

Dean Eckles¹ / Brian Karrer² / Johan Ugander³

Design and Analysis of Experiments in Networks: Reducing Bias from Interference

¹ Massachusetts Institute of Technology, Sloan School of Management, Cambridge, MA, USA, E-mail: eckles@mit.edu.
<http://orcid.org/0000-0001-8439-442X>.

² Facebook, Menlo Park, CA, USA, E-mail: briankarrer@fb.com

³ Department of Management Science & Engineering, Stanford University, Stanford, CA, USA, E-mail: jugander@stanford.edu

Abstract:

Estimating the effects of interventions in networks is complicated due to interference, such that the outcomes for one experimental unit may depend on the treatment assignments of other units. Familiar statistical formalism, experimental designs, and analysis methods assume the absence of this interference, and result in biased estimates of causal effects when it exists. While some assumptions can lead to unbiased estimates, these assumptions are generally unrealistic in the context of a network and often amount to assuming away the interference. In this work, we evaluate methods for designing and analyzing randomized experiments under minimal, realistic assumptions compatible with broad interference, where the aim is to reduce bias and possibly overall error in estimates of average effects of a global treatment. In *design*, we consider the ability to perform random assignment to treatments that is correlated in the network, such as through graph cluster randomization. In *analysis*, we consider incorporating information about the treatment assignment of network neighbors. We prove sufficient conditions for bias reduction through both design and analysis in the presence of potentially global interference; these conditions also give lower bounds on treatment effects. Through simulations of the entire process of experimentation in networks, we measure the performance of these methods under varied network structure and varied social behaviors, finding substantial bias reductions and, despite a bias–variance tradeoff, error reductions. These improvements are largest for networks with more clustering and data generating processes with both stronger direct effects of the treatment and stronger interactions between units.

Keywords: causal inference, field experiments, peer effects, spillovers, social contagion, social network analysis, graph partitioning

DOI: 10.1515/jci-2015-0021

1 Introduction

Many situations and processes of interest to scientists involve individuals interacting with each other, such that causes of the behavior of one individual are also indirect causes of the behaviors of other individuals; that is, there are *peer effects* or *social interactions* [1]. Likewise, in applied work, the policies considered by decision-makers often have many of their effects through the interactions of individuals [2]. Examples of such cases are abundant. In online social networks, the behavior of a single user explicitly and by design affects the experiences of other users in the network. If an experimental treatment changes a user’s behavior, then it is reasonable to expect that this will have some effect on their friends, a perhaps smaller effect on their friends of friends, and so on out through the network. In an extreme case, treating one individual could alter the behavior of everyone in the network.

To see the challenges this introduces, consider what is, in many cases, a primary quantity of interest for experiments in networks – the average treatment effect (ATE) of applying a treatment to all units compared with applying a different (control) treatment to all units.¹ Let Z be a vector of length N giving each unit’s treatment assignment, so that $Y_i(Z = z)$ is the potential outcome of interest for unit i when Z is set to z . Then the ATE is a contrast between two such treatment vectors,

$$\tau(z_1, z_0) = \frac{1}{N} \sum_i E[Y_i(Z = z_1) - Y_i(Z = z_0)], \quad (1)$$

where N is the number of units and z_1 and z_0 are two treatment assignments vectors; the prototypical case has $z_1 = 1$ and $z_0 = 0$, the N -vectors of all ones and of all zeros, such that we would call τ the *global* ATE.

Dean Eckles is the corresponding author.

Authors are listed in alphabetical order.

©2017 by De Gruyter.

Note that each unit's potential outcome is a function of the global treatment assignment vector Z , not just its own treatment Z_i .² Additional assumptions will thus be required for τ to be identifiable.³ The standard approach is to assume that each unit's response is not affected by the treatment of any other units. Versions of this assumption are sometimes called the *stable unit treatment value assumption* (SUTVA) [12], a *no interference* [13] assumption, or an *individualistic treatment response* (ITR) [14] assumption. Combined with random assignment to treatment, this suffices to identify τ . However, for many processes and situations of interest the units are interacting, and SUTVA becomes implausible [7, 15].

Rather than substituting other strong assumptions about interference that may result in point identifying τ , this paper considers how we can reduce bias in estimates of τ through both the choice of experimental design and analysis when interactions among units occur along an observed network.⁴ The design of the experiment dictates how each vertex in the network (i.e., unit) is assigned to a condition, and the analysis says how the observed responses are combined into estimates of causal quantities of interest. We study these methods by formalizing the process of experimentation in networks and proving sufficient conditions for bias reduction through design and analysis. These sufficient conditions for bias reduction are also sufficient for bounding the ATE. We augment these theoretical results with extensive simulations.

We do not consider all possible designs and analysis, but limit this work to some relatively general methods for each. We consider experimental designs that assign clusters of vertices to the same treatment; this is *graph cluster randomization* [16]. Since the counterfactual situations of interest involve all vertices being in the same condition, the intuition is that assigning a vertex and vertices near it in the network into the same condition, the vertex is "closer" to the counterfactual situation of interest.⁵ For analysis methods, we consider methods that define *effective treatments* such that only units that are effectively in global treatment or global control are used to estimate the ATE. For example, an estimator for the ATE might only compare units in treatment that are surrounded by units in treatment with units in control that are surrounded by units in control. The intuition is that a unit that meets one of these conditions is again "closer" to the counterfactual situation of interest.

The rest of the paper is structured as follows. We briefly review some related work on experiments in networks. Section 2 presents a model of the process of experimentation in networks, including initialization of the network, treatment assignment, outcome generation, and analysis. This formalization allows us to develop theorems giving sufficient conditions for bias reduction. To develop further understanding of the magnitude of the bias and error reduction in practice, Section 3 presents simulations using networks generated from small-world models and then degree-corrected blockmodels. While our theoretical sufficient conditions for bias reduction are somewhat restrictive, the simulations include data generating processes that do not meet these sufficient conditions yet still show substantial bias and error reductions, demonstrating that our alternative design and analysis approaches remain useful far outside the range of the theorems.

We find that graph cluster is capable of dramatically reducing bias compared to independent assignment without adding "too much" variance. The benefits of graph cluster randomization are larger when the network has more local clustering and when social interactions are strong. If social interactions are weak or the network has little local clustering, then the benefits of the more complex graph-clustered design are reduced. Finally, we found larger bias and error reductions through design than analysis: analysis strategies using neighborhood-based definitions of effective treatments do further reduce bias, but often at a substantial cost to precision such that the simple estimators were preferable in terms of error. No combination of design and analysis is expected to work well across very different situations, but these general insights from simulation can be a guide to practical real-world experimentation in the presence of peer effects.

1.1 Related work

Much of the literature on interference between units focuses on situations where there are multiple independent groups, such that there are interactions within, but not between, groups, (e.g., [6–8, 17]). Some more recent work has studied interference when the analyst may only observe a single, connected network [9, 14–16, 18–22], where this between-groups independence structure cannot be assumed. Walker and Muchnik [23] review some of this work, including a previous version of this paper.

Two features are common to much of this growing body of work on interference in networks. First, most work has focused on assuming restrictions on the extent of interference (e.g., vertices are only affected by the number of neighbors treated) and then deriving results for designs and estimators motivated by these same assumptions. Aronow and Samii [15] give unbiased estimators for ATEs under these assumptions and derive variance estimators. Also using these assumptions, Ugander et al. [16] show that graph cluster randomization puts more vertices in the conditions required for these estimators, such that the variance of these estimators is bounded for certain types of networks, such that the asymptotic variance is $O(1/N)$. But, as noted by Manski [14] and as we discuss in Section 2.3.2 below, the very processes expected to produce interference also make these assumptions implausible. There are some exceptions. Like the present work, Choi [21] uses monotonicity

assumptions but allows for global interference in the context of inference for attributable effects. Aronow [18] and Athey et al. [22] develop tests for spillovers that are valid under arbitrary global interference. Similar to the present work, the setup in Toulis and Kao [9] presents bias–variance tradeoffs, as they study how network structure constrains how “manipulable” the number of treated neighbors is.

Second, much of the related work has focused on detecting or estimating effects of peer assignments on ego outcomes; that is, estimating the magnitude of local interference (i.e., exogenous peer effects, indirect effects, spillovers) rather than estimating a global ATE. Again, one exception is Choi [21], which considers a contrast between global control (or treatment) and whatever treatment vector was actually observed.

The present work explicitly considers realistic data generating processes that violate these restrictive assumptions. That is, in contrast to prior work, we evaluate design and analysis strategies in the absence of assumptions that deliver particular desirable properties (e.g., unbiasedness, asymptotic consistency). Instead, we settle for reducing the bias and error of our inevitably biased estimators.

2 Model of experiments in networks

We consider experimentation in networks as consisting of four phases: (i) *initialization*, (ii) *treatment assignment*, (iii) *outcome generation*, and (iv) *estimation*. A single run through these phases corresponds to a single instance of the experimental process. Treatment assignment embodies the experimental design, and the estimation phase embodies the analysis of the network experiment. These same phases, shown in Figure 1, are implemented in our simulations in which we instantiate this process many times.

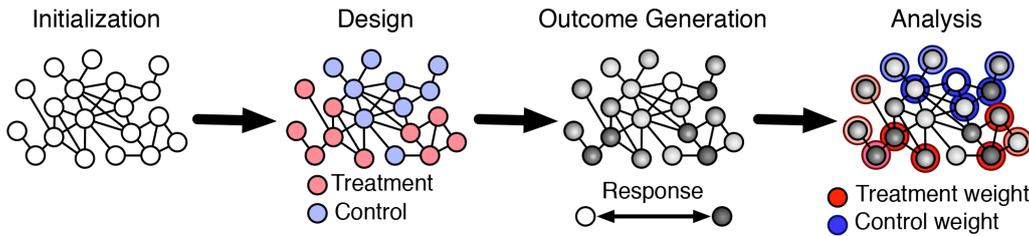


Figure 1 Model of the network experimentation process, consisting of (i) *initialization*, which generates the graph and vertex characteristics, (ii) *design*, which determines the randomization scheme, (iii) *outcome generation*, which observes or simulates behavior, and (iv) *analysis*, which constructs an estimator. We examine the bias and variance of treatment effect estimators under different design and analysis methods for varied initialization and outcome generation processes.

2.1 Initialization

Initialization is everything that occurs prior to the experiment. This includes network formation and the processes that produce vertex characteristics and prior behaviors. In some cases, we may regard this initialization process as random, and so wish to understand design and analysis decisions averaged over instances of this process; for example, we may wish to average over a distribution of networks that corresponds to a particular network formation model. In the simulations later in this paper, we generate networks from small-world models [24] and degree-corrected blockmodels [25]. In other cases, we may regard the outcome of this process as fixed; for example, we may be working with a particular network and vertices with particular characteristics, which we wish to condition on in planning our design and analysis.

When initialization is complete, we have a particular network $G = (V, E)$ with adjacency matrix A .⁶ In addition to producing a graph, the initialization process could also produce a collection of vertex characteristics X that may or may not relate to the structure of the graph, but may play a role in outcome generation.

2.2 Design: Treatment assignment

The treatment assignment phase creates a mapping from vertices to treatment conditions. We only consider a binary treatment here (i.e., an “A/B” test), so the mapping is from vertex to treatment or control. Treatment assignment normally involves independent assignment of units to treatments, such that one unit’s assignment is uncorrelated with other units’ assignments.⁷ In this case, each unit’s treatment is a Bernoulli random variable

$$Z_i \sim \text{Bernoulli}(q)$$

with probability of assignment to the treatment q .

The present work evaluates treatment assignment procedures that produce assignments with network autocorrelation. While many methods could produce such network autocorrelation, we work with graph cluster randomization, in which the network is partitioned into clusters and those clusters are used to assign treatments. Let the vertices be partitioned into N_C clusters C_1, C_2, \dots, C_{N_C} , and define $C(\cdot) : \{1, \dots, N\} \rightarrow \{1, \dots, N_C\}$ as mapping vertex indices to cluster indices. Thus C_j refers to cluster j by its index, while $C(i)$ refers to the cluster containing vertex i .

In standard graph cluster randomization, as presented by Ugander et al. [16], treatments are assigned at the cluster level, where each cluster C_j is assigned a treatment $W_j \sim \text{Bernoulli}(q)$. Thus the treatments assigned to vertices are simply those assigned to their clusters,

$$Z_i = W_{C(i)}.$$

For some estimands and analyses, assigning all vertices in a cluster to the same treatment can make it impossible for some vertices to be observed with, e.g., some particular number of treated peers. This can violate the standard requirement for causal inference that all units have positive probability of assignment to all conditions being compared *positivity*, cf. [15]. For this reason, it can be desirable to use an assignment method that allows for some vertices to be assigned to a different treatment than the rest of its cluster; we describe such a modification in Appendix A.1.

Graph cluster randomization could be applied to any mapping $C(\cdot)$ of vertices to clusters. This could include methods developed for community detection [26]. Many global community detection methods, such as modularity maximization [27], have a resolution limit such that they do not distinguish small clusters [28]; graph cluster randomization applied using these clusters could then introduce too large an increase in variance for the resulting bias reduction. Therefore, local clustering methods (e.g., ϵ -net clustering [16]) may be more appealing for graph cluster randomization. Observed community membership (e.g., current educational institution) or geography (e.g., village of residence) could also be used as this mapping.⁸ In using any graph partitioning method, the experimenter can choose which network to partition; for example, if multiple networks are observed, the experimenter might choose the network on which they believe the relevant social interactions are occurring. In practice, graph partitioning methods are often applied to sparsified networks where edges believed to be irrelevant are removed (i.e., some sparsification is performed on the network). Lastly, independent random assignment can be considered as clustered random assignment where each vertex is in its own cluster.

2.3 Outcome generation and observation

Given the network (along with vertex characteristics and prior behavior) and treatment assignments, some data generating process produces the observed outcomes of interest. In the context of social networks, typically this is the unknown process by which individuals make their decisions. In this work, we consider a variety of such processes. For our simulations, we use a known process meant to simulate decisions, in which units respond to others' prior behaviors. Doing so allows us to understand the performance of varied design and analysis methods, measured in terms of estimators' bias and error, under varied (although simple) decision mechanisms. Before considering these processes themselves, we consider outcomes as a function of treatment assignment.

2.3.1 Treatment response assumptions

In the following presentation, we use the language of "treatment response" assumptions developed by Manski [14] to organize our discussion of outcome generation. Consider vertices' outcomes as determined by a function from the global treatment assignment $Z \in \mathbb{Z}^N$ and an independent stochastic component $U \in \mathbb{U}^N$ to an outcome vector $Y \in \mathbb{Y}^N$:

$$f(\cdot) : \mathbb{Z}^N \times \mathbb{U}^N \rightarrow \mathbb{Y}^N.$$

We then observe $Y = f(Z, U)$. We can decompose this function into a function for each vertex

$$f_i(\cdot) : \mathbb{Z}^N \times \mathbb{U}^N \rightarrow \mathbb{Y}.$$

We can, as we have done above, continue to write $Y_i(Z = z)$ to refer to the outcome for vertex i that would be observed under assignment z ; by suppressing dependence on U , this treats $Y_i(\cdot)$ as a (possibly) stochastic function.

If vertices' outcomes are not affected by others' treatment assignment, then SUTVA is true. Perhaps more felicitously, Manski [14] calls this assumption *individualistic treatment response* (ITR). Under ITR we could then consider vertices as having a function from only their own assignment to their outcome:

$$f_i(\cdot) : \mathbb{Z} \times \mathbb{U}^N \rightarrow \mathbb{Y}.$$

One way for this assumption to hold is if the vertices do not interact.⁹ This specification of $f_i(\cdot)$ corresponds to the assumption that a vertex's outcome is invariant to changes in other vertices' assignments. That is, for any two global assignments $z_0, z_1 \in \mathbb{Z}^N$ and any stochastic component $U \in \mathbb{U}^N$,

$$z_{1,i} = z_{0,i} \Rightarrow f_i(z_1, U) = f_i(z_0, U).$$

ITR is a particular version of the more general notion of *constant treatment response* (CTR) assumptions [14]. More generally, a CTR assumption involves establishing equivalence classes of treatment vectors by defining a function $g_i(\cdot) : \mathbb{Z}^N \rightarrow \mathbb{G}_i$ that maps global treatment vectors to the space \mathbb{G}_i of *effective treatments* for vertex i [14] such that

$$g_i(z_1) = g_i(z_0) \Rightarrow f_i(z_1, U) = f_i(z_0, U)$$

for any two global assignments $z_0, z_1 \in \mathbb{Z}^N$ and any stochastic component $U \in \mathbb{U}^N$. Specifying the functions g_i is then a general way to specify a CTR assumption. Such assumptions, which specify levels sets of $f_i(\cdot)$, can be described as constituting an *exposure model* [15, 16].

Other CTR assumptions have been proposed that allow for some interference. Aronow and Samii [15] simply posit different restrictions on this function, such as that a vertex's outcome only depends on its assignment and its neighbors' assignments. This *neighborhood treatment response* (NTR) assumption has that, for any two global assignments $z_0, z_1 \in \mathbb{Z}^N$ and any stochastic component $U \in \mathbb{U}^N$,

$$z_{1,i} = z_{0,i} \text{ and } z_{1,\delta(i)} = z_{0,\delta(i)} \Rightarrow f_i(z_1, U) = f_i(z_0, U),$$

where $\delta(i)$ are the neighbors of vertex i . Aronow and Samii [15] and Ugander et al. [16] consider further restrictions, such as that a vertex's response only depends on the number of treated neighbors.

2.3.2 Implausibility of tractable treatment response assumptions

How should we select an exposure model? Aronow and Samii [15] suggest that we “must use substantive judgment to fix a model somewhere between the traditional randomized experiment and arbitrary exposure models”. However, it is unclear how substantive judgement can directly inform the selection of an exposure model for experiments in networks – at least when the vast majority of vertices are in a single connected component. Interference is often expected because of social interactions (i.e., peer effects) where vertices respond to their neighbors' behaviors: in discrete time, the behavior of a vertex at t is affected by the behavior of its neighbors at $t - 1$; if this is the case, then the behavior of a vertex at t would also be affected by the behavior of its neighbors' neighbors at $t - 2$, and so forth. Such a process will result in violations of the NTR assumption, and many other assumptions that would make analysis tractable. Manski [14] shows how some, quite specific, models of simultaneous endogenous choice can produce some restrictions on $f_i(\cdot)$.¹⁰ Since many appealing CTR assumptions are violated by the very theories that motivate the expectation of interference, it is useful to evaluate the performance of available design and analysis methods – including estimators that would be motivated by these assumptions – under outcome generating processes consistent with these theories. In particular, we now consider outcome generating processes in which vertices respond to their own treatment and the prior behavior of their neighbors. That is, peer behavior fully mediates the effects of the assignments of an ego's peers on the ego. This is notably different from Aronow and Samii [15] and Ugander et al. [16], where ego response is specified in terms of peer assignments without being mediated through peer behavior.¹¹

We consider a dynamical model with discrete time steps in which a vertex's behavior at time t , denoted by the vector $Y_{i,t}$, is a function h of ego treatment assignment and it and its neighbors' prior behaviors $Y_{\delta(i),t-1}$, such that

$$h_{i,t}(\cdot) : \mathbb{Z} \times \mathbb{Y}^{k_i+1} \times \mathbb{U}^N \rightarrow \mathbb{Y},$$

where k_i is the degree of vertex i and $Y_{\cdot,0}$ is initialized by some prior process. That is, $h_{i,t}(\cdot)$ is the nonparametric structural equation (NPSE) for $Y_{i,t}$.

Together with the graph G , the functions $h_{i,t}(\cdot)$ determine the treatment response function $f_i(\cdot)$. Thus, this outcome generating process implies some CTR assumptions. After the first time step (i.e., at time 1), the effective

treatment for a vertex, the function $g_i(\cdot)$ considered earlier, maps to the space of the vertex's treatment. After the second time step, it maps to the space of the vertex's treatment and its neighbors treatments. After the third time step (i.e., at time 3), the effective treatment is no finer than the treatments of all vertices within distance 2. At time step t , the effective treatment is no finer than the treatments of all vertices within distance $t - 1$. We see here that under such a dynamic outcome generating process, Manski's notion of effective treatment, conceived of to limit the scope of dependence, quickly expands to encompass entire connected components of the network.¹² See Figure 2 for a graphical illustration.

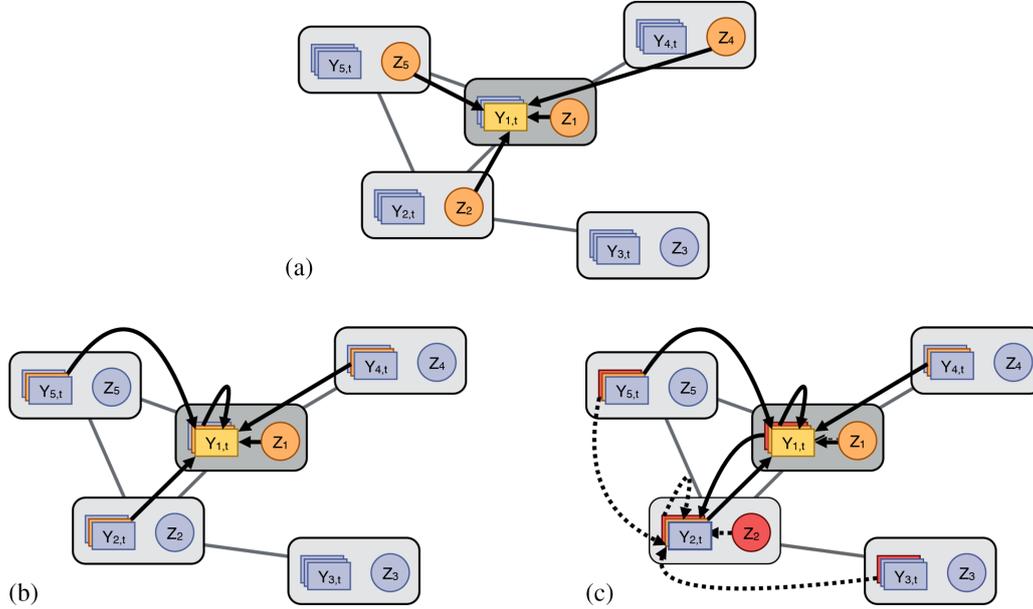


Figure 2 Varieties of interference illustrated in a small social network. (a) Interference under the neighborhood treatment response (NTR) assumption, where the response of vertex 1 at time t (light orange) depends on the treatments of its neighbors and itself (dark orange). (b) Interference due to social interactions (i.e., peer effects, social contagion), where the response of vertex 1 at time t depends on its own treatment Z_1 and the responses of its neighbors at time $t - 1$. (c) Interference caused by social interactions induces long range dependence; for example, the response of vertex 3 at time $t - 1$ in turn depends on vertex 3's treatment and the responses of their neighbors at time $t - 2$ (dashed lines and red objects). Interference under the NTR assumption has no long range dependence, but is much less realistic.

2.3.3 Utility linear-in-means

Many familiar models are included in the above outcome generating process. To make this more concrete, clarify the relationship to work on graphical games, and for our subsequent simulations, we consider a model in which a vertex's behavior is a stochastic function of the mean of neighbors' prior behaviors, so that behavior at some new time step t is generated as:

$$Y_{i,t}^* = \alpha + \beta Z_i + \gamma \frac{A_i' Y_{t-1}}{k_i} + U_{i,t} \quad (2)$$

$$Y_{i,t} = a(Y_{i,t}^*) \quad (3)$$

where A_i is a row of the adjacency matrix and k_i is the degree of vertex i . One interesting case, when the behavior is binary, has $a(x) = 1[x > 0]$ and $U_{i,t} \sim \mathcal{N}(0, 1)$, which is a probit model. Here α is the baseline, where a negative α determines the threshold that must be crossed for $Y_{i,t}^*$ to be positive. Setting β determines the strength of the direct effect of the treatment, while γ is the slope for peer behavior, and therefore determines the strength of the peer effects. For some small value of t , this implies CTR assumptions, as described above.

This can be interpreted as a simple *noisy best response* or *best reply* model [31], when vertices anticipate neighbors taking the same action in the present round as they did in the previous round. In particular, we can interpret $Y_{i,t}^*$ as the payoff for vertex i to adopt behavior 1 at time t . When $\gamma > 0$, then this is a semi-anonymous graphical game with strategic complements ([32], Ch. 9). When $a(x) = x$, this is a linear-in-means model, which is widely studied and used in econometrics, (e.g., [33–36]). This suggests the relevance of theoretical or simulation results with this model.

2.4 Analysis and estimation

We focus on the ATE (the average treatment effect; τ in eq. (1)) of global treatment. Compared with direct and indirect effects separately, this quantity is naturally of interest when considering whether a new treatment would be beneficial if applied to all units, (cf. [10]).

There are many options available for estimating the ATE. For example, if the relevant network is completely unknown or if peer effects are not expected, then one might use estimators for experiments without interference, such as a simple difference-in-means between the outcomes of vertices assigned to treatment and control. To clarify the sources of error in estimation, we begin with the population analogs of these quantities – i.e., the associated estimands defined with respect to the observed and unobserved potential outcomes – and return to the estimators themselves in Section 2.4.3. Consider the simple difference-in-means estimand

$$\tau_{\text{ITR}}^d(1, 0) = \mu_{\text{ITR}}^d(1) - \mu_{\text{ITR}}^d(0) \quad (4)$$

where the μ_{ITR}^d are mean outcomes when a vertex is in treatment and control, i.e.,

$$\mu_{\text{ITR}}^d(z) = \frac{1}{N} \sum_{i=1}^N E^d[Y_i | Z_i = z_i].$$

We index these quantities by both the definition of effective treatments (ITR for “individualistic treatment response”, as in Section 2.3.1) and the experimental design d , since the former determines the conditioning involved and the latter determines the distribution of Z over which we take expectations. Note that the effective treatment definition determines the conditioning, but need not match the true effective treatment definition.

When a vertex’s outcome depends on the treatment assignments of others, these estimands need not equal the quantities of interest. That is, they can suffer from some *estimand bias* (or model bias), such that $\tau_{\text{ITR}}^d(1, 0) - \tau(1, 0)$ is non-zero. Each vertex assigned to treatment contributes to this bias through the difference between its expected outcome when assigned to treatment (given the experimental design) and what would be observed under global treatment. More generally, for some global treatment vector z , vertex i contributes to the bias of $\mu_{\text{ITR}}^d(z)$ through $E^d[Y_i - Y_i(Z = z) | Z_i = z_i]$. If the treatment assignment of other vertices do not affect vertex i ’s behavior much, then this contribution might be quite small. Or this contribution could be more substantial.

2.4.1 Bias reduction through design

We are now equipped to elaborate on the intuition that graph cluster randomization puts vertices in conditions “closer” to the global treatments of interest and thereby reduces bias in estimates of average treatment effects, even if a vertex’s outcome depends on the global treatment vector.

Theorem 2.1 Assume we have a linear outcome model for all vertices $i \in V$ that is monotonically increasing in z ; that is, there exists an N -vector a and $N \times N$ matrix B with non-negative entries $B_{ij} \geq 0$ such that

$$E_U[Y_i(z, U)] = a_i + \sum_{j \in V} B_{ij} z_j. \quad (5)$$

Then for any mapping of vertices to clusters $C(\cdot)$, the absolute bias of $\tau_{\text{ITR}}^d(1, 0)$ when the design d is graph cluster randomization is less than or equal to the absolute bias when d is independent assignment, with a fixed treatment probability p .

Proof. Using the linear model for Y_i and the definition of τ , we have that the true ATE τ is given by

$$\tau(1, 0) = \mu(1) - \mu(0) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N B_{ij} \quad (6)$$

for this outcome model. Under graph cluster randomization,

$$\tau_{\text{ITR}}^{\text{gcr}}(1, 0) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^N B_{ij} \mathbf{1}[C(i) = C(j)]. \quad (7)$$

Then under independent assignment,

$$\tau_{\text{ITR}}^{\text{ind}}(1, 0) = \frac{1}{N} \sum_{i=1}^N B_{ii}. \quad (8)$$

Because $B_{ij} \geq 0$, together this implies that $\tau(1, 0) - \tau_{\text{ITR}}^{\text{gr}}(1, 0) \leq \tau(1, 0) - \tau_{\text{ITR}}^{\text{ind}}(1, 0)$, where monotonicity dictates that each side of this inequality is positive.

The bias of graph cluster randomization in Theorem 2.1 is small when $C(i) = C(j)$ is true for the pairs (i, j) with large coefficients B_{ij} and B_{ji} . This comparison allows seeing how, at least in this linear model, the magnitude of bias reduction from graph cluster randomization depends on the “strength” of the interactions within clusters. That is, this clarifies the intuition that using clusters formed from more distant vertices will not generally reduce bias as much as clusters formed from closer vertices, as is the aim of using graph partitioning methods such as ϵ -net partitioning or community detection methods.¹³ It also highlights that when there are mainly non-zero B_{ij} 's, *ceteris paribus* large clusters result in more bias reduction; of course, there are corresponding costs to precision.

To clarify this further, let's consider the relative bias defined by

$$\tau_{\text{ITR}}^{\text{gr}}(1, 0)/\tau(1, 0) - 1 = \frac{\sum_{i=1}^N \sum_{j=1}^N B_{ij} \mathbf{1}[C(i) = C(j)]}{\sum_{i=1}^N \sum_{j=1}^N B_{ij}} - 1. \quad (9)$$

Assume that there are $O(N)$ clusters of size $O(1)$ used for the graph cluster randomization.¹⁴ Under this condition, the numerator has $O(N)$ terms and the denominator has $O(N^2)$ terms. So unless there is a judicious choice of clustering, the numerator will be overwhelmed by the denominator and the estimator $\tau_{\text{ITR}}^{\text{gr}}(1, 0)$ will be a dramatic underestimate of the true ATE. Meanwhile it is clear that $\tau_{\text{ITR}}^{\text{ind}}(1, 0)$ would be even worse. In order for meaningful relative bias reduction to occur, the clustering must capture the structure of the dependence between units specified by the matrix of coefficients B . Note also that a proof parallel to that of Theorem 2.1, under the same assumptions, shows that $\tau_{\text{ITR}}^{\text{gr}}(1, 0)$ is a lower bound on the ATE.

In Appendix A.2, we derive similar intuitions from an alternative graph cluster randomization that preserves balance between the sizes of the treatment and control group. There graph cluster randomization no longer always achieves bias reduction for every clustering over independent assignment, but meaningful bias reduction is again possible and again depends on how the clustering captures B in an identical way.

This linear outcome model, which has N^2 parameters, has as special cases some simpler models of interest. Let $a(x) = x$ in eq. (3). Then for $t \geq 1$ the quantity $E^U[Y_{i,t}(z)]$ is

$$E^U[Y_{i,t}(z)] = \alpha + \beta z_i + \gamma \frac{A_i' E^U[Y_{t-1}(z)]}{k_i}. \quad (10)$$

The closed form solution for $E^U[Y_t(z)]$ for any $t \geq 0$ is then given by

$$E^U[Y_t(z)] = (\gamma \mathbf{D}^{-1} \mathbf{A})^t E^U[Y_0] + \sum_{q=0}^{t-1} (\gamma \mathbf{D}^{-1} \mathbf{A})^q (\alpha + \beta z) \quad (11)$$

where D^{-1} is the diagonal matrix of inverse degrees, A is the adjacency matrix, and Y_0 is the vector of initial states. This is a linear outcome model with $a_i = \alpha(1 - \gamma^t)/(1 - \gamma) + ((\gamma \mathbf{D}^{-1} \mathbf{A})^t E^U[Y_0])_i$ and $B_{ij} = \beta \sum_{q=0}^{t-1} (\gamma \mathbf{D}^{-1} \mathbf{A})^q_{ij}$.

To be clear, the linearity and monotonicity assumption made in Theorem 2.1 are restrictive assumptions, but it is important to note that they are sufficient, rather than necessary, conditions for graph cluster randomization to reduce bias. The probit utility linear-in-means model introduced in Section 2.3.3 and used for the later simulations will not meet the conditions of the theorem, yet still shows bias reduction.

2.4.2 Bias reduction through analysis

Definitions of effective treatments other than ITR correspond to different estimands. In particular, we can incorporate assumptions about effective treatments into eq. (1). Let

$$\mu_g^d(z) = \frac{1}{N} \sum_{i=1}^N E^d[Y_i | g_i(Z) = g_i(z)] \quad (12)$$

be the mean outcome for the global treatment z when g specifies the effective treatments and d is the experimental design. Then we have

$$\tau_g^d(z_1, z_0) = \mu_g^d(z_1) - \mu_g^d(z_0) \quad (13)$$

as our revised estimand for the ATE.¹⁵ If the effective treatment definition is correct, then this estimand will also be the global ATE. As with the ITR assumption, we can again describe the estimand bias that occurs when effective treatments are incorrectly specified. For some global treatment vector z , vertex i contributes to the bias of $\mu_g^d(z)$ through

$$E^d[Y_i - Y_i(Z = z) | g_i(Z) = g_i(z)], \quad (14)$$

where $g_i(\cdot)$ is the potentially incorrect (i.e., too coarse) specification of effective treatments for vertex i .

Considering two specifications of effective treatments allows us to elaborate on the intuition that using a finer specification of effective treatments will reduce bias by comparing only vertices that are in conditions “closer” to the global treatments of interest. For example, the NTR assumption corresponds to finer effective treatments than the ITR assumption. We can also relax the NTR assumption to a fractional λ -neighborhood treatment in which a vertex is considered effectively in global treatment if a fraction λ of its neighbors are treated (and the same for control) [16].

To show sufficient conditions for bias reduction, we define the following generalization of this relationship among definitions of effective treatments. Consider functions $g_i(\cdot)$ such that $g_i(Z) = g_i(z)$ just implies that for some set of vertices $J_i \in V$ we have that $\sum_{j \in J_i} \mathbf{1}[Z_j = z_j] \geq l_i$ and that $Z_i = z_i$. These are conditions such that any subset of size l_i of a set of vertices J_i has treatment assignment matching that in z , the global treatment vector of interest. The fractional neighborhood treatment response (FNTR) assumption corresponds to such a function with $J_i = \delta(i)$ and $l_i = \lceil \lambda k_i \rceil$, where k_i is vertex i 's degree. This has both ITR and NTR as special cases with $\lambda = 0$ and $\lambda = 1$ respectively.¹⁶

Definition 2.2 If we have two such functions $g_i^A(\cdot)$ and $g_i^B(\cdot)$ with the same J_i , and $g_i^A(z) = g_i^A(z')$ implies $g_i^B(z) = g_i^B(z')$, then we say that $g_i^A(\cdot)$ is *more restrictive* than $g_i^B(\cdot)$.

Theorem 2.3 Let $g_i^A(\cdot)$ and $g_i^B(\cdot)$ be functions such that $g_i^A(\cdot)$ is more restrictive than $g_i^B(\cdot)$ for every vertex i , and let independent random assignment be the experimental design. A sufficient condition for estimand $\tau_{g^A}^{\text{ind}}(1, 0)$ to have less than or equal absolute bias than $\tau_{g^B}^{\text{ind}}(1, 0)$, where these estimands are defined by eq. (13), is that we have monotonically increasing responses or monotonically decreasing responses for every vertex with respect to z .

Proof. Given in Appendix A.3.

As with Theorem 2.1 above, a proof parallel to that of Theorem 2.3 shows that more restrictive functions $g_i^A(\cdot)$ yield estimands that are also increasingly tight (or at least as tight) lower bounds on the ATE. Note that the utility linear-in-means model in eq. (2) satisfies this monotonicity condition if the direct effect β and peer effect γ are both non-negative.

What about the combination of a design using graph cluster randomization with an analysis using these neighborhood-based estimands? As we show in Appendix A.3, similar arguments apply if $g_i^A(\cdot)$ and $g_i^B(\cdot)$ count matching *clusters* instead of vertices, but use of the FTNR estimand with graph cluster randomization is not necessarily bias reducing under monotonic responses without this modification.

2.4.3 Estimators

We now briefly discuss estimators for the estimands considered above. First, we can estimate $\tau_{\text{ITR}}^d(1, 0)$ with the difference in sample means $\hat{\tau}_{\text{ITR},S}(1, 0) = \hat{\mu}_{\text{ITR},S}(1) - \hat{\mu}_{\text{ITR},S}(0)$ where the $\hat{\mu}_{\text{ITR},S}$ are simple sample means, i.e.,

$$\hat{\mu}_{\text{ITR},S}(z) = \frac{1}{\sum_{i=1}^N \mathbf{1}[Z_i = z_i]} \sum_{i=1}^N Y_i \mathbf{1}[Z_i = z_i].$$

Note that these estimators are again indexed by the effective treatment used (i.e., ITR), but, unlike the estimands, they are not indexed by the design, though the design determines their distribution. We additionally distinguish these estimators by the weighting used (discussed below), identifying the simple (i.e., unweighted) means with S . If treatment is randomized, then this estimator will be unbiased for $\tau_{\text{ITR}}^d(1, 0)$ under independent randomization of clusters.¹⁷ More generally, there is a natural correspondence between the conditioning on $g_i(Z) = g_i(z)$ in the estimands and the vertices whose outcomes are used in an estimator. Given some specification of effective treatments g , one could construct an estimator of the ATE as a simple difference in the sample means for vertices in effective treatment and in effective control

$$\hat{\tau}_{g,S}(1, 0) = \hat{\mu}_{g,S}(1) - \hat{\mu}_{g,S}(0)$$

where we have

$$\hat{\mu}_{g,S}(z) = \frac{\sum_{i=1}^N Y_i \mathbf{1}[g_i(Z) = g_i(z)]}{\sum_{i=1}^N \mathbf{1}[g_i(Z) = g_i(z)]}.$$

This estimator will only be unbiased for the corresponding estimand $\mu_g^d(z)$ under certain conditions such that the effective treatments are unconfounded. One way for the effective treatments to be unconfounded is if either $E^d[Y_i | g_i(Z) = g_i(z)]$ or $\Pr^d[g_i(Z) = g_i(z)]$ is the same for all vertices. Usually we would not want to assume that $E^d[Y_i | g_i(Z) = g_i(z)]$ is homogeneous, and $\Pr[g_i(Z) = g_i(z)]$ will not be homogeneous under many relevant effective treatments, such as neighborhood treatment response (NTR), since the distribution of effective treatments for a vertex depends on network structure. For example, as Ugander et al. [16] observe, high degree vertices will generally have low probability of being assigned to some kinds of “extreme” effective treatments, such as having all neighbors treated, while low degree vertices have a much higher probability of being in such an effective treatment.

Observed effective treatments can be made unconfounded by conditioning on the design [15] or sufficient information about the vertices. The experimental design determines the probability of assignment to an effective treatment $\pi_i(z) = \Pr(g_i(Z) = g_i(z))$. In the case of graph cluster randomization and effective treatments determined by thresholds, these probabilities can be computed exactly using a dynamic program [16]. These are generalized propensity scores that can then be used in Horvitz–Thompson estimators or other inverse-probability weighted estimators, such as the Hajek estimator [15] of the ATE. The Horvitz–Thompson estimator will often suffer from excessive variance, so we focus on the Hajek estimator:

$$\hat{\tau}_{g,H}(z_1, z_0) = \left(\sum_{i=1}^N \frac{\mathbf{1}[g_i(Z)=g_i(z_1)]}{\pi_i(z_1)} \right)^{-1} \sum_{i=1}^N \frac{Y_i \mathbf{1}[g_i(Z)=g_i(z_1)]}{\pi_i(z_1)} - \left(\sum_{i=1}^N \frac{\mathbf{1}[g_i(Z)=g_i(z_0)]}{\pi_i(z_0)} \right)^{-1} \sum_{i=1}^N \frac{Y_i \mathbf{1}[g_i(Z)=g_i(z_0)]}{\pi_i(z_0)} \quad (15)$$

The bias of this Hajek estimator for eq. (13) is not zero, but it is typically small and worth the variance reduction, cf. [15].

Beyond bias, we also care about the variance of the estimator as well. Estimators making use only of vertices with all neighbors in the same condition will suffer from substantially increased variance, both because few vertices will be assigned to this effective treatment and because the weights in the Hajek estimator will be highly imbalanced. This could motivate borrowing information from other vertices, such as by using additional modeling or, more simply, through relaxing the definition of effective treatment, such as by using the fractional relaxation of the NTR assumption (FNTR).

The most appropriate effective treatment assumption to use for the analysis of a given experiment is not clear *a priori*. We will consider estimators motivated by two different effective treatments in our simulations.

3 Simulations

In order to evaluate both design and analysis choices, we conduct simulations that instantiate the model of network experiments presented above. First, graph cluster randomization puts more vertices into positions where their neighbors (and neighbors’ neighbors) have the same treatment; this is expected to produce observed outcomes “closer” to those that would be observed under global treatment. Second, estimators using fractional neighborhood treatment restrict attention to vertices that are “closer” to being in a situation of global treatment. Third, weighting using design-based propensity scores adjusts for bias resulting from associations between propensity of being in an effective treatment of interest and potential outcomes. Each of these three changes to design and analysis is expected to reduce bias, potentially at a cost to precision. Under some conditions, we have shown above that these design and analysis methods reduce (or at least do not increase) bias for the ATE. The goal of these simulations then is to characterize the magnitude of this bias reduction, weigh it against increases in variance, and do so specifically under circumstances that do not meet the given sufficient conditions for the theoretical results.

For each run of the simulation, we do the following. First, we construct a small world network with $N = 1,000$ vertices and initial degree parameter $k = 10$. We vary the rewiring probability $p_{rw} \in \{0.00, 0.01, 0.10, 0.50, 1.00\}$, thereby producing both regular powers of the cycle ($p_{rw} = 0$), graphs with “small world” characteristics ($p_{rw} \in \{0.01, 0.10\}$), graphs with many random edges and less clustering ($p_{rw} = 0.50$), and graphs with all random edges ($p_{rw} = 1.00$). The small world model of networks [24] is notable for being able to succinctly introduce clustering into an otherwise complex distribution over random graphs, all featuring a small diameter. The clustering of the graph, typically measured by the clustering coefficient, is a measure of the extent to which adjacent vertices

share many common neighbors in the graph, and many social networks, including online social networks, e.g. [39], have been found to exhibit a high degree of clustering as well as a small diameter.

For graph cluster randomization, we use ϵ -net clustering, as previously considered by Ugander et al. [16]. An ϵ -net in the graph distance metric is a set of vertices such that no two vertices in the set are less than ϵ hops of each other, and every vertex outside the set is within ϵ hops (in fact, $\epsilon - 1$ hops) of a vertex in the set. An ϵ -net can be formed by repeatedly selecting a vertex and removing it and every vertex within distance $\epsilon - 1$ from the network, until all vertices have been removed. Having completed this step, the population of selected vertices forms an ϵ -net. An ϵ -net clustering can be formed by assigning each vertex to the closest vertex in the ϵ -net, and breaking the possible ties through some arbitrary rule. We compare clustered random assignment using ϵ -nets (with $\epsilon = 3$) to independent random assignment, where vertices are independently assigned to treatment and control.¹⁸ We generate the observed outcomes using the probit model in eqs (2) and (3), and set the baseline as $\alpha = -1.5$, making the behavior somewhat rare:

$$Y_{i,t}^* = -1.5 + \beta Z_i + \gamma \frac{A_i' Y_{i,t-1}}{k_i} + U_{i,t}, \quad Y_{i,t} = 1\{Y_{i,t}^* > 0\}. \quad (16)$$

We initialize $Y_{i,0} = 0$ for all vertices, and then run the process for all combinations of $\beta \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$ and $\gamma \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$, up to a maximum time $T = 3$.¹⁹ Note that this data generating process *does not* satisfy the conditions for graph cluster randomization to be bias reducing given by Theorem [2.1], since the outcome model is not linear in Z . We emphasize that this is a specific choice of model, which we selected because it is a widely-used model that does not satisfy those conditions, while being low-dimensional to easily explore via simulation.

Finally, for each simulation, we compute three estimates of the ATE. The individual unweighted estimator (or difference-in-means estimator) $\hat{\tau}_{\text{ITR},S}$ makes no use of neighborhood information. This is the baseline to which we compare the neighborhood unweighted estimator $\hat{\tau}_{\text{FNTR},S}$ and the neighborhood Hajek estimator $\hat{\tau}_{\text{FNTR},H}$, both using a fractional neighborhood treatment response (FTNR) specification of effective treatments with $\lambda = 0.75$. That is, these estimators count a vertex as being in effective treatment or effective control if at least three-fourths of its neighbors have the same assignment. With independent assignment, the conditions for bias reduction given in Theorem 2.3 from using FNTR instead of ITR are satisfied.

We run each of these configurations 5,000 times. We estimate the true ATE with simulations in which all vertices are put in treatment or control. Each configuration is run 5,000 times for the global treatment case and 5,000 times for the global control case.²⁰ Our evaluation metrics are bias and root mean squared error (RMSE) of the estimated ATE.

3.1 Design

First we examine the bias and mean squared error of the estimated ATE for designs using graph cluster randomization compared with independent randomization. In both cases we use the difference-in-means estimator $\hat{\tau}_{\text{ITR},S}$. As expected, using graph cluster randomization reduces bias (Figure 3), especially when the peer effects and direct effects are large relative to the baseline ($\alpha = -1.5$), and when the network exhibits substantial clustering (i.e., the rewiring probability p_{rw} is small). While these bias reductions are small on an absolute scale, note that the true ATEs range from 0 to 0.41. When there are direct effects and peer effects, graph cluster randomization reduces the bias of the estimated ATE from independent randomization by 58% to 99% for $p_{\text{rw}} \in \{0.00, 0.01, 0.10\}$.

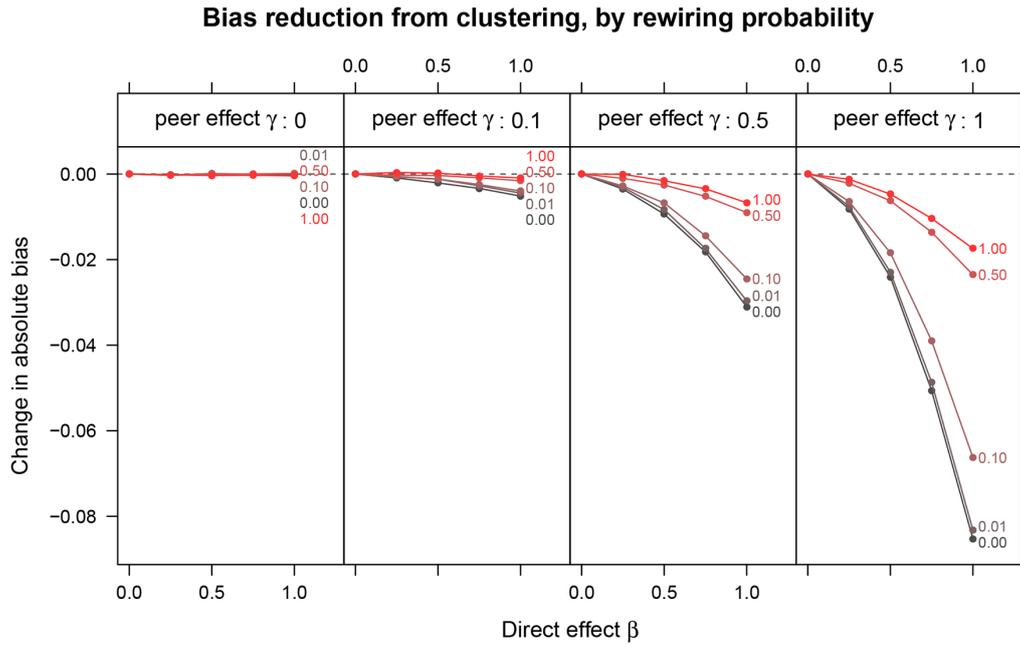


Figure 3 Change in bias due to clustered random assignment as a function of the direct effect of the treatment β , the rewiring p_{rw} probability (different colors), and the strength of the peer effect γ (different panels). Random assignment clustered in the network reduces bias, especially when peer effects are large relative to the baseline ($\alpha = -1.5$) and when the network is more clustered.

Reduction in bias can come with increases in variance, so it is worth evaluating methods that reduce bias also by the effect they have on the error of the estimates. We compare RMSE, which is increased by both bias and variance, between graph cluster randomization and independent assignment in Figure 4. In some cases, the reduction in bias comes with a significant increase in variance, leading to an RMSE that is either left unchanged or even increased. However, in cases where the bias reduction is large, this overwhelms the increase in variance, such that graph cluster randomization reduces not only bias but also RMSE substantially. For example, with substantial clustering ($p_{rw} = 0.01$) and peer effects ($\gamma = 0.5$), we observe approximately 40% RMSE reduction from graph cluster randomization. While the RMSE reduction is strongest under substantial clustering, if both the direct effect strength and peer effect strength are strong, we observe significant universal reductions in RMSE from clustered randomization (though to varied extents), regardless of the clustering structure given by p_{rw} . It is notable that even with small networks (recall that $N = 1,000$), the bias reduction from graph cluster randomization is large enough to reduce RMSE. To further examine the robustness of bias reduction through graph cluster randomization, Appendix A.4 reports on simulations with non-monotonic responses.

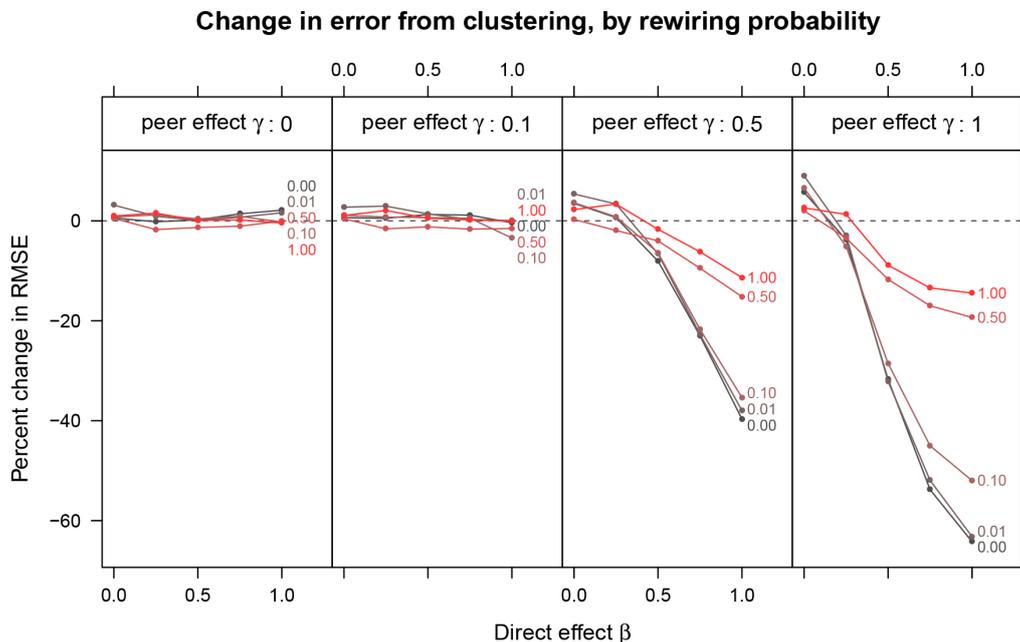


Figure 4 Percent change in root-mean-squared-error (RMSE) from clustered assignment for small world networks. While in some cases graph cluster randomization increases RMSE, in other cases (when bias reduction is large), it quite substantially reduces RMSE.

3.2 Design and analysis

In addition to changes in design, we can also use analysis methods intended to account for interference. We utilize the fractional neighborhood exposure model, which means we only include vertices in the analysis if at least three-quarters of their friends were given the same treatment assignment.²¹ With this neighborhood exposure model, we consider using propensity score weighting, which corresponds to the Hajek estimator, or ignoring the propensities and using unweighted difference-in-means. The second estimator has additional bias due to neglecting the propensity-score weights.

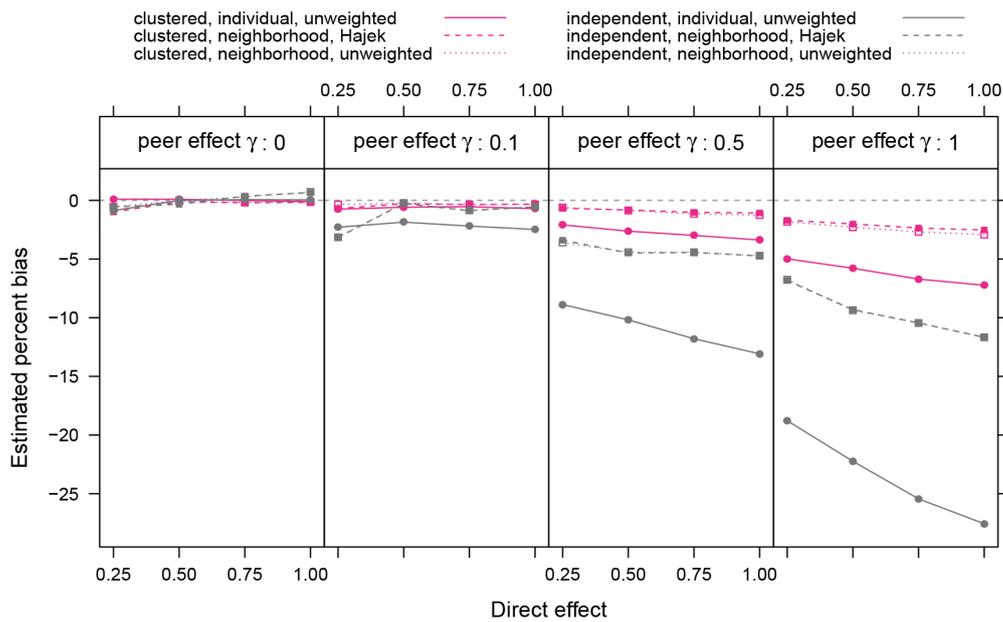


Figure 5 Relative bias in ATE estimates for different assignment procedures, exposure models, and estimation methods. The most striking differences are between the assignment procedures, though the neighborhood exposure model also reduces bias (at the cost of increased variance – see Figure 6). Relative bias is not defined when the true value is zero, so we exclude simulations with the direct effect $\beta = 0$. For all networks, the rewiring probability was $p_{rw} = 0.01$.

Figure 5 shows several combinations of design randomization procedure, exposure model, and estimator. We see that using a neighborhood-based definition of effective treatments further reduces bias, while the impact of using the Hajek estimator is minimal.

The low impact of the Hajek estimator follows understandably from the fact that small-world graphs do not exhibit any notable variation in vertex degree, which is the principle determinant of the propensities used by the Hajek estimator. Thus, for small-world graphs the weights used by the Hajek estimator are very close to uniform. With more degree heterogeneity expected in real networks, the weighting of the Hajek estimator will be more important, especially when these heterogeneous propensities are highly correlated with behaviors. In general, however, the change in bias from adjusting the analysis are not as striking as those from changes due to the experimental design.

Using the neighborhood exposure model means that the estimated average treatment effect is based on data from fewer vertices, since many vertices may not pass the *a priori* condition. So the observed modest changes in bias come with increased variance, as reflected in the change in RMSE compared with independent assignment without using the exposure condition (Figure 6).

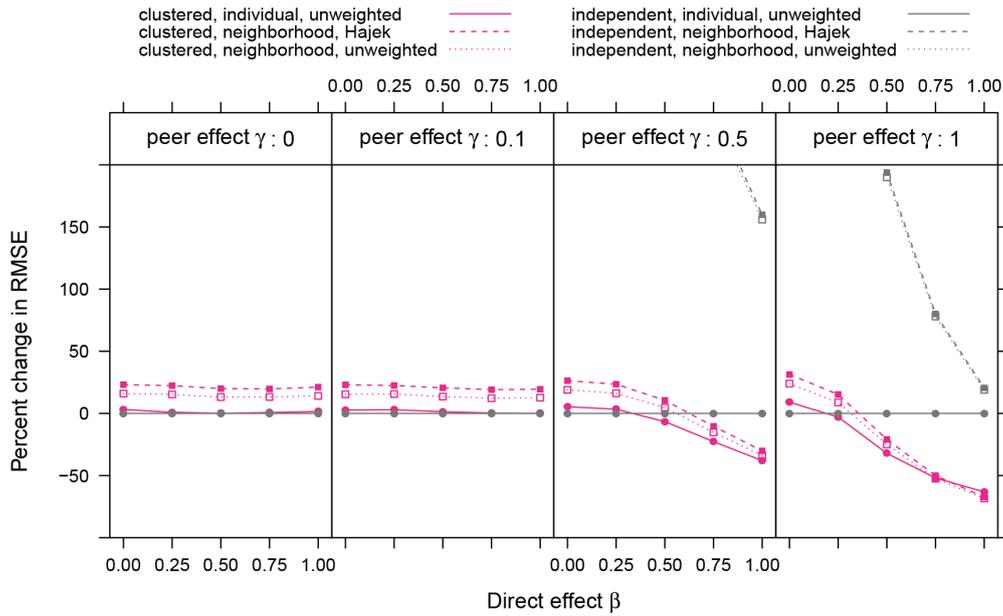


Figure 6 Percent change in root-mean-squared-error (RMSE) compared with independent assignment with the simple difference-in-means estimator. Using the neighborhood condition with independent assignment results in large increases in variance: for the two smaller values of γ , this produces an almost 400% increase in RMSE. For this reason, the y -axis is limited to not show these cases. Rewiring probability $p_{rw} = 0.01$.

3.3 Results with stochastic blockmodels

As a check on the robustness of these results to the specific choice of network model, we also conducted simulations with a degree-corrected block model (DCBM) [25], which provides another way to control the amount of local clustering in a graph and to produce more variation in vertex degree than is possible with small world networks.

In each simulation, the network is generated according to a DCBM with 1,000 vertices and 10 communities. We present results for a subset of the parameter values used with the small-world networks. Instead of varying the rewiring probability p_{rw} to control local clustering, we vary the expected proportion of edges that are within a community $p_{\text{comm}} \in \{0.2, 0.5, 0.8\}$ where vertices are assigned to one of the 10 communities uniformly at random. The distribution of expected degrees is a discretized log-normal distribution with mean 10 (as with the small-world networks) and variance 40. This produces substantially more variation in degree than the small-world networks. Each configuration is repeated 5,000 times.

Figure 7 displays the change in bias and error that results from graph cluster randomization in these simulations. Again, while small on the absolute scale, the observed bias reductions are, with the exception of cases where there are no direct effects, reducing bias by at least 20%. This is reflected in the RMSE reduction results, which include the added error due to variance from graph cluster randomization. The bias and error reduction with the DCBM networks is not as large, for the same values of other parameters, as with the small world networks. We interpret this as a consequence of the presence of higher-degree vertices and of less local clustering, even in the simulations with high community proportion (i.e., $p_{\text{comm}} = 0.8$).²² Qualitative features of these results (e.g., bias and error reduction increase with increases in peer effects and increases in clustering) match those from the small-world networks.

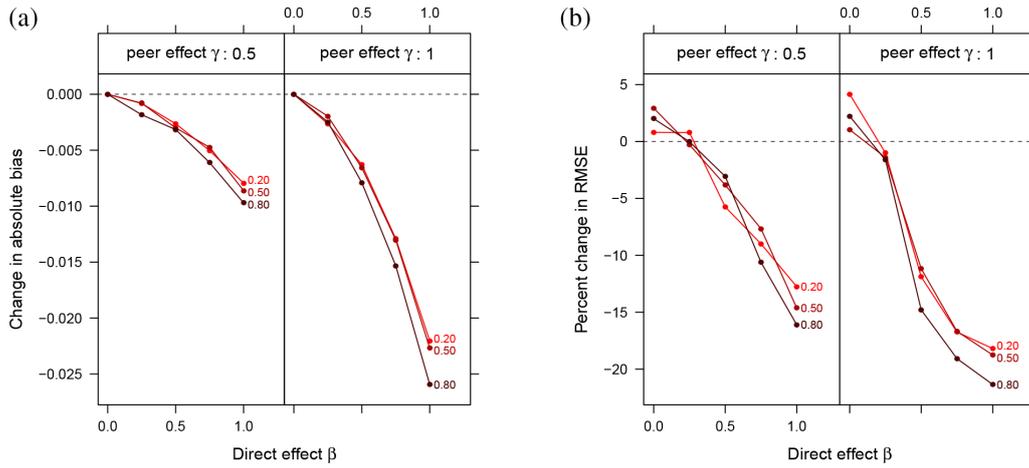


Figure 7 Change in (a) bias and (b) RMSE due to clustered random assignment. Lines are labeled with the expected proportion of edges that are within a community p_{comm} . As before, results vary with the strength of the peer effect γ , and the direct effect of the treatment β . The largest bias and error reductions here are not as substantial as the largest bias reductions with small-world networks.

Figure 8(a) displays bias as a function of both design and analysis decisions. As with the small-world networks, estimators making use of the λ -fractional neighborhood exposure condition reduce bias, whether used with independent or clustered random assignment. This additional bias reduction comes at the cost of additional variance, such that, in terms of MSE, estimators using the exposure condition are worse for many of the parameter values included in these simulations (Figure 8(b)).

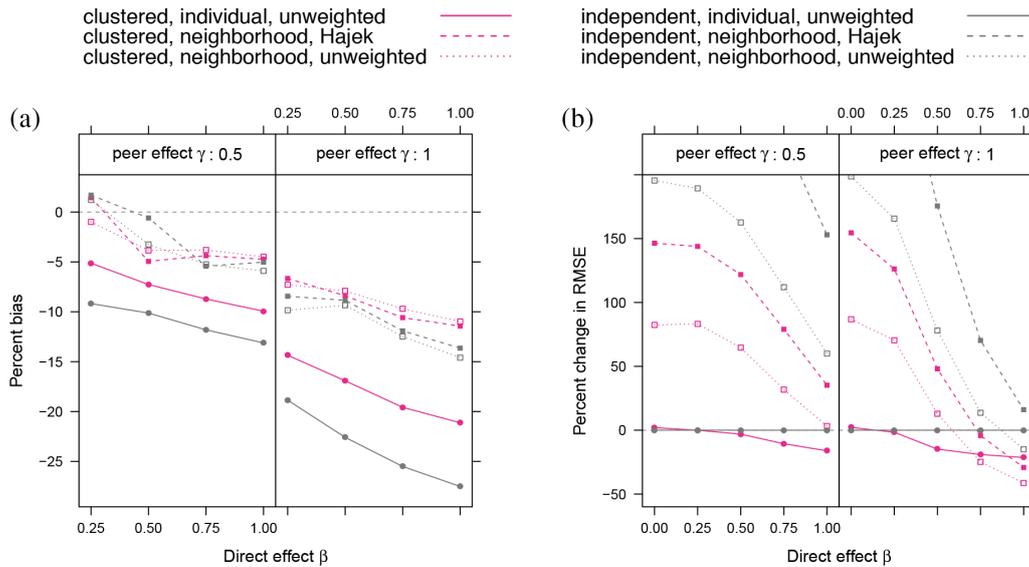


Figure 8 Relative bias (a) and change in RMSE (b) in ATE estimates for different assignment procedures, exposure models, and estimation method, using the degree-corrected block model with community proportion $p_{comm} = 0.8$. Analysis using the exposure model provides additional bias reduction over using graph cluster randomization only – with a cost in variance.

4 Discussion

Recent work on estimating effects of global treatments in networks through experimentation has generally started with a particular set of assumptions about patterns of interference, such as the neighborhood treatment response (NTR) assumption, that make analysis tractable and then developed estimators with desirable properties (e.g., unbiasedness, consistency) under these assumptions [14, 15]. Similarly, Ugander et al. [16] analyzed graph cluster randomization under such assumptions. Unfortunately, these tractable exposure models are also made implausible by the very processes, such as peer effects or social interactions, that are expected to produce interference in the first place. Therefore, we have considered what can be done to reduce bias from interference when such restrictions on interference cannot be assumed to apply in reality.

The theoretical analysis in this paper offers sufficient conditions for this bias reduction through design and analysis in the presence of potentially global interference. To further evaluate how design and analysis decisions can reduce bias, we reported results from simulation studies in which outcomes are produced by a dynamic model that includes peer effects. These results suggest that when networks exhibit substantial clustering and there are both substantial direct and indirect (via peer effects) effects of a treatment, graph cluster randomization can substantially reduce bias with comparatively small increases in variance. Significant error reduction occurred with networks of only 1,000 vertices, highlighting the applicability of these results to experiments with networks of varied sizes – not just massive networks. Since a prior version of this paper, these bias and reduction results have been replicated on a subgraph of a large online social network [41]. Additional reductions in bias can be achieved through the specific estimators used, even though these estimators would usually be motivated by incorrect assumptions about effective treatments.

We have focused on improving estimates of effects of global treatments, but we have not addressed statistical inference that is robust to network dependence. Experimenters will generally want to conduct statistical inference, such as testing null hypotheses of no effects of the treatment or producing confidence intervals for estimated ATEs. Standard methods for randomization inference can be used to test the sharp null hypothesis that the treatment has no effects whatsoever (e.g., permutation tests for clustered randomization). Other hypotheses, such as those about the presence of spillover effects (i.e., SUTVA, the ITR assumption), can be tested exactly either by assuming constant direct effects [19] or using conditional randomization inference for non-sharp null hypotheses [18, 22].

Further work should examine how our results apply to other networks and data-generating processes. The main theoretical analysis and simulations in this paper used models in which outcomes are monotonic in treatment and peer behavior. Such models are a natural choice given many substantive theories; for example, when there are strategic complements in the behavior. On the other hand, non-monotonic responses are expected when having a mix of treated and control neighbors is more different from either having all treated or control neighbors (e.g., inconsistent service offerings to neighbors result in fewer purchases). Other cases where we expect non-monotonic responses include cases where value of engaging in a behavior is determined by a market (e.g., completing a job training program). Our simulations did not include vertex characteristics (besides degree) and prior behaviors, which could play an important role in the bias and variance for different designs and estimators.

Much of the empirical literature that considers peer effects in networks, whether field experiments, e.g. [42–46], or observational studies, e.g. [36, 47], has aimed to estimate peer effects themselves, rather than estimating effects of interventions that work partially through peer effects. This may reflect differences in the quantities of interest for some topics in social science and for decision-making about potential interventions. It is important to note that the intuitions that motivate the clustered designs examined here may not apply to these other estimands (e.g., the case of trying to separately estimate direct and indirect effects). A fruitful direction for future work would involve directly modeling the peer effects involved and then using these models to estimate effects of global treatments, cf. [48].²³ This could substantially expand the range of designs and analysis methods to consider.

Acknowledgement

We are grateful for comments from Edoardo Airoldi, Eytan Bakshy, Thomas Barrios, Guido W. Imbens, Stephen E. Fienberg, Daniel Merl, Cyrus Samii, and participants in the Statistical and Machine Learning Approaches to Network Experimentation Workshop at Carnegie Mellon University and in seminars in the Stanford Graduate School of Business, Columbia University's Department of Statistics, New York University's Department of Politics, UC Davis's Department of Statistics, and Johns Hopkins University's Department of Biostatistics, and from anonymous referees. D.E. was an employee of Facebook while contributing to prior versions of this paper, remains a contractor with Facebook, and has a significant financial interest in Facebook.

Appendix

A Modified graph cluster randomization: Hole punching

We now briefly present a simple modification of graph cluster randomization that adds vertex-level randomness to the treatment assignment, such that some vertex assignments may not match their cluster assignment.

We set

$$W_i \sim \text{Bernoulli}(q_{C(i)})$$

$$X_i \sim \text{Bernoulli}(\eta)$$

$$Z_i = X_i W_{C(i)} + (1 - X_i)(1 - W_{C(i)}).$$

The X_i are independent switching variables that set Z_i to $W_{C(i)}$ with probability η , typically high, and flip the assignment otherwise (“punch a hole”). That is, clusters are assigned to have their vertices predominantly in one of treatment or control. We call this modification *hole punching*, because it inverts the treatment condition of a small fraction of vertices, placing them in a highly isolated treatment position within their cluster. This modification could be useful for estimating differences between direct and peer effects, since it results in many vertices experiencing the direct treatment without peer effects or the peer effects without the direct treatment. It also has the appealing consequence of avoiding exact zero probabilities of assignment to some vectors Z . This is important in cases where one might want to compare outcomes as a function of number of peers assigned to the treatment; otherwise, many of these comparisons would be between conditions to which many vertices could not be assigned. This could also be desirable when testing for interference is a goal, cf. [22], in addition to estimating a global ATE.

B Bias reduction from design: balanced linear case

In this appendix, we consider the linear outcome model under an alternative graph cluster randomization that enforces balance (i.e., equal sample sizes in treatment and control) Assume there is an even number of clusters N_C , each with N/N_C vertices. Pick $N_C/2$ clusters at random and assign them to treatment; assign the remaining clusters to control.

Theorem A.1 Assume we have a outcome model for all vertices $i \in V$ such there exists an N -vector a and $N \times N$ matrix B with non-negative entries $B_{ij} \geq 0$ such that

$$E_U[Y_i(z, U)] = a_i + \sum_{j \in V} B_{ij} z_j. \quad (17)$$

Then for some mapping of vertices to clusters $C(\cdot)$, the absolute bias of $\tau_{\text{ITR}}^d(1, 0)$ when d is graph cluster randomization is less than or equal to the absolute bias when d is independent assignment, with a fixed treatment probability p .

Proof. Using the linear model for Y_i and the definition of τ , we have that the true ATE τ is given by

$$\tau(1, 0) = \mu(1) - \mu(0) = \frac{1}{N} \sum_{ij} B_{ij} \quad (18)$$

for this outcome model. Under balanced graph cluster randomization,

$$\tau_{\text{ITR}}^{\text{bgcr}}(1, 0) = \frac{1}{N} \sum_{ij} B_{ij} \left[\mathbf{1}[C(i) = C(j)] + \mathbf{1}[C(i) \neq C(j)] \left(\frac{N_C/2-1}{N_C-1} - \frac{N_C/2}{N_C-1} \right) \right] \quad (19)$$

$$= \frac{1}{N} \sum_{ij} B_{ij} \left[\mathbf{1}[C(i) = C(j)] - \frac{\mathbf{1}[C(i) \neq C(j)]}{N_C-1} \right]. \quad (20)$$

We can extend this to the case where the mapping of vertices to clusters is random:

$$\tau_{\text{ITR}}^{\text{bgcr}}(1, 0) = \frac{1}{N} \sum_{ij} B_{ij} \left[\Pr(C(i) = C(j)) - \frac{\Pr(C(i) \neq C(j))}{N_C-1} \right]. \quad (21)$$

Separating out B_{ii} :

$$\tau_{\text{ITR}}^{\text{bgcr}}(1, 0) = \frac{1}{N} \left(\sum_i B_{ii} + \sum_{ij; j \neq i} B_{ij} \left(\Pr(C(i) = C(j)) - \frac{\Pr(C(i) \neq C(j))}{N_C-1} \right) \right). \quad (22)$$

If we have uniform probability over all cluster assignments with the same number of vertices per cluster, then for $i \neq j$,

$$\Pr(C(i) = C(j)) = \frac{N/N_C - 1}{N},$$

so

$$\tau_{\text{ITR}}^{\text{bgr}}(1, 0) = \frac{1}{N} \left(\sum_i B_{ii} - \sum_{ij:j \neq i} B_{ij} \frac{N_C}{(N_C-1)N} \right). \quad (23)$$

Under balanced independent assignment, we just have $N_C = N$, so

$$\tau_{\text{ITR}}^{\text{bind}}(1, 0) = \frac{1}{N} \left(\sum_i B_{ii} - \sum_{ij:j \neq i} B_{ij} / (N - 1) \right). \quad (24)$$

Because $B_{ij} \geq 0$, together this implies that $\tau(1, 0) - \tau_{\text{ITR}}^{\text{bgr}}(1, 0) \leq \tau(1, 0) - \tau_{\text{ITR}}^{\text{ind}}(1, 0)$, where monotonicity again dictates that each side of this inequality is positive.

The proof showed that clustering can reduce bias over independent assignment when preserving balance. The relative bias for graph cluster randomization that preserves balance is

$$\begin{aligned} \tau_{\text{ITR}}^{\text{bgr}}(1, 0) / \tau(1, 0) - 1 &= \frac{\sum_{ij} B_{ij} \left[\mathbf{1}[C(i) = C(j)] - \frac{\mathbf{1}[C(i) \neq C(j)]}{N_C - 1} \right]}{\sum_{ij} B_{ij}} - 1 \\ &= \left(1 + \frac{1}{N_C - 1} \right) \left(\frac{\sum_{ij} B_{ij} \mathbf{1}[C(i) = C(j)]}{\sum_{ij} B_{ij}} - 1 \right). \end{aligned} \quad (25)$$

which is the same expression as the relative bias for graph cluster randomization except for the multiplicative factor in the front. For large enough N_C , the relative biases will be identical, and therefore meaningful relative bias reduction occurs depending only on the clustering's relationship to the values B_{ij} , and not whether the sampling scheme preserves balance or not.

C Bias reduction from analysis

Here we restate and prove Theorem 2.3 from the main text. We also consider two possible extensions of this theorem to graph cluster randomization (from independent random assignment), giving a counterexample for one extension and proving an analog of the theorem for the other extension.

Consider functions $g_i(\cdot)$ such that $g_i(Z) = g_i(z)$ just implies that for some subset of vertices J_i we have that $\sum_{j \in J_i} \mathbf{1}[Z_j = z_j] \geq l_i$ and that $Z_i = z_i$. These are conditions such that some subset of size l_i of a set of vertices has treatment assignment matching that in the global treatment vector of interest z . The ITR and NTR assumptions both are of this type, where with ITR J_i is the empty set and with NTR $J_i = \delta(i)$ and $l_i = k_i$, i 's degree. The fractional relaxation of NTR (FNTR) is also of this type, with $J_i = \delta(i)$ and $l_i = \lceil \lambda k_i \rceil$.

Definition 2.2 If we have two such functions $g_i^A(\cdot)$ and $g_i^B(\cdot)$ with the same J_i , and $g_i^A(z) = g_i^A(z')$ implies $g_i^B(z) = g_i^B(z')$, then we say that $g_i^A(\cdot)$ is *more restrictive* than $g_i^B(\cdot)$.

Theorem 2.3 Let $g_i^A(\cdot)$ and $g_i^B(\cdot)$ be functions such that $g_i^A(\cdot)$ is more restrictive than $g_i^B(\cdot)$ for every vertex i , and let independent random assignment be the experimental design. A sufficient condition for estimand $\tau_{g^A}^{\text{ind}}(1, 0)$ to have less than or equal absolute bias than $\tau_{g^B}^{\text{ind}}(1, 0)$, where these estimands are defined by eq. (13), is that we have monotonically increasing responses or monotonically decreasing responses for every vertex with respect to z .

Proof. All expectations are taken with respect to independent random assignment. Assume monotonically increasing responses for every vertex and select an arbitrary vertex i . Let

$$\tilde{Y}_i(z_{J_i}) = E_{Z_{V \setminus J_i}} [Y_i(z_i = 1, Z_{V \setminus J_i} = z_{J_i})]. \quad (26)$$

This quantity is the expectation of the potential outcome for i when $z_i = 1$ and the subset of z corresponding to J_i is set to z_{J_i} . The monotonicity of Y_i carries over to $\tilde{Y}_i(z_{J_i})$.

To reduce the notation in what follows, we define A_i to be the event that $g_i^A(Z) = g_i^A(1)$ and B_i to be the event that $g_i^B(Z) = g_i^B(1)$. We also define $q_i(Z) = \sum_{j \in J_i} \mathbf{1}[Z_j = 1]$. Then

$$\begin{aligned} E[\tilde{Y}_i | A_i] &= \sum_{\substack{|J_i| \\ q \geq l_i^A}} E[\tilde{Y}_i | q_i(Z) = q] P(q_i(Z) = q | A_i), \\ E[\tilde{Y}_i | \neg A_i \wedge B_i] &= \sum_{\substack{l_i^A - 1 \\ q \geq l_i^B}} E[\tilde{Y}_i | q_i(Z) = q] P(q_i(Z) = q | \neg A_i \wedge B_i). \end{aligned} \quad (27)$$

Due to independent random assignment, conditioning on $q_i(Z) = q$ means uniformly sampling a z_{J_i} that has q ones and $|J_i| - q$ zeroes. Consider the following process where $q < |J_i|$. Randomly select a z_{J_i} with q ones and $|J_i| - q$ zeroes. Select at random a 0 element and change it into a 1 to create another vector z'_{J_i} . Record both $\tilde{Y}_i(z_{J_i})$ and $\tilde{Y}_i(z'_{J_i})$ as a pair of values. Due to the monotonicity of \tilde{Y}_i , we have that $\tilde{Y}_i(z_{J_i}) \leq \tilde{Y}_i(z'_{J_i})$.

In this process, z_{J_i} is a uniformly sampled vector that has q ones and $|J_i| - q$ zeroes, and z'_{J_i} is a uniformly sampled vector that has $q + 1$ ones and $|J_i| - (q + 1)$ zeroes. Repeating this process an infinite number of times and using the empirical average of the $\tilde{Y}_i(z_{J_i})$'s computes $E[\tilde{Y}_i | q_i(Z) = q]$. Similarly, the empirical average of the $\tilde{Y}_i(z'_{J_i})$ computes $E[\tilde{Y}_i | q_i(Z) = q + 1]$. Due to the per sample inequality, this shows that $E[\tilde{Y}_i | q_i(Z) = q] \leq E[\tilde{Y}_i | q_i(Z) = q + 1]$. By induction, $E[\tilde{Y}_i | q_i(Z) = q] \leq E[\tilde{Y}_i | q_i(Z) = q']$ when $q < q'$. Combining this with eq. (27),

$$E[\tilde{Y}_i | \neg A_i \wedge B_i] \leq E[\tilde{Y}_i | A_i]. \quad (28)$$

Since the design is independent random assignment, we have that

$$\begin{aligned} E[Y_i | B_i] &= E[\tilde{Y}_i | B_i] \\ &= E[\tilde{Y}_i | A_i] P(A_i | B_i) + E[\tilde{Y}_i | \neg A_i \wedge B_i] P(\neg A_i | B_i). \end{aligned} \quad (29)$$

where in the second equality we have used that g_i^A is more restrictive than g_i^B and that the set J_i is common to both g_i^A and g_i^B . With eq. (28), this implies

$$E[Y_i | B_i] \leq E[\tilde{Y}_i | A_i] = E[Y_i | A_i]. \quad (30)$$

Since this inequality applies for all vertices i , we therefore have that

$$\mu_{g^B}^{\text{ind}}(1) \leq \mu_{g^A}^{\text{ind}}(1), \quad (31)$$

from which we immediately conclude that g^A has less absolute bias for $\mu(1)$ than g^B . An analogous argument applies for $\mu(0)$, proving that $\tau_{g^A}^{\text{ind}}$ has less absolute bias for $\tau(1, 0)$, the average treatment effect.

The proof for monotonically decreasing responses follows when switching the inequalities throughout the above.

This proposition demonstrates how using more restrictive exposure conditions can be helpful in reducing bias, but the proposition just applies to independent assignment, rather than graph cluster randomization. To show why it does not hold for graph cluster randomization, we present the following counterexample with two fractional neighborhood treatment response (FNTR) effective treatments.

Consider some vertex i with no neighbors in its own cluster, and three other clusters present in its neighborhood: one cluster with 10 neighbors, one cluster with one neighbor, and another cluster with one neighbor; call this last neighbor vertex a . Let $Y_i = 1$ when $Z_a = 1$ and $Z_i = 1$, and let $Y_i = 0$ otherwise. Let the less restrictive function $g_i^B(\cdot)$ require that at least 2 neighbors match the global treatment vector, and let the more restrictive function $g_i^A(\cdot)$ require that at least 3 neighbors match; that is, let $l_i^B = 2$ and $l_i^A = 3$. Then under graph cluster randomization, we have $E[Y_i | A_i] \approx 0.5$, but $E[Y_i | B_i] \approx 0.6$. So using the more restrictive function actually increases bias in this somewhat extreme scenario.

While this counterexample demonstrates that using more restrictive exposure conditions of this kind is not always helpful under graph cluster randomization, we do observe bias reduction in our simulations using graph cluster randomization without meeting the sufficient conditions of the theorem. In general, we expect that for bias to increase, there must be heterogeneous effects across heterogeneously sized clusters as in the counterexample above.

In fact, with a redefinition of the exposure conditions, we can provide a similar proposition that does include graph cluster randomization and also encompasses independent assignment as a special case.

Corollary A.2 Consider a fixed set of clusters which will be used for graph cluster randomization. Let function $g_i(\cdot)$, for all vertices i , be such that $g_i(Z) = g_i(z)$ implies that some subset of clusters J_i which do not include i we have that $\sum_{C \in J_i} \mathbf{1}[Z_C = z_C] \geq l_i$ (at least l_i of the clusters in J_i match the global treatment vector z exactly), and $Z_i = z_i$. Consider two such functions where $g_i^A(\cdot)$ is more restrictive than $g_i^B(\cdot)$ for all i . Then a sufficient condition for estimand $\tau_{g^A}^{\text{gcr}}(1, 0)$ to have less than or equal absolute bias than $\tau_{g^B}^{\text{gcr}}(1, 0)$, where these estimands are defined by eq. (13), is that we have monotonically increasing responses or monotonically decreasing responses for every vertex with respect to z .

Proof. This proof is essentially the same as for Theorem 2.3 except \tilde{Y}_i is redefined as

$$\tilde{Y}_i(z_{J_i}) = E_{Z_{V \setminus J_i}} [Y_i(z_{C_i} = 1, Z_{V \setminus J_i}, z_{J_i})], \quad (32)$$

expectations are computed with respect to graph cluster randomization instead of independent treatment assignment, and references to 1's and 0's apply to clusters in J_i .

An important special case of this corollary covers the comparison of FNTR with ITR under graph cluster randomization, since FNTR and ITR can be written as cluster-level exposure conditions of this kind.

D Additional simulations with non-monotonic responses

The simulations reported in the main text, while not satisfying the conditions for graph cluster randomization to be bias reducing given by Theorem 2.1, did nonetheless have responses monotonically increasing with respect to z . We conducted some additional simulations to explore the consequences of graph cluster randomization with non-monotonic responses.

We repeat the simulations with small world networks in Section 3 with the following change. We run the process 1,000 times for all combinations of $\beta \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$ (as before) and $\gamma \in \{0, -0.25, -0.5, -0.75, -1.0\}$; that is, the peer effect now has the opposite sign as the direct effects (when non-zero). We fix $p_{\text{rw}} = 0.01$.

Figure 9 shows the results of these simulations. Graph cluster randomization results in bias reduction, and for larger (in absolute terms) values of γ it again results in error reduction.

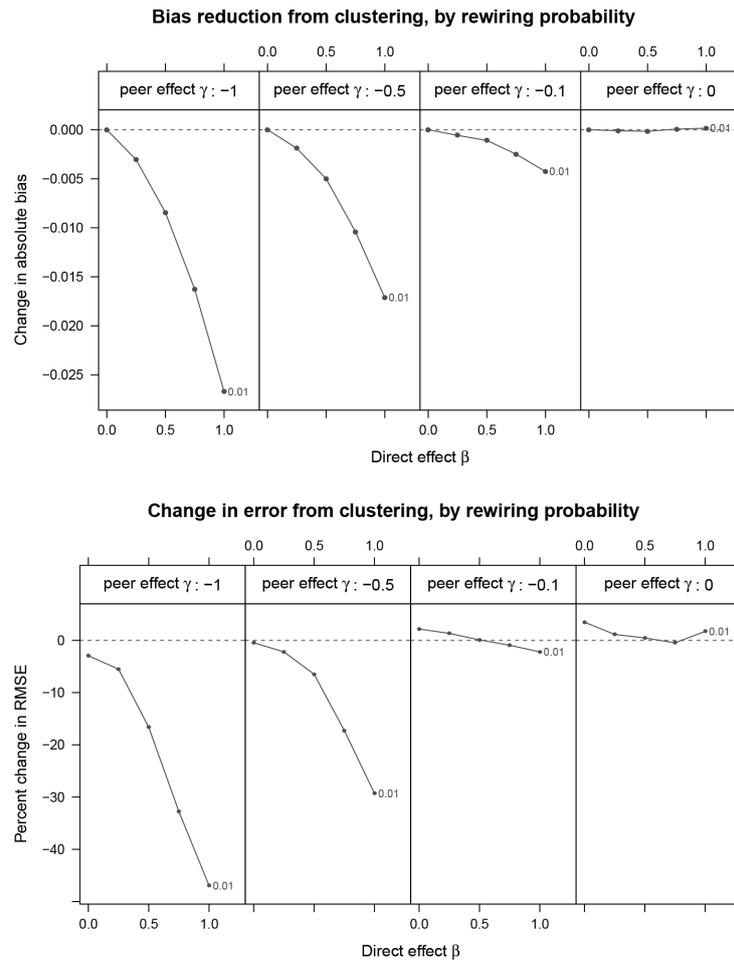


Figure 9 Changes in bias and root-mean-squared-error (RMSE) due from clustered random assignment for small world networks with non-monotonic responses and $p_{rw} = 0.01$.

Notes

¹Estimating this quantity is a common goal. For example, Bond et al. [3] consider the effect of a voter mobilization intervention, such that the aim is to compare voter turnout if everyone (or almost everyone) is assigned to the treatment with turnout if everyone is assigned to the control. Many routine experiments in the Internet industry aim to choose among two or more alternatives based on an estimate of this ATE [4, 5]. Hudgens and Halloran [6] call this the *average total causal effect*. There are other causal quantities that may be of interest, which we do not treat here. Other authors consider decompositions of effects into various direct and indirect effects of the treatment [6–10].

²To allow for other sources of dependence besides interference, we allow Y_i to be random, rather than fixed. In some cases the experimenter may regard these potential outcomes as fixed.

³This is closely connected to what Holland [11] regards as the fundamental problem of causal inference – that one can only observe a unit's response under a single treatment. The difference is that here we can only observe *all* units' responses under a single global treatment.

⁴While we limit the analysis here to cases where the measured network and the network through which the interaction occur are the same, the methods examined here may also substantially reduce bias in when using a network observed with error.

⁵This intuition applies for the global ATE considered here. For estimating other quantities, such as the magnitude of the interference, very different intuitions apply.

⁶For the purposes of this paper, we assume that the network is fixed over the timescale of the experiment.

⁷A normal but minor exception occurs when forcing a specific number of units within a block to be assigned to each of treatment and control; this produces negative dependence between units in the same block. This includes global balancing of sample sizes in treatment and control, as in a *completely randomized experiment*, rather than *simple randomized experiment*, as a special case.

⁸Thus, the present analysis can be applied to cases where the experimenter observes only community membership and so their estimators may be biased by interference via an unobserved network.

⁹The vertices might interact without necessarily violating the ITR assumption. This can occur, for example, when vertices interact in one period, and then are affected by treatment assignment, while no longer interacting. This is why we define $f_i(\cdot)$ as being a function from \mathbb{U}^N rather than just \mathbb{U} .

¹⁰Manski [14] calls these models of simultaneous endogenous choice a “system of structural equations”. But because these equations are simultaneous, they are not structural in the sense of corresponding to a directed acyclic graph (DAG) given a causal interpretation [29]. However, we can regard these equations as specifying an equilibrium that arises out of some unknown dynamic process. We prefer to work with a posited dynamic process, which may or may not be in equilibrium when we observe it, cf. [30].

¹¹Note that this specification in terms of “direct” effect could be compatible with various data-generating processes that involve “indirect” effects – at least on a short time scale.

¹²And similarly for the assumed exposure models in Aronow and Samii [15] and Ugander et al. [16].

¹³Note that in the above treatment, the mapping of vertices to clusters is not random, so any mapping is bias reducing.

¹⁴As shown by Ugander et al. [16], assuming NTR, that the graph satisfies a restricted growth condition, and non-stochastic outcomes, this implies that an experimental design with $O(N)$ clusters of size $O(1)$ will produce NTR-based estimators that are asymptotically consistent.

¹⁵It is precisely the effective treatment assumption that allows generalization from a single sampled z to the behavior at z_1 and z_0 .

¹⁶Of course, ITR can also be analyzed with any choice of J_i , including the empty set.

¹⁷Under complete randomization (i.e., when a fixed number of clusters are randomized to treatment, this estimator can be biased if clusters vary in size [37, 38].

¹⁸Simulations for the Louvain method [40] for community detection, not reported here, are qualitatively similar to those for ϵ -net clustering, but generally resulted in more bias reduction but also larger variance increases, as expected by this method’s resolution limit.

¹⁹We also repeated these simulations with the small-world networks for $T = 10$. The results were qualitatively similar.

²⁰As a variance-reduction strategy for comparisons between designs and true ATE, we use common random numbers throughout the simulations where possible. In particular, for generating observed outcomes, the first instance of each configuration uses the same seed s_1 , the second instance of each configuration uses the same seed s_2 , and so on.

²¹It is possible for no vertices to meet this condition for treatment or for control. In this case, the estimator is undefined. If this occurs, we expect that experimenters would re-randomize or modify their analysis plan. For the results shown here, we exclude simulations where this occurred, which corresponds to re-randomizing. This did not occur for graph clustered randomization. For independent assignment, this occurred for one of the 5,000 simulations for rewiring probability $p = 0.01$ (i.e., the results shown in Figure 5).

²²Note that with $p_{\text{comm}} = 0.8$ and the chosen degree distribution, the DCBM networks have an average clustering coefficient of approximately 0.095 and average transitivity of approximately 0.091. This is similar to that of small-world networks with $p_{\text{rw}} = 0.5$. This observed bias and error reduction is likewise comparable to that observed with those small-world networks.

²³In the language of contemporary econometrics, one could describe the present work as taking a “reduced form” approach to this problem, rather than trying to learn about the underlying data generating process through estimating “structural” parameters.

References

- Manski CF. Economic analysis of social interactions. *J Econ Perspect*. 2000;14:115–136.
- Moffitt RA. Policy interventions, low-level equilibria, and social interactions. In: Durlauf SN, Young HP, editors. *Social Dynamics*. Cambridge, MA: MIT Press, 2001:45–82.
- Bond RM, Fariss CJ, Jones JJ, Kramer AD, Marlow C, Settle JE, et al. A 61-million-person experiment in social influence and political mobilization. *Nature*. 2012;489:295–298.
- Bakshy E, Eckles D, Bernstein MS. Designing and deploying online field experiments. In *Proceedings of the 23rd International Conference on World Wide Web*. 2014;283–292.
- Kohavi R, Longbotham R, Sommerfield D, Henne RM. Controlled experiments on the web: Survey and practical guide. *Data Min Knowl Discovery*. 2009;18:140–181.
- Hudgens MG, Halloran ME. Toward causal inference with interference. *J Am Stat Assoc*. 2008;103(482):832–842.
- Sobel ME. What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *J Am Stat Assoc*. 2006;101:1398–1407.
- Tchetgen EJ, VanderWeele TJ. On causal inference in the presence of interference. *Stat Methods Med Res*. 2012;21:55–75.
- Toulis P, Edward K. Estimation of Causal Peer Influence Effects *Proceedings of The 30th International Conference on Machine Learning*. *Mach Learn Res W&CP*. 2013;28(3):1489–1497.
- Baird S, Bohren JA, McIntosh C, Ozler B. (2014). *Designing Experiments to Measure Spillover Effects* Policy Research Working Paper–sWashington, D.C.The World Bank.
- Holland PW. Causal inference, path analysis, and recursive structural equations models. *Sociol Methodol*. 1988;18:449–484.
- Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol*. 1974;66:688–701.
- Cox DR. *Planning of Experiments*. Hoboken, New Jersey: Wiley, 1958.
- Manski CF. Identification of treatment response with social interactions. *Econom J*. 2013;16 S1–S23.
- Aronow P, Samii C. Estimating average causal effects under general interference (Working Paper). 2014. Available at <http://arxiv.org/abs/1305.6156>.
- Ugander J, Karrer B, Backstrom L, Kleinberg JM Graph cluster randomization: network exposure to multiple universes. In: Dhillon IS, Koren Y, Ghani R, Senator P, Bradley J, Parekh R, et al., editors. *Proceedings of the 19th ACM SIGKDD international conference on knowledge discovery and data mining (KDD ’13)*. New York: ACM, 2013:329–337.
- Rosenbaum PR. Interference between units in randomized experiments. *J Am Stat Assoc*. 2007;102.
- Aronow PM. A general method for detecting interference between units in randomized experiments. *Sociol Methods Res*. 2012;41:3–16.
- Bowers J, Fredrickson MM, Panagopoulos C. [Reasoning about interference between units: a general framework](#). *Polit Anal*. 2013;21:97–124.
- Coppock A, Sircar N. An experimental approach to causal identification of spillover effects under general interference (Working Paper). 2013. Available at https://nsircar.files.wordpress.com/2013/02/coppocksircar_20130718.pdf.
- Choi DS. Estimation of monotone treatment effects in network experiments. *J Am Stat Assoc* 2014. DOI:10.1080/01621459.2016.1194845.
- Athey S, Eckles D, Imbens GW. Exact p-values for network interference. *J Am Stat Assoc*. 2016. Available at <http://arxiv.org/abs/1506.02084>.
- Walker D, Muchnik L. Design of randomized experiments in networks. *Proc IEEE*. 2014;102:1940–1951.

24. Watts DJ, Strogatz SH. Collective dynamics of 'small-world' networks. *Nature*. 1998;393:440–442.
25. Karrer B, Newman ME. Stochastic blockmodels and community structure in networks. *Phys Rev E*. 2011;83:016107.
26. Fortunato S. Community detection in graphs. *Phys Rep*. 2010;486:75–174.
27. Newman ME. Modularity and community structure in networks. *Proc Natl Acad Sci*. 2006;103:8577–8582.
28. Fortunato S, Barthelemy M. Resolution limit in community detection. *Proc Natl Acad Sci*. 2007;104:36–41.
29. Pearl J. *Causality: Models, Reasoning and Inference*. Cambridge, UK: Cambridge University Press, 2009.
30. Young HP. *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ, USA: Princeton University Press, 1998.
31. Blume LE. [The statistical mechanics of best-response strategy revision](#). *Game Econ Behav*. 1995;11:111–145.
32. Jackson MO. *Social and Economic Networks*. Princeton, NJ, USA: Princeton University Press, 2008.
33. Manski CF. Identification of endogenous social effects: The reflection problem. *Rev Econ Stud*. 1993;60:531–542.
34. Lee LF. Identification and estimation of econometric models with group interactions, contextual factors and fixed effects. *J Econom*. 2007;140:333–374.
35. Bramoulle Y, Djebbari H, Fortin B. Identification of peer effects through social networks. *J Econom*. 2009;150:41–55.
36. Goldsmith-Pinkham P, Imbens GW. Social networks and the identification of peer effects. *J Bus & Econ Stat*. 2013;31:253–264.
37. Middleton JA. [Bias of the regression estimator for experiments using clustered random assignment](#). *Stat Probab Lett*. 2008;78:2654–2659.
38. Middleton JA, Aronow PM. [Unbiased Estimation of the Average Treatment Effect in Cluster-Randomized Experiments](#). *Statistics, Politics and Policy*. 2015;6(1-2):39–75.
39. Ugander J, Karrer B, Backstrom L, Marlow C. 2011. Available at <http://arxiv.org/abs/1111.4503>.
40. Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. Fast unfolding of communities in large networks. *J Stat Mech Theor Exp*. 2008;2008:10008.
41. Gui H, Xu Y, Bhasin A, Han J. Network A/B testing: from sampling to estimation. *Proceedings of the 24th international conference on World Wide Web*. New York: ACM, 2015;399–409.
42. Aral S, Walker D. [Creating social contagion through viral product design: a randomized trial of peer influence in networks](#). *Manage Sci*. 2011;57:1623–1639.
43. Bakshy E, Eckles D, Yan R, Rosenn I. Social influence in social advertising: evidence from field experiments. *Proceedings of the 13th ACM Conference on Electronic Commerce (EC '12)*. New York: ACM, 2012;146–161.
44. Bakshy E, Rosenn I, Marlow C, Adamic L. The role of social networks in information diffusion. New York: ACM, 2012;519–528. In *Proceedings of the 21st international conference on World Wide Web (WWW '12)*. DOI:10.1145/2187836.2187907.
45. Bapna R, Umyarov A. [Do your online friends make you pay? A randomized field experiment on peer influence in online social networks](#). *Management Science*. 2015;61(8):1902–1920.
46. Eckles D, Kizilcec RF, Bakshy E. Estimating peer effects in networks with peer encouragement designs. *Proc Natl Acad Sci*. 2016;113:7316–7322.
47. Aral S, Muchnik L, Sundararajan A. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proc Natl Acad Sci*. 2009;106:21544–21549.
48. van der Laan MJ. Causal inference for a population of causally connected units. *J Causal Inference*. 2014;2(1):1–62.