et al. \(2014\)](#), [Miguel and Kremer \(2004\)](#), [Paluck et al. \(2016\)](#).

²E.g. [Aronow and Samii \(2017\)](#), [Basse et al. \(2019\)](#), [Forastiere et al. \(2020\)](#), [Manski \(2013\)](#), [Toulis and Kao \(2013\)](#).

effect of the intervention (being offered to attend an information session on weather insurance), while variation in the second identifies a spillover effect.

Exposure mappings, if correctly specified, substantially reduce the dimensionality of the model since we can reparameterize potential outcomes as

$$Y_i(\mathbf{d}) = \tilde{Y}_i(t) \quad \text{for } t \in \mathcal{T}, \quad (2)$$

where \mathcal{T} is the range of $T(\cdot)$. Interest then centers on “exposure effects”

$$\frac{1}{n} \sum_{i=1}^n (\tilde{Y}_i(t) - \tilde{Y}_i(t')) \quad (3)$$

for $t, t' \in \mathcal{T}$, which measure the average change in potential outcomes in response to counterfactual manipulations of the exposure mapping.

The majority of the literature studies (2) under the assumption that T_i only depends on treatments assigned to the K -neighborhood of i for some small K .³ However, this imposes strong structural restrictions on the underlying outcome process that are incompatible with a variety of models of social interactions studied in the networks literature (Guilbeault et al., 2018; Jackson, 2010; Manski, 1993). In these models, interference can arise from units outside of the ego’s K -neighborhood, for any K . This is the case for models with endogenous peer effects, where outcomes are functions of the outcomes of neighbors (Eckles et al., 2017). Another example is the Cai et al. (2015) setting, where under a simple diffusion model, information obtained by treated units can eventually diffuse to distant alters, which violates (1).

Important recent work studies *misspecified* exposure mappings (Chin, 2019; Sävje et al., 2020; Sävje, 2021). The insight of this literature is that standard estimators for (3) unbiasedly estimate meaningful exposure effects even without imposing (2) to restrict interference, which indicates a certain robustness of the estimator to more general patterns of interference. However, under what conditions *inference* can be made similarly robust is a more challenging question. Without (2), potential outcomes can depend arbitrarily on the entire assignment vector \mathbf{D} , which makes large-sample inference impossible. These papers accordingly propose a variety of high-level conditions weaker than (2) that implicitly restrict interference in order to obtain large-sample results. Unfortunately, the connection between these conditions and the literature

³This is the set of units whose network distance from i is at most K , formally defined in §2.

on social interactions remains unclear. It is an open question whether models in the networks literature violating (2) (e.g. models with endogenous peer effects or other diffusion models) can satisfy these high-level conditions.

Our Contribution. We study inference on exposure effects with misspecified exposure mappings under a new restriction on interference. We propose a model of *approximate neighborhood interference* (ANI), which allows treatments assigned to units further from the ego to have potentially nonzero, but smaller, effects on the ego’s response. Unlike the existing literature, we formally verify ANI in well-known models of social interactions. We also show that, under ANI, the data satisfies ψ -*weak dependence*, a recently proposed notion of network dependence. This enables us to apply limit theorems due to [Kojevnikov et al. \(forthcoming\)](#) to establish that, under restrictions on the network topology, standard inverse probability weighting (IPW) estimators are consistent for certain exposure effects and asymptotically normal.

For inference, we consider a network HAC (heteroskedasticity and autocorrelation consistent) variance estimator and characterize its asymptotic bias under a finite population model. We show the bias can be interpreted as the variance of unit-level exposure effects, which is fundamentally unidentified even under no interference. This generalizes the well-known result on conservativeness of the standard variance estimator of the difference-in-means estimate under no interference (e.g. [Imbens and Rubin, 2015](#), Ch. 6.4.2) to settings with dependence due to interference.

Finally, we propose a novel bandwidth for the network HAC estimator based on the average path length of the observed network. For a given bandwidth rule, the HAC estimator has different rates of convergence depending on whether the average K -neighborhood size in \mathbf{A} grows exponentially or polynomially with K . The utility of the average path length is that its magnitude adapts to the neighborhood growth rate to better trade off bias and variance.

[Kojevnikov et al. \(forthcoming\)](#) and [Kojevnikov \(2021\)](#) respectively provide consistency results for network HAC and bootstrap variance estimators for ψ -weakly dependent network data. Their results pertain to settings in which the data is mean-homogeneous, which often holds in superpopulation models. We extend their results to settings with mean-heterogeneous data, as is the case in finite population models.

[Choi \(2017\)](#) and [Choi \(2018\)](#) study causal inference without imposing an exposure mapping model. These papers focus on different estimands than ours and assume

treatment responses satisfy a monotonicity condition, which we do not require. There is also work on testing for interference, which can be used to test for correct specification of exposure models (e.g. [Athey et al., 2018](#)).

Several papers in econometrics study causal inference under interference ([Baird et al., 2018](#); [He and Song, 2021](#); [Lazzati, 2015](#); [Leung, 2020](#); [Vazquez-Bare, 2020](#); [Viviano, 2021](#)). The second paper studies a dynamic setting, whereas the others focus on a static setting like ours but under correctly specified exposure mappings. Many assume the special case of *stratified interference* under which the data consists of many clusters and interference only operates within clusters. We instead consider a single large cluster with known network structure.

The next section states our basic assumptions. Then §3 presents our model of interference (ANI) and large-sample results. We discuss variance estimation in §4. In §5, we illustrate the performance of our methods in an empirical application and simulation study. Finally, §6 concludes. All proofs are given in Appendix C.

2 Setup

We consider a finite population model in which the only source of randomness is *design uncertainty* ([Abadie et al., 2020](#); [Imbens and Menzel, 2019](#)). That is, \mathbf{D} is the only random quantity. In the special case of no interference, this corresponds to the well-known Neyman causal model. The setup can be viewed as conditioning on the network and potential outcomes, which allows for arbitrary dependence between the two. This allows links to form at higher rates between units with similar unobservables, which corresponds to unobserved homophily, a well-known hindrance to identifying social interactions ([Shalizi and Thomas, 2011](#)).

Let $\mathcal{N}_n = \{1, \dots, n\}$ denote the set of units. We assume \mathbf{A} is an undirected and unweighted network with no self-links, represented as an adjacency matrix with ij th entry $A_{ij} \in \{0, 1\}$ denoting a potential link between units i and j . Let \mathcal{A}_n denote the set of such networks on n units. We assume the components of \mathbf{D} are independent across units but not necessarily identically distributed, which allows for assignment based on unit covariates and network position. For example, treatment may be assigned “optimally” according to these characteristics (e.g. [Viviano, 2021](#)), or randomization may be stratified, as in the empirical application in §5.1.

Let $\mathcal{T} \subseteq \mathbb{R}^{d_T}$ be a discrete set. For any $n \in \mathbb{N}$, an *exposure mapping* is a function

$T \equiv T_n: \mathcal{N}_n \times \{0, 1\}^n \times \mathcal{A}_n \rightarrow \mathcal{T}$. Most of the literature assumes $T(\cdot)$ is *correctly specified* in the sense that (2) holds, which enables a simple definition of exposure effects (3). We instead follow the literature on misspecified exposure mappings and employ $T(\cdot)$ only to define useful estimands that summarize treatment and spillover effects but not to restrict the true interference structure. This is a reasonable solution to the task of parsimoniously summarizing the causal effect of a high-dimensional vector \mathbf{D} on potential outcomes.

For any $n \in \mathbb{N}$ and $i \in \mathcal{N}_n$, a potential outcome $Y_i(\cdot)$ is a mapping from $\{0, 1\}^n$ to \mathbb{R} . Define the *unit-level exposure effect*

$$\tau_i(t, t') = \mu_i(t) - \mu_i(t'), \quad \text{where} \quad \mu_i(t) = \sum_{\mathbf{d} \in \{0, 1\}^n} Y_i(\mathbf{d}) \mathbf{P}(\mathbf{D} = \mathbf{d} \mid T_i = t)$$

and $t, t' \in \mathcal{T}$. This is the difference in unit i 's expected response under two different values of the exposure mapping. The estimand of interest is the average effect

$$\tau(t, t') = \mu(t) - \mu(t'), \quad \text{where} \quad \mu(t) = \frac{1}{n} \sum_{i=1}^n \mu_i(t), \quad (4)$$

what [Sävje \(2021\)](#) calls a “misspecification-robust exposure effect.” This is analogous to estimands proposed by [Hudgens and Halloran \(2008\)](#) but generalized to allow for an incomplete network. We refer to these references for more detailed discussion of interpretation, but the idea is analogous to (3), which is to compare the average outcomes of units under two different values of the exposure mapping.

Recall $T_i \equiv T(i, \mathbf{D}, \mathbf{A})$, and define the generalized propensity score ([Imbens, 2000](#))

$$\pi_i(t) = \mathbf{E}[\mathbf{1}_i(t)], \quad \text{where} \quad \mathbf{1}_i(t) = \mathbf{1}\{T_i = t\}.$$
⁴

We estimate $\tau(t, t')$ using the standard IPW estimator, which is unbiased:

$$\hat{\tau}(t, t') = \hat{\mu}(t) - \hat{\mu}(t'), \quad \text{where} \quad \hat{\mu}(t) = \frac{1}{n} \sum_{i=1}^n Y_i \frac{\mathbf{1}_i(t)}{\pi_i(t)}.$$
⁵

⁴Computation of $\pi_i(t)$ depends on the exposure mapping and design. In §5.1, where treatments are block randomized, it can be easily computed in closed form using the hypergeometric distribution.

⁵The Hájek estimator replaces $\hat{\mu}(t)$ with $\sum_{i=1}^n Y_i \frac{\mathbf{1}_i(t)}{\pi_i(t)} / \sum_{i=1}^n \frac{\mathbf{1}_i(t)}{\pi_i(t)}$, which improves efficiency at the cost of finite-sample bias ([Hirano et al., 2003](#); [Aronow and Samii, 2017](#), §7.2). It is straightforward to extend our results to this estimator by deriving the usual asymptotically linear representation.

Most of the literature focuses on K -neighborhood exposure mappings, requiring $T(\cdot)$ to only be a function of \mathbf{d} and \mathbf{A} through i 's K -neighborhood, denoted

$$\mathcal{N}_{\mathbf{A}}(i, K) = \{j \in \mathcal{N}_n : \ell_{\mathbf{A}}(i, j) \leq K\},$$

where $\ell_{\mathbf{A}}(i, j)$ is the *path distance* between i, j .⁶ (Note that $i \in \mathcal{N}_{\mathbf{A}}(i, K)$.) To formalize this requirement, define $\mathbf{d}_{\mathcal{N}_{\mathbf{A}}(i, K)} = (d_j : j \in \mathcal{N}_{\mathbf{A}}(i, K))$ and $\mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i, K)} = (A_{kl} : k, l \in \mathcal{N}_{\mathbf{A}}(i, K))$, respectively the subvector of \mathbf{d} and subnetwork of \mathbf{A} on $\mathcal{N}_{\mathbf{A}}(i, K)$.

Assumption 1 (Exposure Mappings). *There exists $K \in \mathbb{N}$ such that, for any $n \in \mathbb{N}$ and $i \in \mathcal{N}_n$, $T(i, \mathbf{d}, \mathbf{A}) = T(i, \mathbf{d}', \mathbf{A}')$ for all $\mathbf{d}, \mathbf{d}' \in \{0, 1\}^n$ and $\mathbf{A}, \mathbf{A}' \in \mathcal{A}_n$ such that $\mathcal{N}_{\mathbf{A}}(i, K) = \mathcal{N}_{\mathbf{A}'}(i, K)$, $\mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i, K)} = \mathbf{A}'_{\mathcal{N}_{\mathbf{A}'}(i, K)}$, and $\mathbf{d}_{\mathcal{N}_{\mathbf{A}}(i, K)} = \mathbf{d}'_{\mathcal{N}_{\mathbf{A}'}(i, K)}$.*

This is a weak restriction on $T(\cdot)$ satisfied by most exposure mappings of interest in the literature. See for example (1) or the estimands in §5, where $K = 1$.

The next two assumptions impose uniform boundedness on potential outcomes and the generalized propensity score, which will involve bounding these quantities over all n . For this to be well-defined, we embed the observed network \mathbf{A} in a sequence of networks $\{\mathbf{A}_m\}_{m \in \mathbb{N}}$ where \mathbf{A}_m is a network on m nodes and $\mathbf{A}_n = \mathbf{A}$. We will be imposing conditions on this sequence, which amount to restrictions on the topology of the observed network for large n ; see §3.3 for further discussion.

Assumption 2 (Overlap). $\pi_i(t) \in [\underline{\pi}, \bar{\pi}] \subset (0, 1) \forall n \in \mathbb{N}, i \in \mathcal{N}_n, t \in \mathcal{T}$.

This requires the generalized propensity score to be uniformly bounded away from 0 and 1 over the sequence of networks. While overlap is standard, it can be restrictive. In particular, it restricts the exposure mapping, network sequence, distribution of treatments, and population. For instance, if $T_i = \mathbf{1}\{\sum_j A_{ij} D_j > 0\}$ and treatments are i.i.d., then overlap holds for $\pi_i(1)$ if treatment assignment is nontrivial and degrees $\sum_j A_{ij}$ are uniformly bounded over i, n . However, if i 's degree is large, then $\pi_i(1)$ will be close to one since it is highly likely to have at least one of many neighbors treated. At the cost of changing the estimand, overlap can be restored if we instead randomize

⁶A *path* between i, j is a sequence of links $A_{k_1 k_2}, A_{k_2 k_3}, \dots, A_{k_{m-1} k_m} = 1$ such that $k_1 = i$, $k_m = j$, and $k_a \neq k_b$ for all $a, b \in \{1, \dots, m\}$. The *length* of this path is $m - 1$. The path distance between i, j is the length of the shortest path between them, defined as ∞ if $i \neq j$ and no path exists and defined as 0 if $i = j$.

treatment only to a small subset of “eligible” units and restrict the population to units with an eligible neighbor. Then even with a large degree, the chance of having an *eligible* neighbor treated can be far from 0 and 1. This is the design in §5.1.

Assumption 3 (Bounded Outcomes). $|Y_i(\mathbf{d})| < \bar{Y} < \infty \forall n \in \mathbb{N}, i \in \mathcal{N}_n, \mathbf{d} \in \{0, 1\}^n$.

This assumption is standard and can be generalized to uniformly bounded moments.

3 Approximate Neighborhood Interference

We next present our model of interference. For any $\mathbf{d} \in \{0, 1\}^n$, partition $\mathbf{d} = (\mathbf{d}_{\mathcal{N}_A(i,s)}, \mathbf{d}_{-\mathcal{N}_A(i,s)})$, so that $\mathbf{d}_{-\mathcal{N}_A(i,s)} = (d_j : j \in \mathcal{N}_n \setminus \mathcal{N}_A(i,s))$. Let \mathbf{D}' be an independent copy of \mathbf{D} . Define $\mathbf{D}^{(i,s)} = (\mathbf{D}_{\mathcal{N}_A(i,s)}, \mathbf{D}'_{-\mathcal{N}_A(i,s)})$, obtained by concatenating the subvector of \mathbf{D} on $\mathcal{N}_A(i,s)$ with that of \mathbf{D}' on $\mathcal{N}_n \setminus \mathcal{N}_A(i,s)$. Finally, let

$$\theta_{n,s} \equiv \max_{i \in \mathcal{N}_n} \mathbf{E}[|Y_i(\mathbf{D}) - Y_i(\mathbf{D}^{(i,s)})|].$$

This measures interference from “distant” alters, those more than distance s away from the ego. It is the largest expected perturbation of any unit’s potential outcome due to redrawing the treatment assignments of distant alters.

Assumption 4 (ANI). $\sup_n \theta_{n,s} \rightarrow 0$ as $s \rightarrow \infty$.

This requires interference from distant alters to be negligible for large distances. In the special case of correct specification (2), there is *no* interference from units outside i ’s K -neighborhood, so $Y_i(\mathbf{D}) - Y_i(\mathbf{D}^{(i,s)}) = 0$ for all $s \geq K$ and any i . In contrast, ANI allows this difference to be nonzero for all s but requires that it decays with s . This means interference from alters beyond a unit’s s -neighborhood becomes increasingly negligible as we expand the radius s . Hence, ANI says that a unit’s response is primarily, but not entirely, determined by the assignments of alters close to it.⁷

ANI restricts the network, potential outcomes, and distribution of treatments. The first two are quite evident since the existence of interference is determined by

⁷Assumption 4 has some similarities with Assumption 6 of [Chin \(2019\)](#) in bounding the effect of manipulations of treatment assignments of distant units. For a CLT, we do not need a high-level condition analogous to his Assumption 5, which requires correlations between observed outcomes $\{Y_i\}_{i=1}^n$ to be sufficiently weak.

\mathbf{A} as well as the potential outcomes $Y_i(\cdot)$. For instance, if \mathbf{A} has no links for any n , then ANI holds because for $s > 0$, $Y_i(\mathbf{D}) - Y_i(\mathbf{D}^{(i,s)}) = 0$ for any potential outcome model, as no interference is possible without neighbors. On the other hand, for some $Y_i(\cdot)$, ANI can potentially hold for any \mathbf{A} . In the SUTVA case where $Y_i(\mathbf{d}) = \tilde{Y}_i(d_i)$ for all \mathbf{d} , this is clearly the case since for any \mathbf{A} and $s > 0$, $Y_i(\mathbf{D}) - Y_i(\mathbf{D}^{(i,s)}) = 0$.⁸

3.1 Social Interactions Models

We next verify ANI for two well-known models of social interactions. Our results yield uniform bounds on $\theta_{n,s}$ that decay exponentially with s under restrictions on the strength of social interactions.

Linear-in-Means Model. Consider a network version of the [Manski \(1993\)](#) model

$$Y_i = \alpha + \beta \frac{\sum_j A_{ij} Y_j}{\sum_j A_{ij}} + D_i \gamma + \varepsilon_i, \quad (5)$$

where the unobserved heterogeneity $\{\varepsilon_i\}_{i=1}^n$ is uniformly bounded (to impose Assumption 3) and nonrandom. As usual, to ensure the model is coherent, we assume

$$|\beta| < 1. \quad (6)$$

The model defines potential outcomes $Y_i(\mathbf{D})$ through its reduced form

$$\mathbf{Y} = \frac{\alpha}{1 - \beta} \mathbf{1} + \mathbf{D} \gamma + \gamma \beta \sum_{k=0}^{\infty} \beta^k \tilde{\mathbf{A}}^{k+1} \mathbf{D} + \sum_{k=0}^{\infty} \beta^k \tilde{\mathbf{A}}^k \boldsymbol{\varepsilon}$$

(e.g. [Bramoullé et al., 2009](#), eq. (6), assuming \mathbf{A} is connected), where $\mathbf{Y} = (Y_i(\mathbf{D}))_{i=1}^n$, $\boldsymbol{\varepsilon}$ is similarly defined, and $\tilde{\mathbf{A}}$ is the row-normalized version of \mathbf{A} (divide each row by its sum). The third term roughly says that the impact of treatments assigned to k -neighbors is exponentially down-weighted by β^k . This leads to the following result.

Proposition 1. *If responses are realized according to the linear-in-means model, then there exists $C > 0$ such that Assumption 4 holds with $\theta_{n,s} \leq C|\beta|^s$ for all n, s .*

⁸As pointed out by a referee, since $\theta_{n,s}$ is defined with respect to the design, it may be possible to choose the design to ensure that ANI holds for a given network and potential outcome model. See for example Proposition 2 below, where a sufficient condition for ANI depends on a parameter φ_j that is a function of the distribution of D_j .

Complex Contagion. We next consider a model of “complex contagion,” variants of which have been widely studied in the networks literature.⁹ Initialize a dynamic discrete-time process at period 0 at some binary response vector $\mathbf{Y}^0 \in \{0, 1\}^n$, which may be a function of treatments \mathbf{D} and nonrandom unobserved heterogeneity $\boldsymbol{\varepsilon} = (\varepsilon_i)_{i=1}^n$. For some \mathbb{R} -valued $\phi(\cdot)$, each unit i at period t updates according to

$$Y_i^t = \mathbf{1} \left\{ \beta \frac{\sum_j A_{ij} Y_j^{t-1}}{\sum_j A_{ij}} \geq \phi(D_i, \varepsilon_i) \right\} \quad (7)$$

to obtain new responses $\mathbf{Y}^t = (Y_i^t)_{i=1}^n$ from last period’s responses \mathbf{Y}^{t-1} . The rule says that i selects 1 over 0 if and only if the fraction of neighbors choosing response 1 in the previous period is large enough relative to the heterogeneous threshold $\phi(D_i, \varepsilon_i)$. The parameter β measures the strength of social interactions.

Because the setup in §2 is static, we consider running the dynamic process until the first period T such that $\mathbf{Y}^T = \mathbf{Y}^{T-1}$. To ensure such a T exists for any \mathbf{Y}^0 , we assume $\beta \geq 0$ (Milgrom and Roberts, 1990). We then take \mathbf{Y}^T as the vector of responses \mathbf{Y} observed in the data, which yields outcomes $(Y_i(\mathbf{D}))_{i=1}^n$. Hence, this process implicitly defines potential outcomes.

To verify Assumption 4, we need a condition analogous to (6), which will be more complicated to state since the model is nonlinear. Define a weighted directed network \mathbf{G} on \mathcal{N}_n with ij th entry $G_{ij} = A_{ij} \mathbf{E}[\varphi_j]$ for $\varphi_j = \mathbf{1}\{0 < \phi(D_j, \varepsilon_j) \leq \beta\}$. Let

$$\rho_n(\bar{s}) = \sup_{s \geq \bar{s}} \|\mathbf{G}^s\|_\infty^{1/s} \equiv \sup_{s \geq \bar{s}} \left(\max_{i \in \mathcal{N}_n} \sum_{j=1}^n (\mathbf{G}^s)_{ij} \right)^{1/s}$$

for any $\bar{s} > 0$, so $\|\cdot\|_\infty$ is the matrix norm induced by the vector ∞ -norm.

Proposition 2. *Let $\alpha_{n,s}(\bar{s}) = 2(\rho_n(\bar{s}))^{s-1} \mathbf{1}\{s - 1 \geq \bar{s}\} + \mathbf{1}\{s - 1 < \bar{s}\}$. Suppose responses are realized according to the complex contagion model. If $\sup_n \rho_n(\bar{s}) < 1$ for some $\bar{s} > 0$, then for this \bar{s} , Assumption 4 holds with $\theta_{n,s} \leq \alpha_{n,s}(\bar{s})$ for all n, s .*

We next discuss the interpretation of $\sup_n \rho_n(\bar{s}) < 1$ in relation to (6). Let \mathbf{I} be the $n \times n$ identity matrix. For the linear-in-means model to be coherent, we need $\mathbf{I} - \beta \tilde{\mathbf{A}}$ to be invertible, which is true provided $|\beta| < 1$. Since $\tilde{\mathbf{A}}$ is row-normalized,

⁹E.g. Granovetter (1978), Guilbeault et al. (2018), Jackson (2010), Montanari and Saberi (2010).

this is equivalent to

$$\lambda_{\max}(\beta\tilde{\mathbf{A}}) < 1, \quad (8)$$

where $\lambda_{\max}(\cdot)$ is the spectral radius (Bramoullé et al., 2009). On the other hand, $\|\mathbf{G}^s\|_{\infty}^{1/s} \xrightarrow{s \rightarrow \infty} \lambda_{\max}(\mathbf{G})$ by Gelfand's formula, so for any $\epsilon > 0$, we can choose \bar{s} large enough such that

$$\sup_n \lambda_{\max}(\mathbf{G}) < 1 - \epsilon \quad \text{implies} \quad \sup_n \rho_n(\bar{s}) < 1 \quad (9)$$

(Xu and Lee, 2015). The left-hand side is clearly analogous to (8). The difference is that, in \mathbf{G} , we weight each potential link A_{ij} by $\mathbf{E}[\varphi_j]$, whereas in $\beta\tilde{\mathbf{A}}$, the weight is $\beta/\sum_k A_{ik}$. However, both weights are monotonically increasing in β , so both (8) and (9) implicitly restrict social interactions, as measured by β .

3.2 Weak Dependence

Define $Z_i = (\mathbf{1}_i(t)\pi_i(t)^{-1} - \mathbf{1}_i(t')\pi_i(t')^{-1})Y_i$, so that $\hat{\tau}(t, t') = n^{-1} \sum_{i=1}^n Z_i$. For large-sample inference, we would like the data $\{Z_i\}_{i=1}^n$ to be at most weakly dependent. Recall that $\ell_{\mathbf{A}}(i, j)$ is the path distance between i, j in \mathbf{A} . Since treatments are independent, the indicators are weakly dependent in the sense that $\mathbf{1}_i(t) \perp\!\!\!\perp \mathbf{1}_j(t)$ if $\ell_{\mathbf{A}}(i, j) > 2K$ by Assumption 1. At first glance, $\{Y_i\}_{i=1}^n$ seems to be strongly dependent since Y_i is a function of \mathbf{D} for all i . However, ANI requires Y_i to primarily depend on the treatments of nearby units, which suggests that distant units have weakly dependent outcomes. We next formalize this idea.

We first define a notion of weak network dependence due to Kojevnikov et al. (forthcoming). For any $H, H' \subseteq \mathcal{N}_n$, define $\ell_{\mathbf{A}}(H, H') = \min\{\ell_{\mathbf{A}}(i, j) : i \in H, j \in H'\}$. Let $\mathbf{Z}_H = (Z_i : i \in H)$, \mathcal{L}_d be the set of bounded, \mathbb{R} -valued, Lipschitz functions on \mathbb{R}^d , and

$$\mathcal{P}_n(h, h'; s) = \{(H, H') : H, H' \subseteq \mathcal{N}_n, |H| = h, |H'| = h', \ell_{\mathbf{A}}(H, H') \geq s\}.$$

Definition 1. A triangular array $\{Z_i\}_{i=1}^n$ is ψ -weakly dependent if there exist (a) uniformly bounded constants $\{\tilde{\theta}_{n,s}\}_{s,n \in \mathbb{N}}$ with $\theta_{n,0} = 1 \forall n$ such that $\sup_n \tilde{\theta}_{n,s} \rightarrow 0$ as $s \rightarrow \infty$, and (b) functionals $\{\psi_{h,h'}(\cdot, \cdot)\}_{h,h' \in \mathbb{N}}$ with $\psi_{h,h'} : \mathcal{L}_h \times \mathcal{L}_{h'} \rightarrow [0, \infty)$ such that

$$|\text{Cov}(f(\mathbf{Z}_H), f'(\mathbf{Z}_{H'}))| \leq \psi_{h,h'}(f, f') \tilde{\theta}_{n,s} \quad (10)$$

for all $n, h, h' \in \mathbb{N}$; $s > 0$; $f \in \mathcal{L}_h$; $f' \in \mathcal{L}_{h'}$; and $(H, H') \in \mathcal{P}_n(h, h'; s)$.

This extends temporal ψ -weak dependence (Doukhan and Louhichi, 1999) to network data by using path in place of temporal distance. The concept says two sets of observations \mathbf{Z}_H and $\mathbf{Z}_{H'}$ have small covariance if they are sufficiently distant.

Let $i^*(\cdot)$ be the identity function $x \mapsto x$ on \mathbb{R} and $\mathcal{Z} \subseteq \mathbb{R}$ be any compact set such that $Z_i \in \mathcal{Z}$ for any $n \in \mathbb{N}$ and $i \in \mathcal{N}_n$, which exists by Assumptions 2 and 3. For any Lipschitz function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, let $\text{Lip}(f)$ its Lipschitz constant and $\|f\|_\infty = \sup_{x \in \mathcal{Z}^d} |f(x)|$. Let K be the constant in Assumption 1 and $\lfloor s \rfloor$ be s rounded down to the nearest integer.

Theorem 1 (Weak Dependence). *Under Assumptions 1–4, $\{Z_i\}_{i=1}^n$ is ψ -weakly dependent in that (10) holds with $\tilde{\theta}_{n,s} = \theta_{n, \lfloor s/2 \rfloor} \mathbf{1}\{s > 2 \max\{K, 1\}\} + \mathbf{1}\{s \leq 2 \max\{K, 1\}\}$ for all $n \in \mathbb{N}$ and $s > 0$ and*

$$\psi_{h,h'}(f, f') = 2(\|f\|_\infty \|f'\|_\infty + h \|f'\|_\infty \text{Lip}(f) + h' \|f\|_\infty \text{Lip}(f'))$$

for either $h, h' \in \mathbb{N}$, $f \in \mathcal{L}_h$, and $f' \in \mathcal{L}_{h'}$, or $h = h' = 1$ and $f = f' = i^*$.

3.3 Large-Sample Theory

In what follows, define $\tilde{\theta}_{n,s}$ as in Theorem 1. Having established that $\{Z_i\}_{i=1}^n$ is ψ -weakly dependent, we can apply results due to Kojevnikov et al. (forthcoming) to show that $\hat{\tau}(t, t')$ is consistent and asymptotically normal. Their results require $\tilde{\theta}_{n,s}$ to decay to zero fast enough. The speed of decay depends on the network topology, in particular the growth rate of s -neighborhood sizes. Intuitively, ANI says that Z_i depends primarily on units in $\mathcal{N}_A(i, s)$. Hence, if the typical size of these neighborhoods grows rapidly with s , then weak dependence requires that this be counterbalanced by having the covariances ($\tilde{\theta}_{n,s}$) decay to zero faster with s .

The next assumptions formalize these ideas. To clarify their interpretation, in §A.1, we verify them for networks with polynomial and exponential neighborhood growth rates when $\tilde{\theta}_{n,s}$ decays exponentially with s , as in the examples in §3.1.

Limit Sequence. Our results take $n \rightarrow \infty$ along a sequence of networks $\{\mathbf{A}_n\}_{n \in \mathbb{N}}$ defined prior to Assumption 2. The results hold for any sequence satisfying our assumptions. The design may implicitly depend on n , so the distribution of treatments

may also change along the sequence, so long as treatments remain independent across units. Additionally, potential outcomes $Y_i(\cdot)$ and the exposure mapping $T(\cdot)$ may vary with n (and must do so due to the dimensions of their arguments), so long as they satisfy Assumptions 1 and 3. We emphasize that K and \mathcal{T} do *not* depend on n , but all other quantities may vary along the sequence, within the confines of the stated assumptions.

Let $\mathcal{N}_{\mathbf{A}}^{\hat{\partial}}(i, s) = \{j \in \mathcal{N}_n : \ell_{\mathbf{A}}(i, j) = s\}$ be the s -neighborhood boundary of i , the set of units exactly distance s from i , and $M_n^{\hat{\partial}}(s) = n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}^{\hat{\partial}}(i, s)|$, its average size.

Assumption 5 (Weak Dependence for LLN). $\sum_{s=0}^n M_n^{\hat{\partial}}(s) \tilde{\theta}_{n,s} = o(n)$.

This corresponds to Assumption 3.2 of [Kojevnikov et al. \(forthcoming\)](#). It restricts the network topology through $M_n^{\hat{\partial}}(s)$ and the degree of interference through $\tilde{\theta}_{n,s}$. Since the former grows with s , the latter must decay to zero faster for the sum to be $o(n)$. Furthermore, since $\tilde{\theta}_{n,s} = 1$ for $s \leq 2 \max\{K, 1\}$, this implies $M_n^{\hat{\partial}}(1) = o(n)$, which is a restriction on network density that rules out, for instance, \mathbf{A} being complete.

It is useful to compare this to its analog for α -mixing spatial processes. Consider, for example, Assumption 3(b) of [Jenish and Prucha \(2009\)](#), which essentially requires $\sum_{s=1}^{\infty} s^{d-1} \alpha(s) < \infty$, where d is the dimension of the underlying space and $\alpha(s)$ is the α -mixing coefficient, which measures dependence between sets of observations at spatial distance s apart. In the spatial setting, the s -neighborhood boundary of i is the set of units at any distance $h \in [s, s+1)$ from i . By their Lemma A.1(iii), the size of this set is $O(s^{d-1})$. Thus, we have an analogous trade-off between the sizes of spatial s -neighborhood boundaries and the rate of decay of the mixing coefficient.

Theorem 2 (Consistency). *Under Assumptions 1–5, $|\hat{\tau}(t, t') - \tau(t, t')| \xrightarrow{p} 0$.*

Inspection of the proof shows that we can sharpen the result to $|\hat{\tau}(t, t') - \tau(t, t')| = O_p(n^{-1/2})$ if we strengthen Assumption 5 to $\sum_{s=0}^n M_n^{\hat{\partial}}(s) \tilde{\theta}_{n,s} = O(1)$.

Asymptotic normality requires a stronger version of Assumption 5. Let $M_n(s, k) = n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, s)|^k$, the k th moment of the s -neighborhood size, and

$$\mathcal{H}_n(s, m) = \{(i, j, k, l) \in \mathcal{N}_n^4 : k \in \mathcal{N}_{\mathbf{A}}(i, m), l \in \mathcal{N}_{\mathbf{A}}(j, m), \ell_{\mathbf{A}}(\{i, k\}, \{j, l\}) = s\}.$$

This is the set of paired couples (i, j) and (k, l) such that the units within each couple

are at most path distance m apart from one another, and the two pairs are exactly path distance s apart. Define $\sigma_n^2 = \text{Var}(n^{-1/2} \sum_{i=1}^n Z_i)$.

Assumption 6 (Weak Dependence for CLT). *There exist $\epsilon > 0$ and a sequence of positive constants $\{m_n\}_{n \in \mathbb{N}}$ such that $m_n \rightarrow \infty$ and*

$$\max \left\{ \sigma_n^{-4} \frac{1}{n^2} \sum_{s=0}^n |\mathcal{H}_n(s, m_n)| \tilde{\theta}_{n,s}^{1-\epsilon}, \quad \sigma_n^{-3} n^{-1/2} M_n(m_n, 2), \quad \sigma_n^{-1} n^{3/2} \tilde{\theta}_{n, m_n}^{1-\epsilon} \right\} \rightarrow 0. \quad (11)$$

This strengthens Assumption 5 and corresponds to Assumption 3.4 of [Kojevnikov et al. \(2019\)](#).¹⁰ Similar to Assumption 5, the first term in (11) requires $\tilde{\theta}_{n,s}$ to decay to zero fast enough relative to s -neighborhood sizes. The second term restricts s -neighborhood growth rates, while the third requires sufficiently fast decay of $\tilde{\theta}_{n,s}$. See §A.1 for a discussion of the plausibility of this assumption.

Theorem 3 (Asymptotic Normality). *Under Assumptions 1–4 and 6,*

$$\sigma_n^{-1} \sqrt{n} (\hat{\tau}(t, t') - \tau(t, t')) \xrightarrow{d} \mathcal{N}(0, 1).$$

4 Variance Estimation

For large-sample inference, we consider the variance estimator

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (Z_i - \hat{\tau}(t, t')) (Z_j - \hat{\tau}(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}, \quad (12)$$

where Z_i is defined at the top of §3.2 and $b_n \geq 0$ is a bandwidth parameter discussed below in (13). When $b_n = 0$, this reduces to the sample variance of the Z_i 's, which is only a valid estimator under no interference. Choosing $b_n > 0$ places nonzero weight on pairs at most b_n apart in the network, which accounts for possible autocorrelation.¹¹

Under correctly specified exposure mappings, we can choose $b_n = 2K$, as in [Leung](#)

¹⁰The publication version of their paper ([Kojevnikov et al., forthcoming](#)) formulates this assumption slightly differently.

¹¹The estimator is simple to compute. First calculate the path distance matrix $(\ell_{\mathbf{A}}(i, j))_{i, j \in \mathcal{N}_n}$, which can be done very efficiently for sparse networks using Dijkstra's algorithm (e.g. the Python function `dijkstra` in the `scipy.sparse.csgraph` module). Then for the matrix $\mathbf{P} = (\mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\})_{i, j \in \mathcal{N}_n}$ and vector $\tilde{\mathbf{Z}} = (n^{-1/2}(Z_i - \hat{\tau}(t, t')))_{i=1}^n$, we have $\hat{\sigma}^2 = \tilde{\mathbf{Z}}' \mathbf{P} \tilde{\mathbf{Z}}$.

(2020), since $\mathbf{1}_i(t) \perp \mathbf{1}_j(t)$ if $\ell_{\mathbf{A}}(i, j) > 2K$. For misspecified exposure mappings, b_n must grow with n at a rate depending on the network topology. In this case, (12) corresponds to a HAC estimator familiar from the time series literature but using network distance in place of temporal distance. This estimator has been previously used in practice (e.g. Acemoglu et al., 2015), and its formal properties were first studied by Kojevnikov et al. (forthcoming) in a superpopulation setting. Building on their work, we characterize its behavior in a finite population model.

Remark 1. Kojevnikov et al. (forthcoming) consider general kernel functions that include the uniform kernel in (12). Kojevnikov (2021) proposes alternative weights motivated by network bootstrap procedures, which are guaranteed to be positive semidefinite (PSD) in finite sample, unlike kernel-based weights. A previous version of this paper (Leung, 2019a) considered a modification of his wild bootstrap weights, which are also guaranteed PSD and therefore may be used when (12) is not PSD (which never occurred in our simulations in §5.2). While these weights all have the same asymptotic properties, in simulation experiments, we find that uniform weights better control size in small samples than these alternatives that decay with distance.

Choice of Bandwidth. For $\hat{\sigma}^2$ to have good large-sample properties, we need to restrict the rate at which $|\mathcal{N}_{\mathbf{A}}(i, b_n)|$ diverges with n (see Assumption 7). That is, how fast b_n can diverge depends on how rapidly s -neighborhoods expand with s . In spatial settings, K -neighborhood sizes grow polynomially with K , so b_n is allowed to diverge at a polynomial rate. (A faster rate is better for bias but worse for variance.) However, in network settings, K -neighborhood sizes can grow exponentially with K . Based on the analysis in §A.2, we suggest b_n be chosen as follows. Let $\delta(\mathbf{A}) = n^{-1} \sum_{i,j} A_{ij}$ denote the average degree and $\mathcal{L}(\mathbf{A})$ the *average path length* (APL).¹² We propose

$$b_n = \lfloor \max\{\tilde{b}_n, 2K\} \rfloor \quad \text{for} \quad \tilde{b}_n = \begin{cases} \frac{1}{2} \mathcal{L}(\mathbf{A}) & \text{if } \mathcal{L}(\mathbf{A}) < 2 \frac{\log n}{\log \delta(\mathbf{A})}, \\ \mathcal{L}(\mathbf{A})^{1/3} & \text{otherwise,} \end{cases} \quad (13)$$

where $\lfloor \cdot \rfloor$ means round to the nearest integer.¹³ We suggest in practice the researcher report results for several bandwidths in a neighborhood of (13).

¹²The APL is the average value of $\ell_{\mathbf{A}}(i, j)$ over all pairs in the largest component of \mathbf{A} . A *component* of a network is a connected subnetwork such that all units in the subnetwork are disconnected from those not in the subnetwork.

¹³The formula assumes $\delta(\mathbf{A}) > 1$, which is typically the case in practice.

To account for correlation in $\{\mathbf{1}_i(t)\}_{i=1}^n$ discussed in §3.2, we set b_n at least equal to $2K$. The fractions $1/2$ and $1/3$ are due to Assumption 7 below (see §A.2). The purpose of comparing $\mathcal{L}(\mathbf{A})$ and $\log n / \log \delta(\mathbf{A})$ is to determine whether K -neighborhood sizes grow approximately exponentially or polynomially with K . As discussed in §A.2, in the exponential case, the difference between these two statistics typically converges to zero, whereas in the polynomial case, $\mathcal{L}(\mathbf{A})$ is much larger, having polynomial order. See, for example, the simulations in §5.2, which show at least a four-fold difference in APL between the two regimes. Thus, (13) selects a bandwidth of logarithmic (polynomial) order when neighborhood growth rates are approximately exponential (polynomial).¹⁴

Bias of $\hat{\sigma}^2$. Define

$$\hat{\sigma}_*^2 = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (Z_i - \tau_i(t, t'))(Z_j - \tau_j(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\} \quad \text{and}$$

$$R_n = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (\tau_i(t, t') - \tau(t, t'))(\tau_j(t, t') - \tau(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}.$$

The former is an “oracle” version of $\hat{\sigma}^2$ that replaces $\hat{\tau}(t, t')$ with $\tau_i(t, t')$, while the latter is a bias term. Theorem 4 below establishes that

$$\hat{\sigma}^2 = \hat{\sigma}_*^2 + R_n + o_p(1) \quad \text{and} \quad (14)$$

$$|\hat{\sigma}_*^2 - \text{Var}(\sqrt{n}\hat{\tau}(t, t'))| \xrightarrow{p} 0. \quad (15)$$

Equation (15) says that the oracle estimator is consistent for the variance, and (14) says that our estimator is biased. The source of bias is mean-heterogeneity: $\hat{\tau}(t, t')$ is consistent for $\tau(t, t')$ but not $\tau_i(t, t')$, which is heterogeneous across units.

The bias R_n has the form of a HAC estimate of the variance of the unit-level exposure effects. It is helpful to compare this to the case of no interference, where $T_i = D_i$, $t = 1$, and $t' = 0$, so that $\tau(t, t')$ is the usual average treatment effect (ATE).

¹⁴In the exponential case, we need $b_n = O(\log n)$ for $\text{Var}(\hat{\sigma}^2)$ to be small, which (13) accomplishes since $\mathcal{L}(\mathbf{A}) \approx \log n / \log \delta(\mathbf{A})$ in this regime. In the polynomial case, $b_n = O(\log n)$ is also valid, but the bias then vanishes at an extremely slow rate (see [Kojevnikov et al., forthcoming](#), proof of Proposition 4.1). Our choice of $\mathcal{L}(\mathbf{A})^{1/3}$ substantially improves this rate since $\mathcal{L}(\mathbf{A})$ is polynomial in n in this regime.

Knowing that units are independent, we can choose $b_n = 0$, in which case

$$R_n = \frac{1}{n} \sum_{i=1}^n (\tau_i(1, 0) - \tau(1, 0))^2.$$

This is the well-known asymptotic bias of the standard variance estimator for the difference-in-means estimate of the ATE (e.g. [Imbens and Rubin, 2015](#), Theorem 6.2). It measures the variance of the unit-level treatment effects and is generally impossible to estimate in the finite population setting, so the variance estimator is conservative. In the special case of homogeneous unit-level treatment effects, meaning $\tau_i(t, t')$ does not vary with i , the bias is zero, a property also shared by our R_n . Thus, (14) generalizes Neyman’s well-known result on conservative variance estimation to a setting with interference. The additional covariance terms in R_n weighted by $\mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}$ account for dependence due to interference.

Remark 2. In the Appendix B, we compare R_n with the bias of the [Aronow and Samii \(2017\)](#) estimator for correctly specified exposure mappings. Simulation results there show that our bias is positive but can be notably smaller than theirs. More generally, the asymptotic behavior of R_n depends on the superpopulation model towards which our framework is agnostic. Since R_n has the form of a network HAC, we expect that it typically converges to the population variance of the unit-level exposure effects, although this requires additional weak dependence conditions on the distributions of $Y_i(\mathbf{d})$ and \mathbf{A} . Some such conditions are given in Theorem 4.2 of [Leung \(2019b\)](#).

To show consistency of $\hat{\sigma}^2$, define

$$\mathcal{J}_n(s, m) = \{(i, j, k, l) \in \mathcal{N}_n^4 : k \in \mathcal{N}_{\mathbf{A}}(i, m), l \in \mathcal{N}_{\mathbf{A}}(j, m), \ell_{\mathbf{A}}(i, j) = s\}.$$

This is similar to, and evidently contains, $\mathcal{H}_n(s, m)$ from Assumption 6.

Assumption 7 (Weak Dependence for $\hat{\sigma}^2$). (a) $\sum_{s=0}^n M_n^{\hat{\sigma}}(s) \tilde{\theta}_{n,s}^{1-\epsilon} = O(1)$ for some $\epsilon > 0$, (b) $M_n(b_n, 1) = o(n^{1/2})$, (c) $M_n(b_n, 2) = o(n)$, (d) $\sum_{s=0}^n |\mathcal{J}_n(s, b_n)| \tilde{\theta}_{n,s} = o(n^2)$.

Part (a) strengthens Assumption 5, while (b)–(d) regulate how fast b_n grows relative to neighborhood sizes. In §A.2, we use these to derive (13) and illustrate the plausibility of the conditions for some general classes of graphs. Part (b) allows us to replace

$\hat{\tau}(t, t')$ in $\hat{\sigma}^2$ with its expectation. Part (d) is used to derive the asymptotic bias. It is similar to the first requirement of Assumption 6, with b_n and $\mathcal{J}_n(s, \cdot)$ in place of m_n and $\mathcal{H}_n(s, \cdot)$, respectively.

Theorem 4 (Variance Estimator). *If $b_n \rightarrow \infty$ as $n \rightarrow \infty$, then under Assumptions 1–4 and 7, (14) and (15) hold.*

5 Numerical Illustrations

5.1 Empirical Application

We revisit a network experiment analyzed in [Paluck et al. \(2016\)](#) and [Aronow and Samii \(2017\)](#) that studies the effect of an anti-conflict intervention on adolescent social norms for antagonistic behavior, including harassment, rumor-mongering, social exclusion, and bullying. In the experimental design, 28 of 56 schools are first randomized into treatment. Then within treated schools, a subset of students are selected as eligible for treatment based on covariates, and half of eligibles are block-randomized into treatment. Treated students are invited to participate in bi-monthly meetings that follow an anti-conflict curriculum designed in part by the researchers of the study. At these meetings, a trained adult leader helps students identify social conflicts at their school and design strategies to reduce conflict.

[Aronow and Samii \(2017\)](#) and part of the analysis of [Paluck et al. \(2016\)](#) examine the causal effect of the offer to participate on endorsement of anti-conflict norms. This is measured by self reports of wearing a wristband disseminated as part of the program as a reward to students observed engaging in conflict-mitigating behavior. Through the course of the experiment, over 2500 wristbands were disseminated and tracked. We study exposure effects similar to those used in the application of [Aronow and Samii \(2017\)](#). Unlike their analysis, we restrict the data to the five largest treated schools to illustrate what can be learned from data on a few large networks. In each of our schools, the number of eligibles is exactly 64.

We estimate a treatment and a spillover effect. For the latter, the exposure mapping is $T_i = \mathbf{1}\{\sum_j A_{ij}D_j > 0\}$, an indicator for whether at least one friend is offered treatment. As in [Aronow and Samii \(2017\)](#), to ensure overlap, we restrict to the “spillover population” consisting of students that have at least one eligible friend

($n = 1685$). For the treatment effect, the exposure mapping is $T_i = D_i$, and we restrict to the “treatment population” consisting of students eligible for treatment ($n = 320$). For both, we compute path distances and average degree in the variance estimator (discussed below) using the full network.

Networks are measured by asking students to name up to ten students at the school with “whom they chose to spend time with in the last few weeks, either in school, out of school, or online.” Consequently, \mathbf{A} is directed. When computing the number of treated friends for the exposure mappings, we use the directionality of links. However, when computing network neighborhoods for our variance estimator, we ignore the directionality of links to conservatively define larger neighborhoods and avoid taking a stance on neighborhood definitions for directed networks.

We next provide some summary statistics. Within the treatment population, the average outcome Y_i is 0.16 (SD 0.37), and by block randomization, exactly 50 percent are treated. Within the spillover population, the average outcome is 0.11 (SD 0.32), and 58 percent (SD 0.49) have at least one treated friend. The data includes the blocks in which eligible students are block-randomized, so we can compute the propensity scores $\pi_i(t)$ for each student using the hypergeometric distribution. For the exposure mapping $T_i = \mathbf{1}\{\sum_j A_{ij}D_j > 0\}$, given N eligible neighbors, $\pi_i(0)$ is the chance of having 0 out of N successes when drawing without replacement. We find $n^{-1} \sum_{i=1}^n \pi_i(1) = 0.597$, which is close to empirical proportion of 58 percent.

The average out-degree $n^{-1} \sum_{i,j} A_{ij}$ is 7.96. The APL is small, on average 3.37 across our five schools. Since there are $n = 3306$ students, $\log n / \log \delta(\mathbf{A}) = 3.96$, which is very close to 3.37. Thus, given $K = 1$, our suggested bandwidth (13) is $b_n = 2$. We report results for the range of bandwidths $\{0, \dots, 3\}$, noting that 0 corresponds to standard errors in the no-interference case.

Table I presents the results. The first row is the IPW estimator for the indicated exposure effect, and the last four rows are standard errors for the indicated bandwidths. We find a large treatment effect of 0.15, which is significant at the 5 percent level across all bandwidths. The spillover effect is smaller at 0.04, with larger standard errors, and is statistically insignificant for our suggested bandwidth $b_n = 2$ at the 5 percent level. The small spillover estimate is largely in line with the estimates implied by Figure 3C of [Paluck et al. \(2016\)](#). It does not contradict the overall message of their paper since they find, for example, sizeable spillover effects when comparing treated and untreated schools. In contrast, our analysis so far only makes

Table I. Estimates and SEs

	Treatment	Spillover
$\hat{\tau}(1, 0)$	0.1500	0.0407
$\hat{\mu}(1)$	0.2375	0.1293
$\hat{\mu}(0)$	0.0875	0.0885
$b_n = 0$	0.0443	0.0167
$b_n = 1$	0.0460	0.0184
$b_n = 2$	0.0394	0.0205
$b_n = 3$	0.0470	0.0170

Columns display results for the treatment ($n = 320$) and spillover ($n = 1685$) effects. Rows “ $b_n = k$ ” report SEs for the indicated bandwidths.

comparisons within treated schools, using only a subsample of five schools.

To compare treated and untreated schools, note that for the latter, $Y_i = 0$ for all i by design. Then our estimate of $\hat{\mu}(0)$ for the treatment effect shows that even untreated units are 8.8 percentage points more likely to wear wristbands in treated compared to untreated schools. We obtain standard errors for $\hat{\mu}(0)$ by replacing Z_i and $\hat{\tau}(t, t')$ in (12) with $\mathbf{1}_i(t)\pi_i(t)^{-1}Y_i$ and $\hat{\mu}(0)$. For $b_n = 0, \dots, 3$, the standard errors range from 0.012 to 0.017, all of which indicate a statistically significant effect. Overall, these results indicate that, despite the potential conservativeness of our estimator due to the bias term R_n , they can still deliver reasonable standard errors.

5.2 Monte Carlo

To study the finite sample properties of our estimators, we simulate data from the two response models studied in §3.1 using two models of network formation calibrated to the school data from §5.1. For the linear-in-means model, $Y_i = V_i(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon})$ for

$$V_i(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon}) = \alpha + \beta \frac{\sum_j A_{ij} Y_j}{\sum_j A_{ij}} + \delta \frac{\sum_j A_{ij} D_j}{\sum_j A_{ij}} + D_i \gamma + \varepsilon_i$$

and $(\alpha, \beta, \delta, \gamma) = (-1, 0.8, 1, 1)$. For the complex contagion model, we use the same specification of $V_i(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon})$, except $(\alpha, \beta, \delta, \gamma) = (-1, 1.5, 1, 1)$, and set $Y_i = \mathbf{1}\{V_i(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon}) > 0\}$.

We simulate \mathbf{A} from configuration and random geometric graph (RGG) models. The former is calibrated to the empirical out-degree sequence $(\sum_{j=1}^n A_{ij})_{i=1}^n$ of the schools used in §5.1. This model (approximately) draws an undirected network uniformly at random from the set of all networks with this degree sequence (e.g. Jackson, 2010, Ch. 4.1.4). An RGG is a spatial network where units only link with geographically close alters: $A_{ij} = \mathbf{1}\{\|\rho_i - \rho_j\| \leq r_n\}$ for $\rho_i \stackrel{iid}{\sim} \mathcal{U}([0, 1]^2)$ and $r_n = (\kappa/(\pi n))^2$. Since κ is the limiting expected degree of the model (Penrose, 2003), we set it equal to the average of the empirical out-degree sequence of the schools.

Let $\{\nu_i\}_{i=1}^n \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ be independent of \mathbf{A} . For the configuration model, we take $\varepsilon_i = \nu_i$. For the RGG, we instead use $\varepsilon_i = (\rho_{i1} - 0.5) + \nu_i$, adding the centered first component of i 's "location" ρ_{i1} . Since units with similar ρ_{i1} 's are more likely to form links, this generates unobserved homophily.

Both models yield networks with approximately the same average degree of about 8. For our purposes, the main distinction between them is the APL. The configuration model is theoretically known to have a small APL of logarithmic order in the network size (van der Hofstad, 2016), while the RGG model seems to generate a much larger APL of polynomial order (Friedrich et al., 2013). We choose different bandwidths for the two according to (13).

Note that the value of the peer effect β in both outcome models is actually quite large. The calculations in Appendices A.1 and A.2 suggest that it is larger than what our weak dependence conditions actually require under the configuration model (but not the RGG) since it generates a low APL (and hence approximately exponential neighborhood growth rates). Of course, these conditions are likely stronger than necessary, which may explain the good performance in both designs below.

To show different population sizes, we report results using the largest, two largest, and four largest of the treated schools when calibrating the network models. In all cases, we pool the degree sequences across the schools to treat them as one single network. Following the design in §5.1, we randomly assign treatments to only units classified as eligible in the data with probability 0.5.

We compute estimates and standard errors for the spillover effect $\tau(1, 0)$ in §5.1, whose exposure mapping is $T_i = \mathbf{1}\{\sum_j A_{ij} D_j > 0\}$, again restricting to the population of units with an eligible neighbor. The standard errors use the bandwidth (13). Given the APL differences discussed above, which can also be seen in the tables below, the RGG uses $\tilde{b}_n = \mathcal{L}(\mathbf{A})^{1/3}$, while the configuration model uses $\tilde{b}_n = \mathcal{L}(\mathbf{A})/2$.

Table II. Simulation Results: Configuration Model

# Schools	Linear-in-Means			Complex Contagion		
	1	2	4	1	2	4
$\hat{n}(1)$	199.3	399.7	804.2	199.3	399.7	804.2
$\hat{n}(0)$	151.5	293.1	570.9	151.5	293.1	570.9
$\hat{\tau}(1, 0)$	0.311	0.317	0.313	0.077	0.077	0.080
Our SE	1.310	0.955	0.686	0.109	0.081	0.059
Oracle SE	1.402	0.986	0.694	0.117	0.084	0.060
Our Coverage	0.923	0.938	0.947	0.928	0.937	0.943
Oracle Coverage	0.947	0.950	0.954	0.948	0.950	0.951
Naive Coverage	0.544	0.549	0.563	0.729	0.725	0.730
APL	3.471	3.753	4.070	3.471	3.753	4.070
b_n	2.000	2.000	2.000	2.000	2.000	2.000
Network Size	805	1456	2725	805	1456	2725

Averages over 10k simulations. $\hat{n}(t) = \sum_i \mathbf{1}_i(t)$ is the effective sample size of $\hat{\mu}(t)$. “Coverage” rows display empirical coverage for 95% CIs. “Naive” and “Oracle” respectively correspond to i.i.d. and oracle standard errors.

To illustrate the performance of our bandwidth rule across different networks, we redraw \mathbf{D} , \mathbf{A} , and $\boldsymbol{\varepsilon}$ for each of the 10k simulation draws. This corresponds to a superpopulation design, so we expect our standard errors to yield confidence intervals with coverage close to the nominal rate of 0.95. We report “oracle” standard errors, which correspond to $\text{Var}(\hat{\tau}(1, 0))^{1/2}$, approximated by taking the standard deviation of $\hat{\tau}(1, 0)$ over 10k separate simulation draws. We also report “naive” i.i.d. standard errors to illustrate the degree of dependence in the data.

Tables II and III present results for the configuration and RGG models, respectively. “Our SE” displays the standard error obtained from our variance estimator and “Our Coverage” the empirical coverage of the corresponding two-sided 95% confidence interval. The effective sample size of $\hat{\mu}(t)$ is $\hat{n}(t) = \sum_i \mathbf{1}_i(t)$. “Network Size” is the number of units in the network, whereas the sample size is $\hat{n}(1) + \hat{n}(0)$ since, as in the empirical application, we only use units with an eligible neighbor for overlap.

The results show that the oracle coverage rates are all very close to 0.95, which illustrates the quality of the normal approximation. Our standard errors perform well, with coverage quite close to 0.95. This is in spite of fairly large peer effects β , which can be seen in the magnitudes of the spillover effect estimates. By contrast,

Table III. Simulation Results: RGG Model

# Schools	Linear-in-Means			Complex Contagion		
	1	2	4	1	2	4
$\hat{n}(1)$	210.0	418.2	830.1	210.0	418.2	830.1
$\hat{n}(0)$	155.5	297.8	576.7	155.5	297.8	576.7
$\hat{\tau}(1,0)$	0.702	0.700	0.700	0.088	0.089	0.089
Our SE	1.523	1.107	0.805	0.144	0.105	0.076
Oracle SE	1.601	1.142	0.824	0.152	0.108	0.079
Our Coverage	0.927	0.936	0.936	0.935	0.941	0.943
Oracle Coverage	0.953	0.950	0.947	0.950	0.951	0.954
Naive Coverage	0.519	0.522	0.525	0.626	0.641	0.623
APL	13.996	18.444	24.927	13.996	18.444	24.927
b_n	2.007	3.000	3.000	2.007	3.000	3.000
Network Size	805	1456	2725	805	1456	2725

See table notes of Table II.

naive standard errors are severely anti-conservative.

6 Conclusion

The causal literature on interference typically assumes K -neighborhood exposure mappings are correctly specified. This implies a limited model of interference in which units further than distance K from the ego have no effect on the ego’s response for some small, known K . Such a model is incompatible with those studied in the large theoretical and empirical literature on social interactions, in which units arbitrarily far from the ego can influence the ego’s outcome. This paper proposes a richer model of “approximate neighborhood interference” (ANI) under which treatments assigned to distant alters have a nonzero effect on the ego’s outcome, but the effect declines with network distance. Unlike with existing models of interference, we show that ANI is satisfied by well-known models of social interactions. We also show that ANI is useful for large-sample inference, proving that, under this condition, IPW estimators are consistent and asymptotically normal.

We consider a network HAC variance estimator for inference robust to misspecification of exposure mappings. We show the estimator is biased in a finite population setting. The bias term captures the variance of the unit-level treatment effects, which

generalizes the well-known result on conservative variance estimation under no interference to settings with dependence due to interference. Finally, we propose a new bandwidth rule for the HAC estimator, which trades off bias and variance in a manner that adapts to the growth rate of K -neighborhoods.

A Verifying Assumptions

A.1 Assumption 6

We verify Assumption 6 for networks with polynomial or exponential neighborhood growth rates. We assume a non-degenerate variance ($\sigma_n^{-2} = O(1)$) and $\theta_{n,s} = e^{-c(1-\epsilon)^{-1}s}$ for some $c > 0$ and $\epsilon \in (0, 1)$, which holds for the examples in §3.1. We will choose ϵ in Assumption 6 to be the same as this ϵ .

Polynomial Growth Rate. Suppose

$$\sup_n \max_{i \in \mathcal{N}_n} |\mathcal{N}_{\mathbf{A}}(i, s)| = Cs^d \quad (16)$$

for s sufficiently large and some $C > 0$, $d \geq 1$. Polynomial rates appear to be a property of spatial networks, which are models in which link formation is less likely between spatially distant units.¹⁵ Examples include latent space (Hoff et al., 2002) and RGG models (Penrose, 2003). Appendix A of Leung (2019b) shows that, for the RGG, path distance is of the same order as spatial distance for connected units. Since spatial s -neighborhoods grow polynomially with s , it follows that network s -neighborhoods also grow polynomially, with d determined by the spatial dimension.

We next verify Assumption 6 under this setup. Choose $m_n = n^{1/(\alpha d)}$, $\alpha > 4$. First consider the third term in (11). This is $O(n^{1.5}e^{-cm_n})$, and hence $o(1)$, since m_n is polynomial in n . The second term $n^{-3/2} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, m_n)|^2$ is order $n^{-1/2} m_n^{2d} = n^{2/\alpha - 0.5} = o(1)$. Finally, for the first term in (11), observe that we can bound

$$|\mathcal{H}_n(s, m_n)| \leq \sum_{i=1}^n \sum_{j \in \mathcal{N}_{\mathbf{A}}^2(i, s)} |\mathcal{N}_{\mathbf{A}}(i, m_n)| |\mathcal{N}_{\mathbf{A}}(j, m_n)|. \quad (17)$$

¹⁵This may require replacing the left-hand side of (16) with something weaker such as $\sup_n n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, s)|^p$. The max simplifies the arguments in this section, but we expect higher-order moment conditions of this type can be used to verify the assumptions.

This is $O(n m_n^{2d} s^d) = O(n^{2/\alpha+1} s^d)$, so the first term in (11) is $O(n^{2/\alpha-1} \sum_{s=0}^{\infty} s^d e^{-cs})$.

Exponential Growth Rates. Suppose

$$\sup_n \max_{i \in \mathcal{N}_n} |\mathcal{N}_{\mathbf{A}}(i, s)| = C e^{\beta s} \quad (18)$$

for s sufficiently large and some $C, \beta > 0$. For k -regular tree networks, the left-hand side is exactly $1 + \sum_{m=1}^s k^m$. Inhomogeneous random graphs, a large class of models that includes the Erdős-Rényi and stochastic block models, have a similar property (see Bollobás et al., 2007, proofs of Lemmas 14.2 and 14.3). The usual intuition (Barabási, 2015, Ch. 3.8) is that the average number of units in a unit's 1-neighborhood is the average degree $\delta(\mathbf{A}) = n^{-1} \sum_{i,j} A_{ij}$, so the typical number in its s -neighborhood is approximately $\delta(\mathbf{A})^s$, in which case $\beta \approx \log \delta(\mathbf{A})$ and $C \approx 1$.

We next verify Assumption 6. Choose $m_n = \alpha \beta^{-1} \log n$, $\alpha \in (1.5\beta c^{-1}, 0.5)$, with c from the definition of $\theta_{n,s}$ above. Such an α exists only if $c > 3\beta$, which requires $\tilde{\theta}_{n,s}$ to decay sufficiently fast relative to neighborhood growth rates. The third term in (11) is order $n^{1.5} e^{-cm_n} = n^{1.5-c\alpha\beta^{-1}} = o(1)$. The second term is order $n^{-1/2} e^{2\beta m_n} = n^{2\alpha-0.5} = o(1)$. Finally, using (17), the first term of (11) is conservatively order

$$n^{-1} e^{2\beta m_n} \sum_{s=0}^n e^{\beta s} e^{-cs} = n^{2\alpha-1} \sum_{s=0}^n e^{(\beta-c)s} = o(1).$$

A.2 Choice of Bandwidth

We use a mix of formal and heuristic arguments to show that our bandwidth (13) satisfies Assumption 7(b)–(d) under different neighborhood growth rates. As in §A.1, we suppose that $\sigma_n^{-2} = O(1)$ and $\theta_{n,s} = e^{-c(1-\epsilon)^{-1}s}$ for some $c > 0$ and $\epsilon \in (0, 1)$.

Polynomial Growth Rates. We first consider the case in which s -neighborhood sizes grow polynomially with s in the sense of (16). Let $\nabla(\mathbf{A})$ be the diameter of \mathbf{A} , which is the maximum path length between pairs in the largest component. Then

$$\max_i |\mathcal{N}_{\mathbf{A}}(i, \nabla(\mathbf{A}))| = \alpha n, \quad (19)$$

where α is the fraction of units in the largest component of \mathbf{A} . Most real-world networks have a “giant component” (Barabási, 2015) meaning $\alpha/n \rightarrow c \in (0, \infty)$.

Furthermore, the left-hand side of (19) is $C\nabla(\mathbf{A})^d$ under (16). Hence, $\nabla(\mathbf{A}) = O(n^{1/d})$. This is a well-known heuristic argument for the asymptotic behavior of the diameter or average path length (Barabási, 2015, Ch. 3.8).¹⁶

Our bandwidth rule (13) is based on the APL $\mathcal{L}(\mathbf{A})$ rather than the diameter since the former is considered a more robust measurement of network width.¹⁷ While the heuristics above pertain to the diameter, the derived rate is usually more accurate for the APL. As written in Barabási (2015), Ch. 3.8, “for most networks [these heuristics offer] a better approximation to the average distance between two randomly chosen nodes, than to [the diameter]. This is because [the diameter] is often dominated by a few extreme paths, while [the APL] is averaged over all node pairs, a process that suppresses [sic] the fluctuations.” Hence, our calculations below take $\mathcal{L}(\mathbf{A}) \approx n^{1/d}$.

We can now verify Assumption 7(b)–(d) (the argument for (a) is similar). For (b), using our bandwidth $b_n = \mathcal{L}(\mathbf{A})^{1/3} \approx n^{1/(3d)}$, $n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, b_n)| \approx b_n^d \approx n^{1/3} = o(n^{1/2})$. For (c), $n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, b_n)|^2 \approx b_n^{2d} \approx n^{2/3} = o(n)$. Part (d) is the same as the first part of (11) but with b_n and $\mathcal{J}_n(s, \cdot)$ in place of m_n and $\mathcal{H}_n(s, \cdot)$. Note that (17) applies with $\mathcal{J}_n(s, \cdot)$ in place of $\mathcal{H}_n(s, \cdot)$, so

$$\frac{1}{n^2} \sum_{s=0}^n |\mathcal{J}_n(s, b_n)| \tilde{\theta}_{n,s} \approx n^{-1} b_n^{2d} \sum_{s=0}^n s^d e^{-c(1-\epsilon)^{-1}s} = O(n^{-1/3}). \quad (20)$$

Exponential Growth Rates. Now suppose (18), and take the typical case of $\beta = \log \delta(\mathbf{A})$ and $C = 1$ discussed after (18). The heuristics following (19) yield $\nabla(\mathbf{A}) = \log n / \log \delta(\mathbf{A})$, as in (3.18) of Barabási (2015). As discussed above, this heuristic is actually more accurate for the APL, so we take $\mathcal{L}(\mathbf{A}) \approx \log n / \log \delta(\mathbf{A})$.¹⁸

Choosing b_n according to (13) yields $b_n \approx 0.5 \log n / \log \delta(\mathbf{A})$. Strictly speaking, we will actually need $b_n \approx \alpha \log n / \log \delta(\mathbf{A})$ for some $\alpha < 0.5$, which we will assume next, but since (13) rounds to the nearest integer, there is no difference setting $\alpha = 0.5$ in practice. For Assumption 7(b), $n^{-1} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}(i, b_n)| \leq e^{b_n \log \delta(\mathbf{A})} \approx e^{\alpha \log n} = o(n^{1/2})$, and (c) is similar. For (d), if $(1 - \epsilon)^{-1}c > \beta = \log \delta(\mathbf{A})$, which is weaker than the requirement $c > 3\beta$ in §A.1, then as in (20),

$$\frac{1}{n^2} \sum_{s=0}^n |\mathcal{J}_n(s, b_n)| \tilde{\theta}_{n,s} \approx n^{-1} e^{2b_n \log \delta(\mathbf{A})} \sum_{s=0}^n e^{s \log \delta(\mathbf{A})} e^{-c(1-\epsilon)^{-1}s} = O(n^{2\alpha-1}).$$

¹⁶See Friedrich et al. (2013) for a formal argument for the RGG.

¹⁷I thank a referee for this suggestion.

¹⁸For a formal argument for inhomogeneous graphs, see Theorem 3.14 of Bollobás et al. (2007).

B AS Variance Estimator

We provide a theoretical comparison of our variance estimator and that of [Aronow and Samii \(2017\)](#) (henceforth AS) and provide simulation evidence on differences in conservativeness. Since their estimator is only valid under correctly specified exposure mappings, we now assume (2). The estimand is the variance

$$\begin{aligned}\sigma_n^2 &= \text{Var}(\sqrt{n}\hat{\tau}(t, t')) = V(t) + V(t') + 2C(t, t'), \quad \text{where} \\ V(t) &= \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i(t)^2 \frac{1 - \pi_i(t)}{\pi_i(t)} + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \tilde{Y}_i(t) \tilde{Y}_j(t) \frac{\pi_{ij}(t, t) - \pi_i(t)\pi_j(t)}{\pi_i(t)\pi_j(t)}, \\ C(t, t') &= -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \tilde{Y}_i(t) \tilde{Y}_j(t') \frac{\pi_{ij}(t, t') - \pi_i(t)\pi_j(t')}{\pi_i(t)\pi_j(t')},\end{aligned}$$

and $\pi_{ij}(t, t') = \mathbf{E}[\mathbf{1}_i(t)\mathbf{1}_j(t')]$. The AS estimator given in their equation (11) is

$$\begin{aligned}\hat{\sigma}_{AS}^2 &= \hat{V}(t) + \hat{V}(t') + 2\hat{C}(t, t'), \quad \text{where} \\ \hat{V}(t) &= \frac{1}{n} \sum_{i=1}^n \frac{Y_i^2 \mathbf{1}_i(t)}{\pi_i(t)} \frac{1 - \pi_i(t)}{\pi_i(t)} + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \frac{Y_i \mathbf{1}_i(t) Y_j \mathbf{1}_j(t)}{\pi_{ij}(t, t)} \frac{\pi_{ij}(t, t) - \pi_i(t)\pi_j(t)}{\pi_i(t)\pi_j(t)} \mathbf{1}_{\{\pi_{ij}(t, t) \neq 0\}} \\ &\quad + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \left(\frac{Y_i^2 \mathbf{1}_i(t)}{2\pi_i(t)} + \frac{Y_j^2 \mathbf{1}_j(t)}{2\pi_j(t)} \right) \mathbf{1}_{\{\pi_{ij}(t, t) = 0\}}, \\ \hat{C}(t, t') &= -\frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} \frac{Y_i \mathbf{1}_i(t) Y_j \mathbf{1}_j(t')}{\pi_{ij}(t, t')} \frac{\pi_{ij}(t, t') - \pi_i(t)\pi_j(t')}{\pi_i(t)\pi_j(t')} \mathbf{1}_{\{\pi_{ij}(t, t') \neq 0\}} \\ &\quad + \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \left(\frac{Y_i^2 \mathbf{1}_i(t)}{2\pi_i(t)} + \frac{Y_j^2 \mathbf{1}_j(t')}{2\pi_j(t')} \right) \mathbf{1}_{\{\pi_{ij}(t, t') = 0\}}.\end{aligned}$$

Noting that $\pi_{ii}(t, t') = 0$, this is conservative for σ_n^2 with bias

$$\begin{aligned}R_{n,AS} &\equiv \mathbf{E}[\hat{\sigma}_{AS}^2] - \sigma_n^2 = \frac{1}{n} \sum_{i=1}^n (\tilde{Y}_i(t) - \tilde{Y}_i(t'))^2 \\ &\quad + 0.5 \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} ((\tilde{Y}_i(t) + \tilde{Y}_j(t))^2 \mathbf{1}_{\{\pi_{ij}(t, t) = 0\}} + 2(\tilde{Y}_i(t) - \tilde{Y}_j(t'))^2 \mathbf{1}_{\{\pi_{ij}(t, t') = 0\}} \\ &\quad \quad \quad + (\tilde{Y}_i(t') + \tilde{Y}_j(t'))^2 \mathbf{1}_{\{\pi_{ij}(t', t') = 0\}}).\end{aligned}$$

The asymptotic bias of our estimator $\hat{\sigma}^2$ is R_n given in (14) with $b_n = 2K$ since

we are assuming correct specification. It does not appear that R_n and $R_{n,AS}$ can generally be ordered. However, we can consider a few instructive special cases.

First consider **no interference** ($K = 0$). Then $R_n = n^{-1} \sum_{i=1}^n (\tau_i(t, t') - \tau(t, t'))^2$. For the AS estimator, since $\pi_{ij}(t, t') \neq 0$ for all $i \neq j$, $R_{n,AS} = n^{-1} \sum_{i=1}^n \tau_i(t, t')^2$, which is strictly larger than R_n whenever the treatment effect $\tau(t, t')$ is nonzero. Indeed $R_{n,AS}$ is larger by the *square* of the treatment effect, so that for large effects and relatively small sample sizes, $\hat{\sigma}_{AS}^2$ can be substantially more conservative than $\hat{\sigma}^2$. Next consider **homogeneous unit-level exposure effects**, which corresponds to $\tau_i(t, t') = \tau(t, t')$ for all i . In this case, $R_n = 0$, whereas $R_{n,AS}$ does not simplify in general and is strictly positive if, for example, $\tau_i(t, t') \neq 0$ for all i . Finally, consider the **has-treated-neighbor spillover effect** in §5, where $t = 1$, $t' = 0$, and $T_i = \mathbf{1}\{\sum_j A_{ij} D_j > 0\}$. Call a unit i “eligible” if $\mathbf{E}[D_i] > 0$. Let \mathcal{E}_{ij} be the event that the set of i ’s eligible neighbors in \mathbf{A} equals the set of eligible neighbors i has in common with j . Then $\pi_{ij}(1, 0) = 0$ if and only if \mathcal{E}_{ij} occurs, and $\pi_{ij}(d, d) \neq 0$ for all $d \in \{0, 1\}$ and units i and j who have at least one eligible neighbor. Thus, assuming, as in §5, that the population consists of all units with eligible neighbors,

$$R_{n,AS} = \frac{1}{n} \sum_{i=1}^n \tau_i(1, 0)^2 + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} (\tilde{Y}_i(1) - \tilde{Y}_j(0))^2 \mathbf{1}\{\mathcal{E}_{ij}\}, \quad \text{whereas} \quad (21)$$

$$R_n = \frac{1}{n} \sum_{i=1}^n (\tau_i(1, 0) - \tau(1, 0))^2 + \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} (\tau_i(1, 0) - \tau(1, 0))(\tau_j(1, 0) - \tau(1, 0)) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq 2\}. \quad (22)$$

While the first terms of both expressions can be ordered as in the no-interference case, the second terms cannot. As discussed in §4, the second term in (22) accounts for dependence between unit-level exposure effects due to interference, but the second term in (21) does not have a clear interpretation. However, we can provide simulation evidence on magnitudes.

We simulate networks calibrated to the data in the empirical application, as in §5.2, but use a different outcome model to impose correct specification: $\tilde{Y}_i(0) = \varepsilon_i + \sum_j A_{ij} \varepsilon_j / \sum_j A_{ij}$ and $\tilde{Y}_i(1) = \beta_i + \tilde{Y}_i(0)$, where $\{\varepsilon_i\}_{i=1}^n \stackrel{iid}{\sim} \mathcal{N}(0, 1)$ is independent of $\{\beta_i\}_{i=1}^n \stackrel{iid}{\sim} \mathcal{N}(1, 1)$. The fraction in $\tilde{Y}_i(0)$ can be interpreted as exogenous peer effects in unobservables and serves to generate network autocorrelation in baseline

outcomes. As in §5.2, the sample is the set of units with eligible neighbors. The left half of Table IV reports the average values of $R_{n,AS}$ and R_n across 5000 simulation draws. We see that $R_{n,AS}$ is about 3–5 times larger than R_n on average.

In the previous design, $\tau_i(1, 0) = \beta_i$ is independent across i . To introduce network autocorrelation in unit-level exposure effects, we instead take $\tilde{Y}_i(1) = \beta_i + \sum_j A_{ij}\beta_j / \sum_j A_{ij} + \tilde{Y}_i(0)$. The right half of Table IV reports the results. For the configuration model, $R_{n,AS}$ is twice as large as R_n on average, and for the RGG model, it is about 25 percent larger.

Table IV. Comparison of Average Bias

	Independent Effects				Autocorrelated Effects			
	Configuration		RGG Model		Configuration		RGG Model	
# Schools	1	2	1	2	1	2	1	2
$R_{n,AS}$	9.4	10.2	7.1	7.6	154.6	158.7	170.0	180.0
R_n	1.9	1.9	2.0	2.0	67.6	73.8	136.6	144.1
n	350.8	692.8	365.5	716.1	350.8	692.8	365.5	716.1

Cells are averages over 5k simulations. $n = \#$ units with eligible neighbors.

C Proofs

PROOF OF PROPOSITION 1. The reduced form of the linear-in-means model is

$$\begin{aligned}
 Y_i(\mathbf{D}) = & \alpha + D_i\gamma + \varepsilon_i + \mathbf{1} \left\{ \sum_j A_{ij} > 0 \right\} \times \\
 & \left(\frac{\alpha\beta}{1-\beta} + \gamma\beta \sum_{k=0}^{\infty} \beta^k \left(\sum_{j_1=1}^n \frac{A_{ij_1}}{\sum_{\ell} A_{i\ell}} \sum_{j_2=1}^n \frac{A_{j_1j_2}}{\sum_{\ell} A_{j_1\ell}} \cdots \sum_{j_{k+1}=1}^n \frac{A_{j_kj_{k+1}}}{\sum_{\ell} A_{j_k\ell}} D_{j_{k+1}} \right) \right. \\
 & \left. + \sum_{k=1}^{\infty} \beta^k \left(\sum_{j_1=1}^n \frac{A_{ij_1}}{\sum_{\ell} A_{i\ell}} \sum_{j_2=1}^n \frac{A_{j_1j_2}}{\sum_{\ell} A_{j_1\ell}} \cdots \sum_{j_k=1}^n \frac{A_{j_{k-1}j_k}}{\sum_{\ell} A_{j_{k-1}\ell}} \varepsilon_{j_k} \right) \right).
 \end{aligned}$$

Consider a counterfactual linear-in-means model in which the set of units is $\mathcal{N}_{\mathbf{A}}(i, s)$ rather than \mathcal{N}_n . That is, outcomes are realized according to (5) but with primitives $(\mathbf{D}_{\mathcal{N}_{\mathbf{A}}(i, s)}, \mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i, s)}, \boldsymbol{\varepsilon}_{\mathcal{N}_{\mathbf{A}}(i, s)})$ rather than $(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon})$, where $\boldsymbol{\varepsilon} = (\varepsilon_i)_{i=1}^n$ and $\boldsymbol{\varepsilon}_{\mathcal{N}_{\mathbf{A}}(i, s)} = (\varepsilon_i : i \in \mathcal{N}_{\mathbf{A}}(i, s))$. Let $Y_i^{(s)}(\mathbf{D})$ be unit i 's outcome in this counterfactual model. To

compare $Y_i^{(s)}(\mathbf{D})$ with $Y_i(\mathbf{D})$, consider the k th term in the first series of the previous equation (the term multiplying $\gamma\beta \cdot \beta^k$). We can rewrite this as $\sum_{j=1}^n \omega_{ij}^k D_j$, where ω_{ij}^k is the following weighted sum of all walks of length $k+1$ from i to j :

$$\omega_{ij}^k = \sum_{j_1=1}^n \sum_{j_2=1}^n \cdots \sum_{j_k=1}^n \frac{A_{ij_1}}{\sum_{\ell} A_{i\ell}} \frac{A_{j_1 j_2}}{\sum_{\ell} A_{j_1 \ell}} \cdots \frac{A_{j_k j}}{\sum_{\ell} A_{j \ell}}.$$

In going from $Y_i(\mathbf{D})$ to $Y_i^{(s)}(\mathbf{D})$, we lose terms in ω_{ij}^k involving walks that traverse paths of length greater than s . We can conservatively bound this loss by including *all* walks with length exceeding s in the loss. Then since $\sum_{j=1}^n \omega_{ij}^k = 1$ and D_j is binary,

$$|Y_i(\mathbf{D}) - Y_i^{(s)}(\mathbf{D})| \leq \gamma\beta \sum_{k=s+1}^{\infty} |\beta|^k + \sup_i |\varepsilon_i| \sum_{k=s+1}^{\infty} |\beta|^k$$

for $s \geq 1$. Since ε_i is uniformly bounded, the right-hand side is bounded by a constant times $|\beta|^s$. Furthermore, the argument above is the same if we replace \mathbf{D} with $\mathbf{D}^{(i,s)}$, the latter defined prior to Assumption 4. Since $Y_i^{(s)}(\mathbf{D}) = Y_i^{(s)}(\mathbf{D}^{(i,s)})$, the result follows from the triangle inequality. \blacksquare

PROOF OF PROPOSITION 2. The arguments that follow borrow ideas from the proofs of Proposition 1 of [Xu and Lee \(2015\)](#) and Theorem 6.1 of [Leung \(2019b\)](#). Let $\tilde{\mathbf{G}}$ be the directed network with ij entry $A_{ij}\varphi_j$, C_i the set of units in the strongly connected component of $\tilde{\mathbf{G}}$ containing i , and

$$C_i^+ = C_i \cup \{k \in \mathcal{N}_n : \exists j \in C_i \text{ such that } A_{jk}(1 - \varphi_i) = 1\}.$$

Let $Y_i(\mathbf{D})$ be i 's outcome under the complex contagion model with n units. As in the proof of our Proposition 1, consider a counterfactual model in which the set of units is $\mathcal{N}_{\mathbf{A}}(i, s)$ instead of \mathcal{N}_n , meaning outcomes are realized according to the same model and same initial condition \mathbf{Y}^0 but with primitives $(\mathbf{D}_{\mathcal{N}_{\mathbf{A}}(i,s)}, \mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i,s)}, \boldsymbol{\varepsilon}_{\mathcal{N}_{\mathbf{A}}(i,s)})$ instead of $(\mathbf{D}, \mathbf{A}, \boldsymbol{\varepsilon})$. Let $Y_i^{(s)}(\mathbf{D})$ be i 's outcome in this counterfactual model. Key to our argument is the fact that

$$Y_i(\mathbf{D}) = Y_i^{(s)}(\mathbf{D}) \quad \text{if} \quad \mathbf{A}_{C_i^+} \subseteq \mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i,s)}, \quad (23)$$

where $\mathbf{A}_{C_i^+} \subseteq \mathbf{A}_{\mathcal{N}_{\mathbf{A}}(i,s)}$ means the former is a subnetwork of the latter.

To show (23), let $Y_i^t(\mathbf{D})$ be i 's outcome at time t of the dynamic process described prior to the proposition for the model with n units. If $\varphi_i = 0$, then i has a dominant strategy $Y_i^*(\mathbf{D})$, so $Y_i^t(\mathbf{D}) = Y_i^*(\mathbf{D})$ for all $t > 0$. If instead $\varphi_i = 1$, then i 's outcome may potentially change at any period $t > 0$ in the process, depending on the outcomes of neighboring units at $t - 1$. Consider any path in \mathbf{A} connecting units i and j . If $\varphi_k = 1$ for all units k along that path, then unit i 's outcome may change at any period t in the dynamic process, depending on the outcome of j at some prior period. However, if for all such paths, there exists some unit k along that path such that $\varphi_k = 0$, then unit i 's outcome will never be affected by unit's j outcome at any past period. Now, if $j \notin C_i^+$, then by construction, there exists such a unit k along any path connecting i and j . Therefore, $Y_i(\mathbf{D})$ is invariant to the removal of units $\mathcal{N}_n \setminus C_i^+$ from the model in the sense of (23). It follows that

$$\begin{aligned} \mathbf{E}[|Y_i(\mathbf{D}) - Y_i^{(s)}(\mathbf{D})|] &\leq \mathbf{P}(Y_i(\mathbf{D}) \neq Y_i^{(s)}(\mathbf{D})) \leq \mathbf{P}(\mathbf{A}_{C_i^+} \not\subseteq \mathbf{A}_{\mathcal{N}_A(i,s)}) \\ &\leq \sum_{j_1 \neq \dots \neq j_{s-1}} A_{ij_1} \mathbf{E}[\varphi_{j_1}] A_{j_1 j_2} \mathbf{E}[\varphi_{j_2}] \cdots A_{j_{s-3} j_{s-2}} \mathbf{E}[\varphi_{j_{s-2}}] = \sum_j (\mathbf{G}^{s-1})_{ij}. \end{aligned} \quad (24)$$

The third inequality follows from the union bound, independence of treatments, and the fact that $\mathbf{A}_{C_i^+} \not\subseteq \mathbf{A}_{\mathcal{N}_A(i,s)}$ implies there exists a path of length at least $s - 1$ starting from i such that $\varphi_k = 1$ for all units k on that path.

For all s , $|Y_i(\mathbf{D}) - Y_i^{(s)}(\mathbf{D})| \leq 1$. For $s - 1 \geq \bar{s}$, (24) $\leq \rho_n(\bar{s})^{s-1}$. Furthermore, the argument holds if we replace \mathbf{D} with $\mathbf{D}^{(i,s)}$, the latter defined prior to Assumption 4. Since $Y_i^{(s)}(\mathbf{D}) = Y_i^{(s)}(\mathbf{D}^{(i,s)})$, the result follows from the triangle inequality. \blacksquare

PROOF OF THEOREM 1. Let either $h, h' \in \mathbb{N}$, $f \in \mathcal{L}_h$, and $f' \in \mathcal{L}_{h'}$, or $h = h' = 1$ and $f = f' = i^*$. Let $s > 0$ and $(H, H') \in \mathcal{P}_n(h, h'; s)$. Define $\xi = f(\mathbf{Z}_H)$ and $\zeta = f'(\mathbf{Z}_{H'})$. Let \mathbf{D}' , \mathbf{D}'' each be independent copies of \mathbf{D} . Define $\mathbf{D}^{(i,s,\xi)} = (\mathbf{D}_{\mathcal{N}_A(i,s)}, \mathbf{D}'_{-\mathcal{N}_A(i,s)})$, $\mathbf{D}^{(i,s,\zeta)} = (\mathbf{D}_{\mathcal{N}_A(i,s)}, \mathbf{D}''_{-\mathcal{N}_A(i,s)})$, and

$$\begin{aligned} Z_i^{(s,\xi)} &= Y_i(\mathbf{D}^{(i,s,\xi)}) \left(\frac{\mathbf{1}\{T(i, \mathbf{D}^{(i,s,\xi)}, \mathbf{A}) = t\}}{\pi_i(t)} - \frac{\mathbf{1}\{T(i, \mathbf{D}^{(i,s,\xi)}, \mathbf{A}) = t'\}}{\pi_i(t')} \right), \\ Z_i^{(s,\zeta)} &= Y_i(\mathbf{D}^{(i,s,\zeta)}) \left(\frac{\mathbf{1}\{T(i, \mathbf{D}^{(i,s,\zeta)}, \mathbf{A}) = t\}}{\pi_i(t)} - \frac{\mathbf{1}\{T(i, \mathbf{D}^{(i,s,\zeta)}, \mathbf{A}) = t'\}}{\pi_i(t')} \right). \end{aligned}$$

Finally, let $\xi^{(s)} = f((Z_i^{(s,\xi)} : i \in H))$ and $\zeta^{(s)} = f'((Z_i^{(s,\zeta)} : i \in H'))$.

Since Z_i is uniformly bounded by Assumptions 2 and 3, $|\text{Cov}(\xi, \zeta)| \leq 2\|f\|_\infty\|f'\|_\infty$, so for $s \leq 2\max\{K, 1\}$, we have $|\text{Cov}(\xi, \zeta)| \leq \psi_{h,h'}(f, f')$. Now consider $s > 2\max\{K, 1\}$, so that $\ell_{\mathbf{A}}(H, H') > 2\max\{K, 1\}$. By Assumption 1, $(Z_i^{(\lfloor s/2 \rfloor, \xi)} : i \in H) \perp\!\!\!\perp (Z_j^{(\lfloor s/2 \rfloor, \zeta)} : j \in H)$. Then

$$\begin{aligned} |\text{Cov}(\xi, \zeta)| &\leq |\text{Cov}(\xi - \xi^{(\lfloor s/2 \rfloor)}, \zeta)| + |\text{Cov}(\xi^{(\lfloor s/2 \rfloor)}, \zeta - \zeta^{(\lfloor s/2 \rfloor)})| \\ &\leq 2\|f'\|_\infty \mathbf{E}[|\xi - \xi^{(\lfloor s/2 \rfloor)}|] + 2\|f\|_\infty \mathbf{E}[|\zeta - \zeta^{(\lfloor s/2 \rfloor)}|] \\ &\leq 2(h\|f'\|_\infty \text{Lip}(f) + h'\|f\|_\infty \text{Lip}(f')) \theta_{n, \lfloor s/2 \rfloor}. \end{aligned}$$

The last line uses the fact that, by Assumption 1,

$$\frac{\mathbf{1}\{T(i, \mathbf{D}^{(i, \lfloor s/2 \rfloor, \xi)}, \mathbf{A}) = t\}}{\pi_i(t)} = \frac{\mathbf{1}\{T(i, \mathbf{D}^{(i, \lfloor s/2 \rfloor, \zeta)}, \mathbf{A}) = t\}}{\pi_i(t)} = \frac{\mathbf{1}\{T(i, \mathbf{D}, \mathbf{A}) = t\}}{\pi_i(t)},$$

for any $t \in \mathcal{T}$, and by Assumption 4, $\max_{i \in \mathcal{N}_n} \mathbf{E}[|Y_i(\mathbf{D}) - Y_i(\mathbf{D}^{(i, \lfloor s/2 \rfloor, \xi)})|] \leq \theta_{n, \lfloor s/2 \rfloor}$. ■

PROOF OF THEOREM 2. Since $\mathbf{E}[\hat{\tau}(t, t')] = \tau(t, t')$, we only need to show that $\text{Var}(\hat{\tau}(t, t')) = o(1)$. Since treatments are independent across units, by Assumption 4, $\text{Cov}(Z_i, Z_j) = 0$ if $\ell_{\mathbf{A}}(i, j) > n - 1$. Hence,

$$\text{Var}(\hat{\tau}(t, t')) = \frac{1}{n^2} \sum_{i=1}^n \text{Var}(Z_i) + \sum_{s=1}^{n-1} \frac{1}{n^2} \sum_{i=1}^n \sum_{j \neq i} \mathbf{1}\{\ell_{\mathbf{A}}(i, j) = s\} \text{Cov}(Z_i, Z_j).$$

Using Theorem 1 and uniform boundedness of Z_i (Assumptions 2 and 3), the right-hand side is bounded above by $C(n^{-1} + n^{-2} \sum_{s=1}^{n-1} \tilde{\theta}_{n,s} \sum_{i=1}^n |\mathcal{N}_{\mathbf{A}}^{\hat{\sigma}}(i, s)|)$ for some universal $C > 0$, and this is $o(1)$ by Assumption 5. ■

PROOF OF THEOREM 3. Apply our Theorem 1 and Theorem 3.2 of [Kojevnikov et al. \(2019\)](#). ■

PROOF OF THEOREM 4. Observe that (15) follows from Proposition 4.1 of [Kojevnikov et al. \(forthcoming\)](#) since our Assumption 7 implies their Assumption 4.1. To establish (14), note that there are two parts of alleged $o_p(1)$ term in (14). The first is due to replacing $\hat{\tau}(t, t')$ with $\tau(t, t')$ in the formula for $\hat{\sigma}^2$. This replacement

creates a remainder term of the form

$$r_n = \frac{2}{n} \sum_{i=1}^n \sum_{j=1}^n (Z_i - \tau(t, t')) (\tau(t, t') - \hat{\tau}(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\} \\ + (\tau(t, t') - \hat{\tau}(t, t'))^2 \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}.$$

Since Z_i is uniformly bounded (Assumptions 2 and 3), for some $C > 0$ and any n ,

$$|r_n| \leq C |\tau(t, t') - \hat{\tau}(t, t')| \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}.$$

The summation term is equal to $M_n(b_n, 1)$. The term in the absolute value is $O_p(n^{-1/2})$ since Assumption 7(a) implies that $\text{Var}(\hat{\tau}(t, t')) = O(n^{-1})$ (see the proof of Theorem 2). Hence, the previous display is $o_p(1)$ by Assumption 7(b).

The remaining parts of the alleged $o_p(1)$ term in (14) are the cross-terms

$$\frac{2}{n} \sum_{i=1}^n \sum_{j=1}^n (Z_i - \tau_i(t, t')) (\tau_j(t, t') - \tau(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}.$$

We show this is $o_p(1)$. For $W_i = \sum_{j=1}^n (\tau_j(t, t') - \tau(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\}$,

$$\mathbf{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n (Z_i - \tau_i(t, t')) (\tau_j(t, t') - \tau(t, t')) \mathbf{1}\{\ell_{\mathbf{A}}(i, j) \leq b_n\} \right| \right] \\ \leq \mathbf{E} \left[\left(\frac{1}{n} \sum_{i=1}^n (Z_i - \tau_i(t, t')) W_i \right)^2 \right]^{1/2} \\ \leq \left(\frac{1}{n^2} \sum_{i=1}^n \text{Var}(Z_i) W_i^2 + C \frac{1}{n^2} \sum_{s=0}^n \tilde{\theta}_{n,s} \sum_{i=1}^n \sum_{j \neq i}^n \mathbf{1}\{\ell_{\mathbf{A}}(i, j) = s\} |W_i W_j| \right)^{1/2}$$

for some $C > 0$ by Theorem 1. Since Z_i is uniformly bounded, for some $C' > 0$, $n^{-2} \sum_{i=1}^n \text{Var}(Z_i) W_i^2 \leq C' n^{-1} M_n(b_n, 2)$, which is $o(1)$ by Assumption 7(c). Likewise,

$$\frac{1}{n^2} \sum_{s=0}^n \tilde{\theta}_{n,s} \sum_{i=1}^n \sum_{j \neq i}^n \mathbf{1}\{\ell_{\mathbf{A}}(i, j) = s\} |W_i W_j| \leq \frac{C''}{n^2} \sum_{s=0}^n \tilde{\theta}_{n,s} \mathcal{J}_n(s, b_n)$$

for some $C'' > 0$, and this is $o(1)$ by Assumption 7(d). ■

References

- Abadie, A., S. Athey, G. Imbens, and J. Wooldridge**, “Sampling-Based vs. Design-Based Uncertainty in Regression Analysis,” *Econometrica*, 2020, *88* (1), 265–296.
- Acemoglu, D., C. Garcia-Jimeno, and J. Robinson**, “State Capacity and Economic Development: A Network Approach,” *American Economic Review*, 2015, *105* (8), 2364–2409.
- Aronow, P. and C. Samii**, “Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment,” *Annals of Applied Statistics*, 2017, *11* (4), 1912–1947.
- Athey, Susan, D. Eckles, and G. Imbens**, “Exact P-Values for Network Interference,” *Journal of the American Statistical Association*, 2018, *113* (521), 230–240.
- Baird, S., J. Bohren, C. McIntosh, and B. Özler**, “Optimal Design of Experiments in the Presence of Interference,” *Review of Economics and Statistics*, 2018, *100* (5), 844–860.
- Bandiera, O., I. Barankay, and I. Rasul**, “Social Connections and Incentives in the Workplace: Evidence from Personnel Data,” *Econometrica*, 2009, *77* (4), 1047–1094.
- Barabási, A.**, *Network Science*, Cambridge University Press, 2015.
- Basse, G., A. Feller, and P. Toulis**, “Randomization Tests of Causal Effects Under Interference,” *Biometrika*, 2019, *106* (2), 487–494.
- Bollobás, B., S. Janson, and O. Riordan**, “The Phase Transition in Inhomogeneous Random Graphs,” *Random Structures and Algorithms*, 2007, *31* (1), 3–122.
- Bond, R., C. Fariss, J. Jones, A. Kramer, C. Marlow, Jaime S., and James F.**, “A 61-Million-Person Experiment in Social Influence and Political Mobilization,” *Nature*, 2012, *489* (7415), 295–298.
- Bramoullé, Y., H. Djebbari, and B. Fortin**, “Identification of Peer Effects Through Social Networks,” *Journal of Econometrics*, 2009, *150* (1), 41–55.

- Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman**, “Understanding Mechanisms Underlying Peer Effects: Evidence From a Field Experiment on Financial Decisions,” *Econometrica*, 2014, *82* (4), 1273–1301.
- Cai, J., A. De Janvry, and E. Sadoulet**, “Social Networks and the Decision to Insure,” *American Economic Journal: Applied Economics*, 2015, *7* (2), 81–108.
- Chin, A.**, “Central Limit Theorems via Stein’s Method for Randomized Experiments Under Interference,” *arXiv preprint arXiv:1804.03105*, 2019.
- Choi, D.**, “Estimation of Monotone Treatment Effects in Network Experiments,” *Journal of the American Statistical Association*, 2017, *112* (519), 1147–1155.
- , “Using Exposure Mappings as Side Information in Experiments with Interference,” *arXiv preprint arXiv:1806.11219*, 2018.
- Doukhan, P. and S. Louhichi**, “A New Weak Dependence Condition and Applications to Moment Inequalities,” *Stochastic Processes and their Applications*, 1999, *84* (2), 313–342.
- Eckles, D., B. Karrer, and J. Ugander**, “Design and Analysis of Experiments in Networks: Reducing Bias from Interference,” *Journal of Causal Inference*, 2017, *5* (1).
- Forastiere, L., E. Airoidi, and F. Mealli**, “Identification and Estimation of Treatment and Interference Effects in Observational Studies on Networks,” *Journal of the American Statistical Association*, 2020, pp. 1–18.
- Friedrich, T., T. Sauerwald, and A. Stauffer**, “Diameter and Broadcast Time of Random Geometric Graphs in Arbitrary Dimensions,” *Algorithmica*, 2013, *67* (1), 65–88.
- Granovetter, M.**, “Threshold Models of Collective Behavior,” *American Journal of Sociology*, 1978, *83* (6), 1420–1443.
- Guilbeault, D., J. Becker, and D. Centola**, “Complex Contagions: A Decade in Review,” in “Complex Spreading Phenomena in Social Systems,” Springer, 2018, pp. 3–25.

- He, X. and K. Song**, “Measuring Diffusion over a Large Network,” *arXiv preprint arXiv:1812.04195*, 2021.
- Hirano, K., G. Imbens, and G. Ridder**, “Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score,” *Econometrica*, 2003, 71 (4), 1161–1189.
- Hoff, P., A. Raftery, and M. Handcock**, “Latent Space Approaches to Social Network Analysis,” *Journal of the American Statistical Association*, 2002, 97 (460), 1090–1098.
- Hudgens, M. and M. Halloran**, “Toward Causal Inference with Interference,” *Journal of the American Statistical Association*, 2008, 103 (482), 832–842.
- Imbens, G.**, “The Role of the Propensity Score in Estimating Dose-Response Functions,” *Biometrika*, 2000, 87 (3), 706–710.
- Imbens, G and D. Rubin**, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press, 2015.
- Imbens, G. and K. Menzel**, “A Causal Bootstrap,” *arXiv preprint arXiv:1807.02737*, 2019.
- Jackson, M.**, *Social and Economic Networks*, Princeton University Press, 2010.
- Jenish, N. and I. Prucha**, “Central Limit Theorems and Uniform Laws of Large Numbers for Arrays of Random Fields,” *Journal of Econometrics*, 2009, 150 (1), 86–98.
- Kojevnikov, D.**, “The Bootstrap for Network Dependent Processes,” *arXiv preprint arXiv:2101.12312*, 2021.
- , **V. Marmer, and K. Song**, “Limit Theorems for Network Dependent Random Variables,” *arXiv preprint arXiv:1903.01059*, 2019.
- , – , **and** – , “Limit Theorems for Network Dependent Random Variables,” *Journal of Econometrics*, forthcoming.
- Lazzati, N.**, “Treatment Response with Social Interactions: Partial Identification via Monotone Comparative Statics,” *Quantitative Economics*, 2015, 6 (1), 49–83.

- Leung, M.**, “Causal Inference Under Approximate Neighborhood Interference,” *arXiv preprint arXiv:1911.07085v1*, 2019.
- , “Inference in Models of Discrete Choice with Social Interactions Using Network Data,” *arXiv preprint arXiv:1911.07106v1*, 2019.
- , “Treatment and Spillover Effects Under Network Interference,” *Review of Economics and Statistics*, 2020, *102*, 368–380.
- Manski, C.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *Review of Economic Studies*, 1993, *60* (3), 531–542.
- , “Identification of Treatment Response with Social Interactions,” *The Econometrics Journal*, 2013, *16* (1), S1–S23.
- Miguel, E. and M. Kremer**, “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities,” *Econometrica*, 2004, *72* (1), 159–217.
- Milgrom, P. and J. Roberts**, “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities,” *Econometrica*, 1990, pp. 1255–1277.
- Montanari, A. and A. Saberi**, “The Spread of Innovations in Social Networks,” *Proceedings of the National Academy of Sciences*, 2010, *107* (47), 20196–20201.
- Paluck, E., H. Shepherd, and P. Aronow**, “Changing Climates of Conflict: A Social Network Experiment in 56 Schools,” *Proceedings of the National Academy of Sciences*, 2016, *113* (3), 566–571.
- Penrose, M.**, *Random Geometric Graphs*, Oxford University Press, 2003.
- Sävje, F.**, “Causal Inference with Misspecified Exposure Mappings,” *arXiv preprint arXiv:2103.06471*, 2021.
- , **P. Aronow, and M. Hudgens**, “Average Treatment Effects in the Presence of Unknown Interference,” *arXiv preprint arXiv:1711.06399*, 2020.
- Shalizi, C. and A. Thomas**, “Homophily and Contagion Are Generically Confounded in Observational Social Network Studies,” *Sociological Methods and Research*, 2011, *40* (2), 211–239.

Toulis, P. and E. Kao, “Estimation of Causal Peer Influence Effects,” in “International Conference on Machine Learning” 2013, pp. 1489–1497.

van der Hofstad, R., *Random Graphs and Complex Networks*, Vol. 1, Cambridge University Press, 2016.

Vazquez-Bare, G., “Identification and Estimation of Spillover Effects in Randomized Experiments,” *arXiv preprint arXiv:1711.02745*, 2020.

Viviano, D., “Policy Targeting under Network Interference,” *arXiv preprint arXiv:1906.10258*, 2021.

Xu, X. and L. Lee, “Estimation of a Binary Choice Game Model with Network Links,” *Ohio State University working paper*, 2015.