

SCIENTIFIC REPORTS



OPEN

Incorporating Contact Network Structure in Cluster Randomized Trials

Patrick C. Staples¹, Elizabeth L. Ogburn² & Jukka-Pekka Onnela¹

Received: 14 June 2015

Accepted: 27 October 2015

Published: 03 December 2015

Whenever possible, the efficacy of a new treatment is investigated by randomly assigning some individuals to a treatment and others to control, and comparing the outcomes between the two groups. Often, when the treatment aims to slow an infectious disease, clusters of individuals are assigned to each treatment arm. The structure of interactions within and between clusters can reduce the power of the trial, i.e. the probability of correctly detecting a real treatment effect. We investigate the relationships among power, within-cluster structure, cross-contamination via between-cluster mixing, and infectivity by simulating an infectious process on a collection of clusters. We demonstrate that compared to simulation-based methods, current formula-based power calculations may be conservative for low levels of between-cluster mixing, but failing to account for moderate or high amounts can result in severely underpowered studies. Power also depends on within-cluster network structure for certain kinds of infectious spreading. Infections that spread opportunistically through highly connected individuals have unpredictable infectious breakouts, making it harder to distinguish between random variation and real treatment effects. Our approach can be used before conducting a trial to assess power using network information, and we demonstrate how empirical data can inform the extent of between-cluster mixing.

In order to determine how effective a treatment is, it is common to randomly assign test subjects to different treatment arms. In one arm, subjects receive the experimental treatment, and subjects in the other arm receive usual care or a placebo. Randomization helps to ensure that the treatment is the cause of any difference in outcomes between the subjects in the two treatment arms, as opposed to some pre-treatment characteristics of the individuals. If the treatment is effective, the probability that a trial will find a statistically significant difference attributed to the treatment is called the *power* of the trial¹. Adequate power requires a sufficiently large number of subjects to be tested, which can be expensive or infeasible. Underpowered studies are not only less likely to find a true relationship if one exists, but they are also more likely to erroneously conclude that an effect exists when it does not^{2,3}. In order to control the probability of these errors, it is important to be able to accurately assess power before conducting a study.

When designing a randomized trial, we may not want or be able to randomly assign individuals to treatment. Individuals may be members of a cluster with complex interactions, which makes it infeasible or unethical to assign some individuals within a cluster to treatment and others to control. For example, the spread of HIV from infected to uninfected individuals in a small village might be slowed by offering its members information about safer sexual practices. In this case, it may be difficult or unethical to keep treated individuals' sex partners from sharing information or resources. We may instead choose to randomly select villages to participate in this regime, where villages correspond to naturally occurring clusters, and to compare HIV infection rates between treatment and control villages. This type of experiment is called a *Cluster Randomized Trial* (CRT)⁴⁻⁷.

¹Department of Biostatistics, Harvard University, Boston, MA 02115, USA. ²Department of Biostatistics, Johns Hopkins University, Baltimore MD, 21205, USA. Correspondence and requests for materials should be addressed to P.C.S. (email: patrickstaples@fas.harvard.edu)

The correlation in outcomes of individuals within a cluster (e.g. HIV infection statuses) is known to reduce the power of a trial⁵. This correlation is generally summarized by a single parameter, called the *Intraclass Correlation Coefficient* (ICC)⁴, which is the average pairwise correlation of outcomes within clusters. This measure assumes that the correlation in outcomes for any two individuals within a cluster is identical. However, the structure of relationships within a cluster can be heterogeneous, and power may depend on that structure, which is not captured by the ICC. Usually, this structure is either ignored⁸ or analysis is performed using methods that allow it to be left unspecified⁹. Furthermore, individuals are often likely to interact with others not only in the same cluster but also in other clusters. *Cross-contamination*¹⁰ or *interference*¹¹ occurs when subjects' outcomes depend not only on the treatment to which they are assigned, but on the treatment assignments of other subjects as well. This can reduce the difference in outcomes between treated and untreated clusters, thereby decreasing power¹². For example, economic ties may exist between villages, the residents of which might then share information related to the treatment. In the context of infectious spread through clusters, cross-contamination would occur if infectious contact takes place across clusters. If the treatment succeeds in slowing the infection rate in the treatment cluster, mixing between clusters will decrease the difference between outcomes across clusters, so the power to detect a treatment effect will decrease and the probability of a false discovery will increase. This must be addressed either by adding more clusters to the trial or increasing cluster sizes, both of which could be difficult and costly. This issue is also often left unaddressed^{13,14}.

The effect of within-cluster structure and between-cluster mixing may depend on the type of infection spreading through each cluster. For example, a highly contagious infectious disease like the flu can spread more efficiently through more highly connected individuals¹⁵. Other infectious diseases, such as a sexually transmitted disease, can only be transmitted to one person at a time, no matter how many partners one has. The number of individuals whom an infected person may infect at a given time is the person's *infectivity*. This quantity likely differs from person to person, and it depends crucially on the transmission dynamics of the disease.

In this paper, we study, via simulation, the effect of within-cluster structure, the extent of between-cluster mixing, and infectivity on statistical power in CRTs. We simulate the spread of an infectious process and investigate how power is affected by features of the process. Specifically, we consider two infections with different infectivities spreading through a collection of clusters. We use a *matched-pairs design*, wherein clusters in the study are paired, and each pair has one cluster assigned to treatment one to control⁷. We model the complex within-cluster correlation structure as a network in which edges represent possible transmission pathways between two individuals, comparing results across three different well-known network models. To model one type of cross-contamination, we introduce a single parameter γ that summarizes the extent of mixing between the two clusters comprising each cluster pair. This approach departs from standard power calculations for CRTs, in which the researcher applies a formula that determines the required sample size as a function of the number and size of clusters, the ICC, and the effect size¹⁶. Figure 1 depicts the different assumptions behind these two approaches. We show that our measure of mixing between clusters can have a strong effect on experimental *power*, or the probability of correctly detecting a real treatment effect. We also show that within-cluster structure can affect power for certain kinds of infectivity. We contrast this method to standard power calculations. We end by demonstrating how to assess between-cluster mixing before designing a hypothetical CRT, using a network dataset of inter-regional cell phone calls.

Results

Simulation of cluster randomized trials. We simulate both within-cluster structure and between-cluster mixing using network models. We simulate pairs of clusters with each cluster in each pair initially generated as a stand-alone network. We examine the Erdős-Rényi (ER)¹⁷, Barabási-Albert (BA)¹⁸, and stochastic blockmodel (SBM)¹⁹ random networks, and we simulate 2C clusters comprised of n nodes each. In order to explicitly allow for between-cluster mixing, we define a between-cluster mixing parameter γ as the number of network edges between the treatment cluster and the control cluster, divided by the total number of edges in the cluster pair. To ensure that proportion γ of the edges are shared across clusters, we perform degree-preserving rewiring²⁰ within each of the C cluster-pairs until proportion γ edges are shared between clusters. We then use a compartmental model to simulate the spread of an infection across each cluster pair²¹. All nodes are either susceptible (S) or infected (I), and nodes may only transition from S to I . The number of neighbors each node can potentially infect at any given time is called its *infectivity*. We consider both unit and degree infectivity, for which infected nodes may contact one or all of their neighbors at a given time, respectively. Treated and control clusters infect their neighbors with equal probability under the null hypothesis, and infected individuals in treatment clusters infect with reduced probability under the alternative hypothesis. Finally, we analyze the resulting trial under two different analysis scenarios, and we juxtapose our findings with a standard power calculation¹⁶. Full simulation details are found in Methods.

We begin by showing the effect of the mixing parameter γ on the infection risk ratios (see methods) between treated and untreated clusters. The means and standard deviations of simulated risk ratios observed under Scenario 1 are presented in Fig. 2.

For both kinds of infectivity, neither the heavy-tailed degree distribution of the BA network nor the within-cluster community structure of the SBM network dramatically impacts the differences between

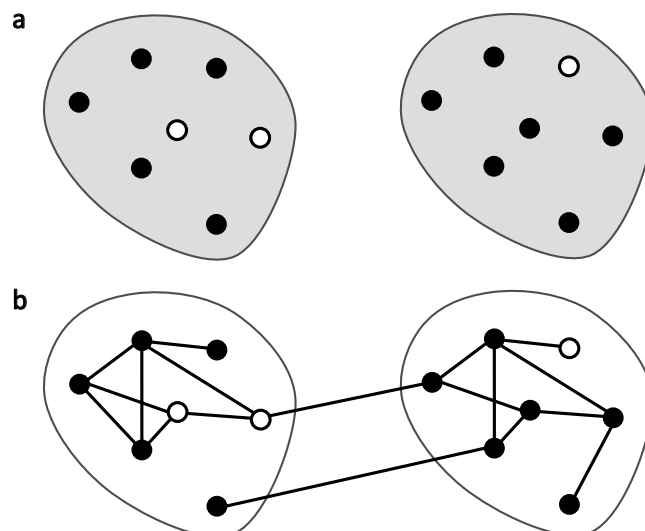


Figure 1. A schematic comparing the Intracluster Correlation Coefficient (ICC) approach to the design of this study. Each panel shows a cluster pair, and each enclosure represents a cluster. Panel (a) depicts cluster pair outcomes (circle colors) which are correlated (gray shading) within each cluster according to the ICC. In contrast, Panel (b) shows specific relationships (contact network ties) among individuals both within and between the two clusters, and outcomes among them will depend on an infection spreading only through these ties. We show that modeling both contact network structure and the spreading process explicitly rather than modeling correlations across outcomes results in new findings about power in CRTs.

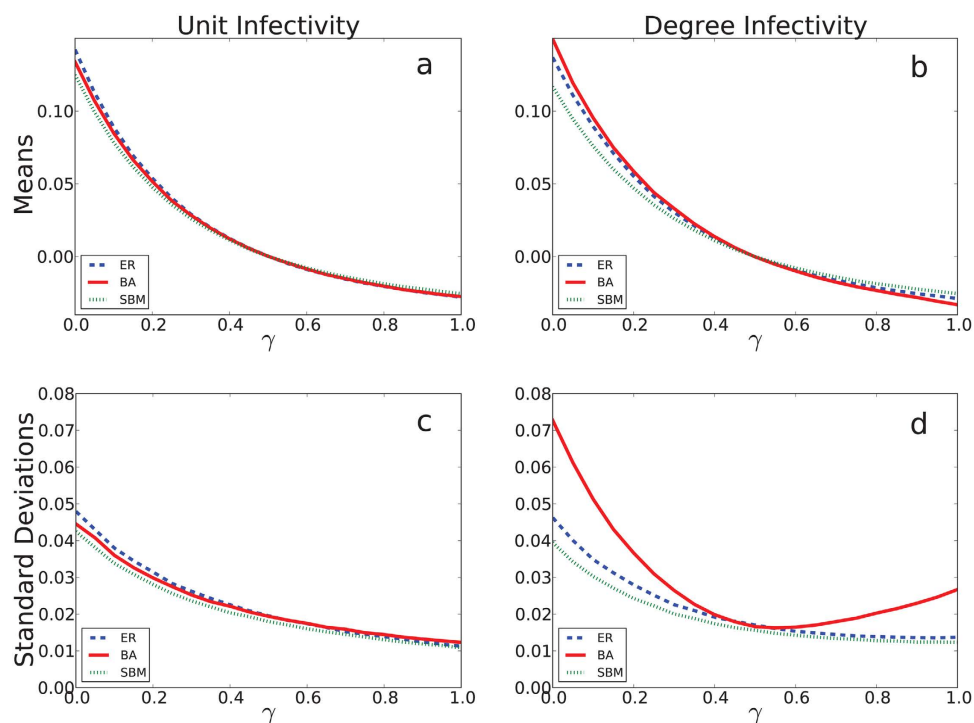


Figure 2. The log risk ratio means and standard deviations under Scenario 1. The rows correspond to the means (Panels (a,b)) and standard deviations of the log risk ratio (Panels (c,d)), shown on the y axis. The x -axis is the value of the mixing parameter γ , and each curve represents one of the three within-cluster network structures. The left column shows the spread of an infection in which an infected node may only infect one neighbor per time step (unit infectivity), whereas the right column assumes one may spread an infection to each of their neighbors (degree infectivity). We see that network topology has an effect on the variation of the log rate ratio only in the latter case.

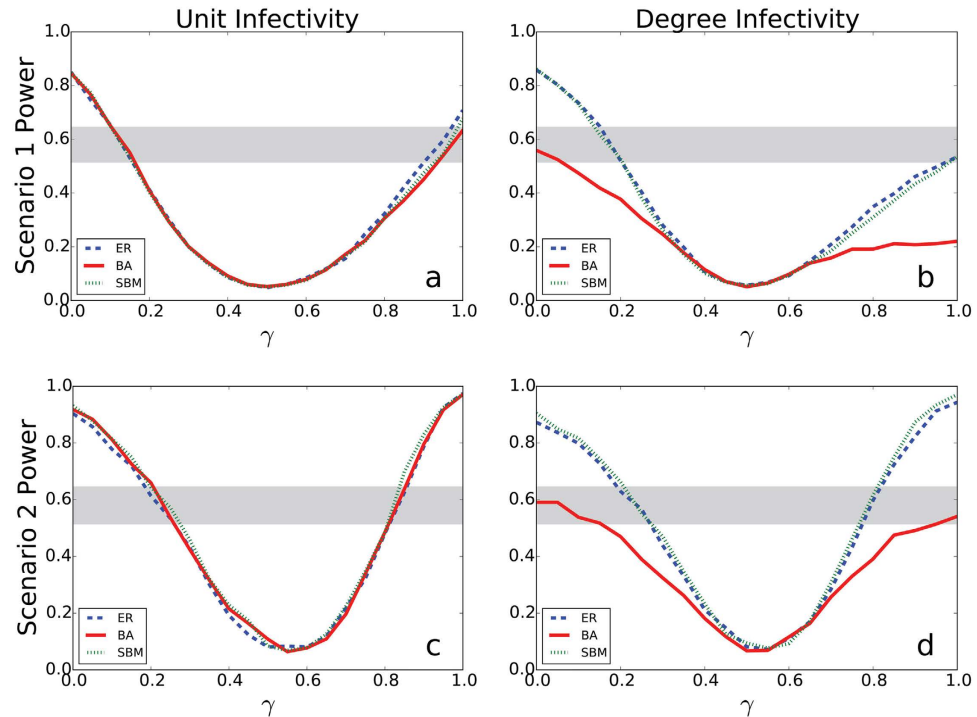


Figure 3. Estimated power for each scenario. The blue (thick dashed), red (solid), and green (thin dashed) lines represent the ER, BA, and SBM network models, respectively. The top row shows results for Scenario 1, and the bottom row shows results for Scenario 2. The left column shows unit infectivity, and the right column shows degree infectivity. The horizontal gray bars represent the expected power using the standard approach for a range of plausible values for the ICC (see Methods for details).

the proportion of infections in the treated and controlled clusters in each pair (top row) compared to the ER network. The differences between the risk of infections in the treated and untreated cluster pairs decreases as mixing increases, and reverses direction when $\gamma > 1/2$. This is expected because for this range of between-cluster mixing, infected individuals in the treatment cluster are more likely to contact members of the untreated cluster and vice versa, which is unlikely in practice but is included here for completeness. In almost all cases, the variation in the simulated studies' average log risk ratio decreases uniformly as γ increases, which suggests that increasing the amount of mixing across communities results in less variation in the average rate of infections. However, the BA network is an exception. Under degree infectivity, when individuals can infect everyone to whom they are connected in a single time step, an infected node with large degree may spread its infection to each of its contacts at a single time point, which can cause a very fast outbreak. However, highly-connected individuals are rare, so in this case outbreaks are large but infrequent, increasing the variation in observed differences between treated and untreated clusters. This variation means that more clusters are required to estimate the average treatment effect with any precision. In other words, rare outbreaks make it harder to distinguish whether differences between the treatment arm and control arm are due to treatment or to a chance outbreak occurring in either arm. Therefore, under degree infectivity, the BA network results in less power than the SBM or ER networks, which shows that within-cluster network structure can impact the power to detect treatment effects in CRTs for certain kinds of infections.

For the two analysis scenarios described in Methods, we can directly estimate empirical power as the proportion of simulations resulting in the rejection of the null hypothesis at the $\alpha = 0.05$ level under the alternative for a range of mixing values γ . Our results, as well as a comparison with the standard approach, are summarized in Fig. 3.

In all settings, power is lowest when $\gamma \approx 1/2$, with approximately the same number of edges between clusters as within them. Scenarios 1 and 2 (the top and bottom rows, respectively) show few differences from one another, which suggests that the two strategies for significance testing tend to give qualitatively similar results. Unit infectivity (lefthand column) shows no differences in power among network types. This is not the case for degree infectivity (righthand column), in which the BA network shows less power than the other networks, for the reasons discussed above. Finally, the gray bars indicate that when no mixing is present, standard power calculations are conservative for all network types we studied, and no sample size adjustment may be needed. However, moderate to severe between-cluster mixing can greatly overestimate expected power. In the case of the BA network and degree infectivity, the standard approach always overestimates trial power.

	C = 5						C = 10						C = 20						
	Identical			Degree			Identical			Degree			Identical			Degree			
n = 100	0.13	0.14	0.14	0.12	0.10	0.12	0.22	0.22	0.23	0.19	0.16	0.20	0.39	0.39	0.38	0.34	0.28	0.37	$\gamma = 0.0$
	0.10	0.10	0.10	0.11	0.09	0.10	0.15	0.17	0.15	0.18	0.12	0.16	0.27	0.27	0.24	0.26	0.21	0.26	$\gamma = 0.1$
	0.08	0.08	0.07	0.09	0.07	0.09	0.09	0.13	0.11	0.11	0.11	0.13	0.15	0.15	0.15	0.22	0.15	0.18	$\gamma = 0.2$
	0.07	0.06	0.06	0.07	0.06	0.07	0.08	0.09	0.06	0.08	0.08	0.08	0.10	0.09	0.10	0.11	0.11	0.11	$\gamma = 0.3$
n = 300	0.34	0.33	0.32	0.33	0.20	0.33	0.57	0.55	0.57	0.59	0.34	0.56	0.85	0.86	0.86	0.87	0.57	0.87	$\gamma = 0.0$
	0.21	0.22	0.21	0.25	0.13	0.27	0.39	0.39	0.39	0.46	0.28	0.46	0.63	0.65	0.65	0.71	0.44	0.70	$\gamma = 0.1$
	0.16	0.15	0.15	0.17	0.14	0.19	0.23	0.22	0.22	0.28	0.19	0.30	0.41	0.43	0.40	0.52	0.38	0.44	$\gamma = 0.2$
	0.08	0.08	0.08	0.12	0.10	0.10	0.13	0.14	0.13	0.14	0.13	0.17	0.21	0.20	0.19	0.28	0.25	0.27	$\gamma = 0.3$
n = 1000	0.78	0.80	0.76	0.84	0.39	0.81	0.97	0.97	0.97	0.98	0.75	0.98	1.00	1.00	1.00	1.00	0.95	1.00	$\gamma = 0.0$
	0.61	0.57	0.59	0.69	0.38	0.67	0.85	0.86	0.85	0.93	0.61	0.91	0.99	0.99	0.99	1.00	0.91	1.00	$\gamma = 0.1$
	0.39	0.37	0.36	0.47	0.31	0.51	0.62	0.60	0.59	0.76	0.52	0.76	0.89	0.90	0.87	0.97	0.83	0.96	$\gamma = 0.2$
	0.15	0.19	0.16	0.30	0.21	0.28	0.31	0.33	0.36	0.49	0.36	0.45	0.58	0.56	0.57	0.74	0.62	0.73	$\gamma = 0.3$

ER BA SBM

Table 1. Experimental power in our simulation framework for different sizes and numbers of cluster pairs, n and C , respectively, for Scenario 1. Each cell shows output for 3,000 simulations of each combination of n and C , all three within-cluster structures, various values of mixing parameter γ , and both unit and degree infectivity. The results are similar for Scenario 2.

1) For all clusters in the study:
a) Ascertain or conjecture within-cluster network structure and between-cluster mixing for clusters.
2) Repeat several times:
a) Simulate a collection of networks consistent with cluster structure and mixing properties.
b) Propagate an infectious spread through networks.
3) Assess the empirical power of the simulation using the outcomes from the spreading process.

Table 2. Our simulation algorithm used to assess the effect of within-cluster structure, between-cluster mixing and infectivity on statistical power.

Size and number of study clusters. Our results so far have shown how power in CRTs is affected by between-cluster mixing, within-cluster structure, and infectivity. Next, we show how power relates to other trial features, namely the size and number of clusters, n and C , respectively. The results are qualitatively similar for Scenarios 1 and 2, and the results shown in Table 1 are for Scenario 1. Table 2 shows results for each combination of a range of cluster sizes $n = \{100, 300, 1000\}$ and numbers $C = \{5, 10, 20\}$ as a 3×3 grid of pairs of cells. Each cell pair is a side-by-side comparison of results for unit infectivity (lefthand cell) and degree infectivity (righthand cell). Each cell shows simulated results for within-cluster structure (columns) as well as amount of between-cluster mixing (rows). Considering the case of $C = 10, n = 300$ (the middle-most cell pair), we notice a few trends. We see that increasing mixing (looking down each column) decreases power in all cases. We can directly compare the two types of infectivity (comparing cells in the pair), and see that all the entries are similar except for the BA network (middle column). For BA networks, power is much lower for degree infectivity spreading compared to unit infectivity. This suggests that CRTs with network structure similar to BA networks can have substantially less power when the infection spreads in proportion to how connected each node is. Finally, we may compare studies of differing cluster numbers and sizes (comparing cell pairs), and see qualitatively similar results: in each case, more or larger clusters in the study (cell pairs further down or right) result in more power overall. When power is very high (bottom-right cell pair), within-cluster structure affects results less. Therefore, careful consideration of expected power is most important when trial resources are limited, which is often the case in practice.

Real-world data and the extent of mixing. Finally, we show how our mixing parameter can be estimated using data in the planning stages of an idealized CRT. Sometimes the entire network structure between individuals in a prospective trial is known beforehand, such as the sexual contact network on Likoma Island²². In this case, between-cluster mixing can be estimated using Equation 3. In other trials, perhaps only partial information is known, like the degree distribution¹⁸ and/or the proportion of ties between clusters. In this case, clusters can be generated that preserve partial network information such as degree distribution^{23,24}, and degree-preserving rewiring can be performed until proportion γ of ties between clusters is observed, where this quantity is estimated from the network data, if possible.

The structure of calls between cell phones is often persistent over time²⁵ and indicative of actual social relationships²⁶. We use a network of cell phone calls <http://www.pnas.org/content/104/18/7332.abstract>

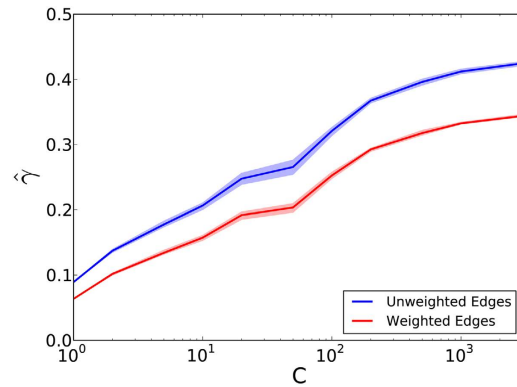


Figure 4. A log-linear plot displaying empirical values of mixing parameter γ . The y axis shows the mean and (2.5, 97.5) quantiles of these estimates. The x axis in each panel corresponds to a range of cluster numbers C .

as a proxy for a contact network, and use our definition of between-cluster mixing to estimate the amount of mixing between hypothetical clusters. The dataset consists of calls made between cell phones of a large mobile carrier within a quarter year, comprising 2,386,888 nodes (individuals) and 19,616,208 edges. Individual phone numbers were anonymized, and we only report results for the number of individuals and calls within or between billing zip codes.

The dataset contains phone calls originating from $Z = 3806$ different zip codes, and we define a cluster as a collection of zip codes that are spatially close to one another. Because zip codes are numerically assigned according to spatial location, we assume that zip codes that are numerically contiguous to each other are also close to each other spatially. Therefore, zip code $z = 1, \dots, Z$ assigned to cluster $c_z = 1, \dots, 2C$ is

$$c_z := \left\lceil \frac{z}{Z} 2C \right\rceil \quad (1)$$

where $2C$ is the total number of clusters in the trial, and $\lceil \cdot \rceil$ is the ceiling function. Once the number of clusters $2C$ is specified, clusters may be paired, with one cluster in each pair randomized to a hypothetical treatment, and the other to the control condition.

Next, we estimate mixing parameter γ for this dataset. We consider two definitions for the number of edges shared between individuals, one in which they are unweighted and one in which they are weighted by the number of calls between them. We define between-cluster mixing parameter γ in terms of these edges and cluster membership (see Methods). For a range of numbers of cluster pairs C , we cluster all Z zip codes into $2C$ clusters, and randomize one cluster in each pair to a hypothetical treatment, and the other to a control. For 200 randomizations, we calculate the between-cluster mixing parameter γ . We examine the relationship between γ and the number of clusters C . The mean and (2.5, 97.5) percentiles of these estimates as a function of the number of clusters number C are shown in Fig. 4.

Figure 4 displays a number of distinct trends. As the number of clusters increases, fewer of the total zip codes are included in each cluster, and the number of calls between clusters increases. This means that individuals are more likely to call others in zip codes geographically closer to them, which has been confirmed in other phone communication networks²⁷. Between-cluster mixing unweighted by the number of calls (blue) results in higher estimates of γ than weighted (red), which means that when individuals call others outside their cluster, they tend to call those people less than others they call within their cluster. There is significant between-cluster mixing for all values of C , implying that between-cluster mixing would significantly decrease the power of a trial that assumes each cluster to be independent ($\gamma = 0$). Furthermore, as the number of clusters increases, the average cluster size decreases, and mixing reaches a maximum of $\gamma = 0.45$. Extrapolating from our simulation framework, power could be reduced dramatically in this case.

Discussion

Before conducting a trial, it is important to have an estimate of statistical power in order to assess the risks of failing to find true effects and of spurious results. If individuals belong to interrelated clusters, randomly assigning them to treatment or control may not be a palatable option, and CRTs can be used to test for treatment effects. Power in CRTs is known to depend on the number and size of clusters, as well as the amount of correlation within each cluster. However, within-cluster correlation structure is often measured by a single number and clusters are usually assumed to be independent of one another. Unfortunately, these assumptions can produce misleading estimates of power.

To investigate this problem, we studied the effects of complex within-cluster structure, a measure of between-cluster mixing strength, and infectivity on power by simulating a matched-pairs CRT for an infectious process. We simulated a collection of cluster pairs as a network, controlling the proportion of edges shared across each pair. We then simulated an *SI* infectious process on each cluster pair, with one cluster assigned to treatment and the other assigned to control. The effect of treatment in this simulation lowered the probability that an infected individual succeeds at infecting a susceptible neighbor. We also considered two types of infectivity: unit and degree.

We found that between-cluster mixing had a profound effect on statistical power, no matter what network or infectious process was simulated. As the number of edges shared across clusters in different treatment groups increased to 1/2, on average the two clusters were nearly indistinguishable, and thus power fell to nearly zero. This is not surprising, but most power calculations assume clusters are independent, and this issue is usually left unaddressed. We compared these findings to the ICC approach, and found it will significantly overestimate expected power if the extent of between-cluster mixing is moderate to severe.

The effect of within-cluster structure was more nuanced. For degree infectivity, the spread of infection was less predictable if the network contained some highly-connected nodes, due to the variation in and strong effects of these hubs becoming infected. We did not observe this level of variability for networks without highly-connected hub nodes. We also did not observe this level of variability for unit infectivity, regardless of how many hubs were present in the network. Taken together, we found that for the network structures we studied, within-cluster structure had a significant impact on power only when the infectious process exhibited degree infectivity. The effect of within-cluster structure and between-cluster mixing on statistical power are qualitatively similar for a range of cluster sizes and numbers, although (as is well known) an increase in either results in more power overall.

Our simulation framework, outlined in the pseudo-algorithm in Methods, can be used to estimate power before an actual trial. If partial or full network information is available, it can be used to simulate an infectious processes using a compartmental model, and analyze the resulting outcomes as we have described. We demonstrated how to estimate between-cluster mixing using a dataset composed of cell phone calls from a large mobile carrier, which are taken to represent a contact network. For a hypothetical prospective trial on the individuals in this dataset, we defined a cluster as a group of individuals within a collection of numerically contiguous zip codes. We then grouped clusters into pairs, randomly assigned one cluster in each pair to a hypothetical treatment condition and the other to a control, and estimated mixing parameter γ for each simulation. We found substantial between-cluster mixing for all choices of cluster numbers, and mixing increased when clusters were chosen to be more numerous but smaller. Estimates of between-cluster mixing ranged from moderate to severe, regardless of whether the estimation adjusted for the frequency of calls or not.

We have shown that our simulation-based method of calculating power can differ quantitatively from the formula-based method (see Fig. 3). The two differ qualitatively as well. Traditional formula-based power calculations have been developed outside the context of network theory and consequently they do not take either within-cluster structure or between-cluster mixing into account. Furthermore, although we selected a restrictively simple simulation for clarity of demonstration, simulations for an actual prospective trial could include a much higher level of study-specific realistic detail, making a simulation-based power calculation more appropriate to the given study. The methods that we propose are most appropriate for studies in which the outcome is infectious, spreading through the population via person-to-person contacts. We leave it to subject matter experts to recognize when this condition is satisfied.

Our study invites several investigations and extensions. First, we have employed restrictively simple network models and infectious spreading process, and more nuanced generalizations are available. While our work shows how infectious spreading and complex structure can affect expected results in CRTs, more specific circumstances require extensions with more tailored network designs and infection types for power to be properly estimated. Second, we have focused our attention on matched-pair CRTs, and our framework should be extended to other CRT designs used in practice⁷. Third, these findings should be replicated in data for which both network structure and infectious spread are available.

Methods

Networks. Infectious disease dynamics have been studied extensively using deterministic ordinary differential equations²⁸ as well as network simulations²⁹. Using networks to simulate the spread of infection allows rich epidemic detail, and this added complexity facilitates exploration of the effect of cluster structure on power in CRTs. A brief treatment of these features using differential equations is in the supplement (S1).

A simple *network* \mathcal{G} consists of a set of n nodes (individuals) and a set of binary pairwise edges (relationships) between the nodes. This structure can be compactly expressed by a symmetric *adjacency matrix* $\mathbf{A}_{n \times n}$. If an edge exists between individuals i and j then $A_{ij} = A_{ji} = 1$ and 0 otherwise. The *degree* of node i , denoted by k_i , is the number of edges connecting node i to other nodes in the network. Networks can be used to describe complex systems like social communities, the structure of metabolic pathways, and the World Wide Web; many reviews of this work are available^{30–33}.

A *random graph ensemble* is a collection of all possible networks specified either by a probability model or a mechanistic model³¹. The simplest and most studied random network is the Erdős-Rényi (ER) model¹⁷, which assumes that each potential edge between any pair of nodes in a network occurs independently with fixed probability. Nodes in an ER network tend to have degrees close to their shared expected value, while in real-world social and contact networks, the distribution of node degrees is typically *heavy-tailed*: a few nodes are very highly connected (“hubs”), but most have small degree. To capture degree heterogeneity, we also simulate networks from the Barabási-Albert (BA) model^{18,34}. These networks are generated beginning with a small group of connected nodes and successively adding nodes one at a time, connecting them to the nodes in the existing network with probability proportional to the degree of each existing node. This mechanism has been shown to yield a power-law degree distribution¹⁸: $P(k) \sim k^{-\alpha}$ with $\alpha = 3$. This distribution is heavy-tailed, so the probability that some individuals are highly connected is more likely than in other network models like the ER. While it can be difficult to assess whether an observed network has a power-law degree distribution³⁵, the BA model comes closer to capturing the heavy-tailed degree distributions observed in social networks than the ER model. Another hallmark of real-world social networks is that individuals tend to cluster together into *communities*, or groups of individuals who share more edges with each other than between them³⁶ <http://www.sciencedirect.com/science/article/pii/S0370157309002841>. We use *stochastic blockmodels* (SBMs)¹⁹ to model within-cluster communities by assuming that each node is a member of a one block in a partition of blocks \mathcal{B} comprising all nodes in the network, and that the probability of an edge between two nodes depends only on block membership (see supplementary material S3 for additional details). Other popular families of random networks include Exponential Random Graphs (ERGMs)³⁷ and Small-World network of Watts and Strogatz, among others³⁸. We leave their implications for CRTs for future research. Network instances generated using Python’s networkx library. Each node within each cluster has the same expected number of edges $\langle k \rangle = 4$. For Figs 2 and 3, we chose $C = 20$ and $n = 300$, because for $\gamma = 0$ these parameters yield empirical power within 0.8–0.9, which is a typical range used in cluster randomized trials.

Network mixing. In each cluster pair, one cluster is randomly assigned to treatment. The mixing parameter γ can be expressed in terms of the entries in the adjacency matrix, \mathbf{A} , and the treatment assignment of clusters:

$$\gamma := \frac{\sum_{ij} A_{ij} (1 - \delta(r_i, r_j))}{\sum_{ij} A_{ij}} \quad (2)$$

$$= 1 - \frac{1}{2m} \sum_{ij} A_{ij} \delta(r_i, r_j). \quad (3)$$

Here, $m := \sum_{i < j} A_{ij}$ is the total number of edges in the study, $r_i = 1$ if node i is in the treatment arm and $r_i = 0$ otherwise, and $\delta(a, b)$ is equal to 1 when $a = b$ and 0 otherwise. This definition of between-cluster mixing is closely related to the concept of *modularity*, used extensively in network community detection (see supplementary material S2). If $\gamma = 0$, the two clusters share no edges with each other. If $\gamma = 1/2$, there are as many edges reaching across two clusters as exist within them. Finally, if $\gamma = 1$, edges are only found between clusters, and the cluster pair network is said to be *bipartite*. A schematic of network mixing is shown in Fig. 5.

Network rewiring. We first simulate two random networks from the same network model and with the same number of nodes and edges, each corresponding to a cluster in a pair of clusters. Then, we randomly select one edge from each cluster in the pair and remove these two edges. Finally we create two new edges among the four nodes such that the two edges reach across the cluster pair. This process is called *degree-preserving rewiring*²⁰ because it preserves the degrees of all the nodes involved. The process is depicted in Fig. 6. We repeat the rewiring process until proportion γ of the total edges are rewired. The result is a single cluster pair in our simulated CRT, and the pair-generating process is repeated until we have generated our target number of cluster pairs.

Infectious spread. Compartmental models assume that each node in a population is in one of a few possible states, or compartments, and that individuals switch between these compartments according to some rules. Although more realistic models include more states³⁹, we will assume for simplicity that nodes are in only one of two states: uninfected but susceptible (S), and infected and contagious (I). We assume that the network structure of each cluster pair represents the possible transmission paths from infected nodes to susceptible ones.

Let I_{irct} represent the infectious status for node i in treatment arm $r = \{0, 1\}$ and cluster pair $c = 1, \dots, C$ at discrete time $t = 1, \dots, T_c$, with $I_{irct} = 1$ if the node is infected and 0 otherwise. We define $r = 0$ if node i is in the control arm, and $r = 1$ if i is in the treatment arm. Let $I_{rct} := \langle I_{irct} \rangle$ represent the proportion of infected nodes in cluster pair c at discrete time t . At the beginning of the study, 1% of individuals

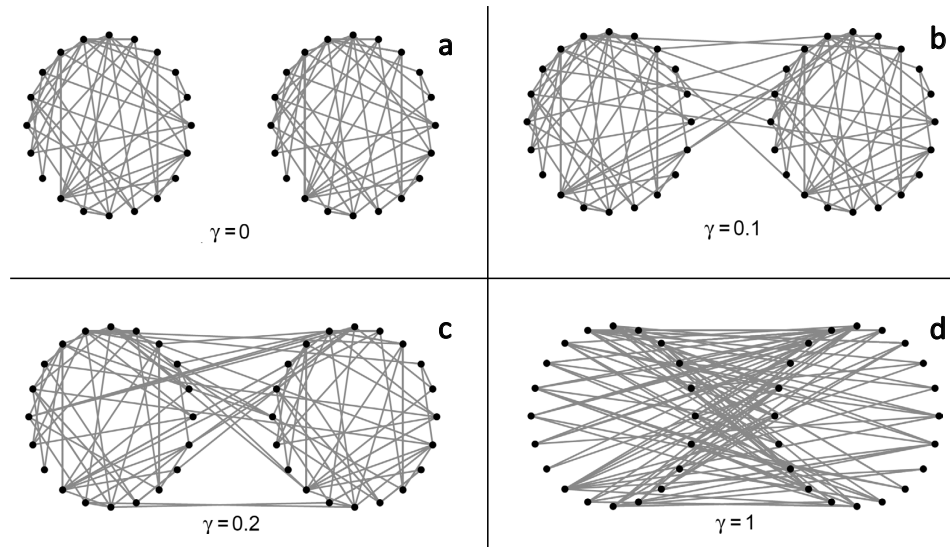


Figure 5. A diagram showing two clusters with various proportions of mixing.

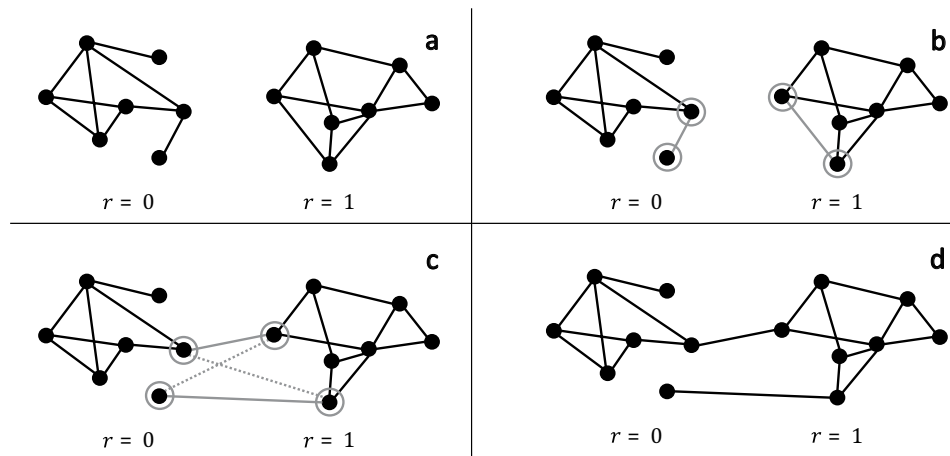


Figure 6. Degree-preserving rewiring is performed by selecting an edge within each cluster, and swapping them to reach across the cluster pair. The dashed gray lines represent another way the edges could have been rewired while still preserving degree; either rewiring is chosen with equal probability.

selected at random in each cluster is infected, i.e. $I_{rc0} = 0.01$. For each time step t , each node i selects q_i network neighbors at random, and infects each one with probability p_i . Because different infectious diseases have different infectivity behavior, we study both unit and degree infectivity, or $q_i = 1$ and $q_i = k_i$, respectively. We assume that the infection probability depends only on the treatment arm membership of each node r_i , thus $p_i = p_{r_i}$. Treatment reduces the probability p_{r_i} of infection. If two clusters in a pair have the same infection rate, the treatment has no effect and $p_{r_i} = p_0$. This is the *null hypothesis* under examination in our hypothetical study. When we simulate trials under the null hypothesis we set $p_0 = 0.30$ in every cluster. The *alternative hypothesis* holds if the treatment succeeds in reducing the infection rate, $p_1 < p_0$. When we simulate under the alternative hypothesis, $p_0 = 0.30$ and $p_1 = 0.25$. The trial ends when the cumulative incidence of infection grows to 10% of the population, i.e., when the cluster pair infection rate $\langle I_{ircT_c} \rangle = 0.1$ for some time T_c .

Analysis. At the end of the simulation, we test whether the treatment was effective by comparing the number of infections between treated and control clusters according to two analysis scenarios. In real-world CRTs, the most efficient and robust way to compare the two groups depends on what information about the infection can feasibly be gathered from the trial. In some trials, surveying the infectious

status of individuals is difficult, and therefore this information is only available for the beginning and end time points of the trial. In others, the times to infection for each node are available. In addition to what information is available, the researcher must choose a statistical test according to which assumptions they find suitable to their study. A *model-based test* assumes that the data are generated according to a particular model, which can be more powerful than other tests if the model is true⁴⁰. Alternatively, a *permutation test*⁴¹ does not make any assumptions about how the data were generated. To show how to conduct an analysis suited to different scenarios based on available data, we analyzed our simulated trial using two different sets of assumptions. In Scenario 1, we assume that outcomes are only known at the end of the trial, and perform a model-based test. In Scenario 2, we assume that the time to each infection is known, and perform a permutation test. We show that the results of the simulation are qualitatively similar under both scenarios. (Note that it is possible to use a permutation test for Scenario 1 or a model-based test for Scenario 2, which would create two new analyses.) For both scenarios, a description of how to carry out a simulation-based power calculation for a CRT studying an infectious spread through networks is as follows:

Scenario 1: The *log risk ratio* is the logarithmic ratio of infected individuals in the treatment clusters to the control clusters at the end of study. For simulation m , let $I_m^{(0)} = \left\langle \log \frac{I_{0cT_c}}{I_{1cT_c}} \right\rangle = \langle \log I_{0cT_c} - \log I_{1cT_c} \rangle$ be the difference in the number of infections between two clusters in a pair averaged over each of the C cluster pairs at the trial end T_c . The simulation was repeated 20,000 times under the null hypothesis and cutoff values $I_{2.5}^*$ and $I_{97.5}^*$ were established such that $P(I_{2.5}^* < I_m^{(0)} < I_{97.5}^*) = \alpha$ for significance level $\alpha = 0.05$. We repeated this process under the alternative 20,000 times, and the proportion of these trials with statistics $I_m^{(1)}$ more extreme than $(I_{2.5}^*, I_{97.5}^*)$ is the simulated power or empirical power.

Scenario 2: We pool the individual infection times for the treatment arm and the control arm, and summarize the difference between the two arms' infection times using an appropriate statistic (e.g. the logrank statistic⁴²). The permutation test is performed by comparing the observed logrank statistic to the distribution of log-rank statistics when the treatment labels are permuted, or switched, for each cluster pair. The p -value for this analysis is the proportion of times the log-rank statistic with the observed labels is more extreme than the permuted log-rank statistics. Because the permutation test is computationally expensive, this entire process is repeated 2,000 times, and we calculate the proportion of permutation p -values below 0.05, which is the empirical or simulated power.

We also compare this formulation to traditional methods. From Hayes and Bennett¹⁶, the number of clusters required for power β in a CRT with binary outcomes is:

$$C = 2 + (z_{\alpha/2} + z_{\beta})^2 \times \left[\frac{\pi_0(1 - \pi_0)}{n} + \frac{\pi_1(1 - \pi_1)}{n} + k^2(\pi_0^2 + \pi_1^2) \right] / (\pi_0 - \pi_1)^2 \quad (4)$$

To calculate power, we fix $n = 300$, $C = 20$, and $\alpha = 0.05$, and solve for power β . In this formula, π_0 and π_1 are the mean proportion of outcomes within control and treated clusters, and k is the coefficient of variation, which is directly related to the ICC ρ ^{6,43}:

$$k = \sqrt{\rho \times \frac{1 - \hat{\pi}}{\hat{\pi}}} \quad (5)$$

where $\hat{\pi}$ is the overall prevalence by study end. This calculation assumes that the log risk ratio by study end $\log\left(\frac{\pi_0}{\pi_1}\right)$ takes on the values observed in our simulation setting 0.135 for no between-cluster mixing $\gamma = 0$, and the overall prevalence is 10%, both assumed to be accurately estimated from a small pilot study. The value for the ICC must also be assumed beforehand or estimated in a small pilot study. To compare this approach with our simulation design, we assumed that the ICC took on a range of plausible empirical values 0.0–0.1 reported in the literature^{7,43,44}. For more details, see supplementary material S4.

Application. For the calling dataset, we consider two definitions for an edge A_{ij} between individuals i and j , belonging to clusters c_i and c_j respectively. The number of calls between i and j over the period of investigation is defined as d_{ij} . For the unweighted case, we assume an edge exists between the two individuals if they have called each other at least once, $A_{ij} = \mathbb{I}(d_{ij} \geq 1)$, and otherwise no edge exists between them $A_{ij} = 0$. For the weighted case, we assume an edge between them may be weighted by the number of total calls made between them, $A_{ij} = d_{ij}$. Using both definitions, we found the degree distribution of each cell phone to be heavy-tailed (see supplementary material S5).

References

1. Lipsey, M. W. *Design sensitivity: Statistical power for experimental research* vol. 19 (Sage, 1990).
2. Button, K. S. *et al.* Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience* **14**, 365–376 (2013).
3. Ioannidis, J. P. Why most published research findings are false. *PLoS medicine* **2**, e124 (2005).
4. Murray, D. M. *Design and analysis of group-randomized trials* vol. 29 (Oxford University Press, 1998).

5. Donner, A. & Klar, N. *Design and analysis of cluster randomization trials in health research* vol. 220 (London Arnold Publishers, 2000).
6. Hayes, R. & Moulton, L. *Cluster randomised trials* (Chapman & Hall/CRC, 2009).
7. Eldridge, S. & Kerry, S. *A practical guide to cluster randomised trials in health services research* vol. 120 (John Wiley & Sons, 2012).
8. Crespi, C. M., Wong, W. K. & Mishra, S. I. Using second-order generalized estimating equations to model heterogeneous intraclass correlation in cluster-randomized trials. *Statistics in medicine* **28**, 814–827 (2009).
9. Wang, R., Goyal, R., Lei, Q., Essex, M. & De Gruttola, V. Sample size considerations in the design of cluster randomized trials of combination hiv prevention. *Clinical Trials* 1740774514523351 (2014).
10. Gilbody, S., Bower, P., Torgerson, D. & Richards, D. Cluster randomized trials produced similar results to individually randomized trials in a meta-analysis of enhanced care for depression. *Journal of clinical epidemiology* **61**, 160–168 (2008).
11. Sobel, M. E. What do randomized studies of housing mobility demonstrate? causal inference in the face of interference. *Journal of the American Statistical Association* **101**, 1398–1407 (2006).
12. Hemming, K., Girling, A. J., Sitch, A. J., Marsh, J. & Lilford, R. J. Sample size calculations for cluster randomised controlled trials with a fixed number of clusters. *BMC Medical Research Methodology* **11**, 1471–2288 (2011).
13. Ukoumunne, O. C. *et al.* Methods in health service research: evaluation of health interventions at area and organisation level. *BMJ: British Medical Journal* **319**, 376 (1999).
14. Hayes, R., Alexander, N. D., Bennett, S. & Cousens, S. Design and analysis issues in cluster-randomized trials of interventions against infectious diseases. *Statistical Methods in Medical Research* **9**, 95–116 (2000).
15. Zhou, T., Liu, J.-G., Bai, W.-J., Chen, G. & Wang, B.-H. Behaviors of susceptible-infected epidemics on scale-free networks with identical infectivity. *Physical Review E* **74**, 056109 (2006).
16. Hayes, R. & Bennett, S. Simple sample size calculation for cluster-randomized trials. *International journal of epidemiology* **28**, 319–326 (1999).
17. Erdős, P. & Rényi, A. The evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.* **5**, 17–61 (1960).
18. Barabási, A.-L. & Albert, R. Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
19. Anderson, C. J., Wasserman, S. & Faust, K. Building stochastic blockmodels. *Social Networks* **14**, 137–161 (1992). Special Issue on Blockmodels.
20. Newman, M. E. J. Mixing patterns in networks. *Phys. Rev. E* **67**, 026126 (2003).
21. Anderson, R. M. & May, R. M. *Infectious Diseases of Humans* (Oxford University Press, Oxford, UK, 1991).
22. HELLERINGER, S. & KOHLER, H.-P. Sexual network structure and the spread of hiv in africa: evidence from likoma island, malawi. *Aids* **21**, 2323–2332 (2007).
23. Goyal, R., Blitzstein, J. & De Gruttola, V. Sampling networks from their posterior predictive distribution. *Network Science* **2**, 107–131 (2014).
24. Molloy, M. & Reed, B. A critical point for random graphs with a given degree sequence. *Random structures & algorithms* **6**, 161–180 (1995).
25. Hidalgo, C. A. & Rodriguez-Sickert, C. The dynamics of a mobile phone network. *Physica A: Statistical Mechanics and its Applications* **387**, 3017–3024 (2008).
26. Eagle, N., Pentland, A. S. & Lazer, D. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences* **106**, 15274–15278 (2009).
27. Onnela, J.-P., Arbesman, S., González, M. C., Barabási, A.-L. & Christakis, N. A. Geographic constraints on social network groups. *PLoS one* **6**, e16939 (2011).
28. Keeling, M. J. & Rohani, P. *Modeling infectious diseases in humans and animals* (Princeton University Press, 2008).
29. Pastor-Satorras, R., Castellano, C., Van Mieghem, P. & Vespignani, A. Epidemic processes in complex networks. *Rev. Mod. Phys.* **87**, 925–979 (2015).
30. Newman, M. E. J. The structure and function of complex networks. *SIAM Review* **45**, 167–256 (2003).
31. Newman, M. E. J. *Networks: An Introduction* (Oxford University Press, Inc., New York, NY, USA, 2010).
32. Kolaczyk, E. D. *Statistical analysis of network data: methods and models* (Springer Science & Business Media, 2009).
33. Mieghem, P. V. *Performance Analysis of Complex Networks and Systems* (Cambridge University Press, 2014).
34. de Solla Price, D. J. Networks of scientific papers. *Science* **149**, 510–515 (1965).
35. Clauset, A., Shalizi, C. R. & Newman, M. E. J. Power-law distributions in empirical data. *SIAM Review* **51**, 661–703 (February 2009).
36. Porter, M. A., Onnela, J.-P. & Mucha, P. J. Communities in networks. *Notices of the AMS* **56**, 1082–1097 (2009).
37. Frank, O. & Strauss, D. Markov graphs. *Journal of the American Statistical Association* **81**, 832–842 (1986).
38. Watts, D. J. & Strogatz, S. H. Collective dynamics of ‘small-world’ networks. *Nature* **393**, 440–442 (1998).
39. van Kleef, E., Robotham, J. V., Jit, M., Deeny, S. R. & Edmunds, W. J. Modelling the transmission of healthcare associated infections: a systematic review. *BMC infectious diseases* **13**, 294 (2013).
40. Murray, D. M., Varnell, S. P. & Blitstein, J. L. Design and analysis of group-randomized trials: a review of recent methodological developments. *American Journal of Public Health* **94**, 423–432 (2004).
41. Good, P. I. *Resampling methods* (Springer, 2001).
42. Harrington, D. Linear rank tests in survival analysis. *Encyclopedia of biostatistics* (2005).
43. Pagel, C. *et al.* Intracluster correlation coefficients and coefficients of variation for perinatal outcomes from five cluster-randomised controlled trials in low and middle-income countries: results and methodological implications. *Trials* **12**, 151 (2011).
44. Turner, R. M., Thompson, S. G. & Spiegelhalter, D. J. Prior distributions for the intracluster correlation coefficient, based on multiple previous estimates, and their application in cluster randomized trials. *Clinical Trials* **2**, 108–118 (2005).

Acknowledgements

The authors would like to acknowledge Victor De Gruttola and Paige Williams for their helpful suggestions. We thank Telenor Research and Senior Data Scientist Kenth Engø-Monsen for supplying the telecom dataset, and helping with initial processing and answering questions related to the dataset. We also thank the National Institutes of Health for funding support for PCS (NIH-5T32ES007142-32, PI Coull). JPO is supported by the Harvard T.H. Chan School of Public Health McLennan Dean’s Challenge Grant, NIAID R37 51164 (De Gruttola), and UG01GH000447 (Essex).

Author Contributions

PCS conducted the simulations and analysis under supervision of ELO and JPO. All authors conceived the work and reviewed the manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Staples, P. C. *et al.* Incorporating Contact Network Structure in Cluster Randomized Trials. *Sci. Rep.* **5**, 17581; doi: 10.1038/srep17581 (2015).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>