Harvard University
Harvard University Biostatistics Working Paper Series

Year 2014 *Paper* 180

Estimation of the Overall Treatment Effect in the Presence of Interference in Cluster-randomized Trials of Infectious Disease Prevention

Nicole Bohme Carnegie* Rui Wang†
Victor De Gruttola‡

*University of Wisconsin - Milwaukee, carnegin@uwm.edu
†Brigham and Women's Hospital, ruiwang29@gmail.com
‡Harvard School of Public Health, degrut@hsph.harvard.edu

http://biostats.bepress.com/harvardbiostat/paper180

Estimation of the overall treatment effect in the presence of interference in cluster-randomized trials of infectious disease prevention.

Nicole Bohme Carnegie[1], Rui Wang[2,3], and Victor De Gruttola[3]

[1]Zilber School of Public Health, University of Wisconsin-Milwaukee
[2]Division of Sleep Medicine, Brigham and Women's Hospital
[3]Department of Biostatistics, Harvard School of Public Health

September 18, 2014

Abstract

Causal inference has progressed dramatically in the past few decades, but one critical assumption that remains difficult to relax is the assumption of no interference between units, one component of the single unit treatment value assumption (SUTVA). Interference occurs when the treatment of one unit can affect the outcome of another, making it difficult to justify the assumption of no interference for outcomes that depend on social interactions, such as occurrence of infectious disease. Significant progress has been made in dealing with this issue in design of vaccine trials that incorporate two-stage randomization of clusters and of individuals within them. Nonetheless there is considerable need for development of methods to relax the no-interference assumption in many settings, such as the one we describe.

In cluster-randomized trials, SUTVA is often thought to be justified by design of studies with large distances (geographic or otherwise) between units, but this condition may not be feasible or even sufficient in some settings. In trials of infectious disease prevention, the estimate of the treatment effect will be attenuated (relative to that which would occur if the treatment were implemented population-wide) if a fraction of the exposures in the treatment clusters come from individuals who are outside these clusters.

This source of interference – the fraction of contacts sufficient for transmission that are within the cluster – is, however, potentially measurable. In this manuscript, we make use of the rich history of epidemic modeling to infer the way in which a given level of interference affects the force of infection upon members of a cluster. This model allows us to develop a weighting factor on the treatment indicator that recovers a counterfactual treatment effect in the absence of interference in a proportional hazards or frailty model.

1 Introduction

Causal inference has progressed dramatically in the past few decades, but one critical assumption that remains difficult to relax is the assumption of no interference between units, one component of the single unit treatment value assumption (SUTVA) in Rubin's potential outcomes framework for causal inference. Interference arises when the treatment status of one unit can influence the outcome of another. This issue arises in areas as diverse as social interventions [1], education [2], or vaccine testing [3], among others.

When interference is present, we can no longer assume that the outcomes of the control units are representative of the outcomes of the treated units had they not received treatment, and the observed treatment effect is now a function of the direct effect of treatment as well as indirect effects from spillover or contamination. This paper deals specifically with cluster-randomized trials of an intervention to prevent the transmission of an infectious disease. A reduction in incidence could occur via an individual-level *direct effect* of the intervention, or through the reduction in the force of infection in the cluster by reducing the overall burden. This *indirect effect* will be attenuated if cluster members have contacts outside of the cluster.

In many cluster-randomized trials, ensuring that clusters are sufficiently distant from each other by some metric (usually geographic), can help support the assumption of no interference (see, e.g. [4–6]). This design strategy may not be sufficient in all cases. In trials of infectious disease prevention, the estimate of the randomized treatment effect will be attenuated relative to that which would occur if the treatment were implemented population-wide if a fraction of the exposures in the treatment clusters come from outside the cluster, regardless of the distances between clusters [7, 8].

Indirect effects at the individual level are identifiable using two-stage randomization [3, 9]. For example, randomizing first clusters to treatment or control, and then individuals within treatment clusters to treatment or control allows for comparison of the outcomes of non-treated individuals in non-treated clusters to those of non-treated individuals in treated clusters to understand effects of mixing across clusters. When the desired outcome of the trial is at the cluster level, however, such a design would involve creating super-clusters to be randomized to treatment or control, then randomizing clusters within treated super-clusters to either receive treatment or control. The required resources to perform such a design make it generally infeasible.

Given this concern, other approaches are required to understand indirect effects between clusters. In this paper, we leverage epidemic models to account for one form of interference: cross-cluster exposures due to mixing across clusters. This leads to spillover effects in control clusters, which consequently receive less infectious pressure from treatment clusters than they would in the absence of an intervention. Conversely, treatment clusters receive infectious pressure from control clusters that are unaffected by treatment. Assuming that the prevention intervention is effective, we expect that it would be rolled out population-wide, so an estimand of interest is the treatment effect in the absence of this mixing equivalent to the a setting where the infectious pressure in the treatment cluster comes solely from treated clusters. This is termed the *overall treatment effect*. In these settings, it is possible to measure the source of the interference –

the proportion of contacts sufficient to transmit the disease that occur outside the cluster. We leverage mathematical models of epidemics to understand how a given level of mixing between clusters is likely to affect the transmission dynamics in the presence of an intervention.

We are particularly motivated by a large community-randomized trial of a combination HIV prevention intervention in Botswana, the Botswana Combination Prevention Project (BCPP) [8, 10]. In this study, the intervention comprises efforts to increase testing coverage (and thereby treatment coverage for those who qualify), treatment for subjects with high viral load regardless of CD4 count (the usual metric used to determine eligibility), prevention of mother-to-child transmission, and promotion of male circumcision. Of these, only the last will have a protective effect for the individual receiving the treatment; the rest operate by reducing the viral exposure of susceptibles by treating those who are already infected. Thus, the bulk of the effect will be due to indirect effects. We also expect a significant fraction of sexual relationships to be with members of other communities, thus attenuating the estimate of the overall treatment effect.

Section 2 outlines a simple epidemic model for disease transmission within and between clusters. Section 3 relates this epidemic model to a proportional hazards model for survival and derive a weighting factor for treatment that allows us to recover the overall treatment effect. We discuss adjustment for mismeasurement of the weighting factor in Section 5. The effectiveness of this adjustment is shown via simulation in Section 4, and an application to design studies for the BCPP appears in Section 6.

2 A simple epidemiological model for disease transmission between clusters

Suppose that we have clusters $1, \cdots, C$ that form a closed population, with $1, ..., c_1$ receiving the control and the rest treatment. We begin by proposing a simple SEIR model of disease transmission between clusters, defined by the following system of equations:

$$\begin{aligned}
\frac{dS_i}{dt} &= \mu - \sum_{j=1}^{C} \alpha_{ij} S_i I_j - \mu S_i \\
\frac{dE_i}{dt} &= \sum_{j=1}^{C} \alpha_{ij} S_i I_j - \upsilon E_i - \mu E_i \\
\frac{dI_i}{dt} &= \upsilon E_i - \gamma I_i - \mu I_i \\
\frac{dR_i}{dt} &= \gamma I_i - \mu R_i
\end{aligned}$$

Here S_i, E_i, I_i and R_i are the proportions of the entire population that are susceptible, exposed, infectious, and recovered individuals in cluster i with parameters:

- μ, the birth and death rate (assuming no population growth)
- α_{ij}, the rate of new infections among susceptibles in i from infecteds in j
- υ, the rate of transition from the exposed to infectious state
- γ, the rate of recovery for infected individuals

If there were no cross-cluster mixing, $\alpha_{ij} = 0 \ \ \forall i \neq j$, and the remaining α_{ii} would be equivalent to the total force of infection in their respective clusters. We break α_{ij} down into two pieces: $\alpha_{ij} = m_{ij}\eta_j$, where m_{ij} is the percentage of partnerships of susceptibles in cluster i that are with infecteds in cluster j and η_j is the force of infection in cluster j (per-contact transmission probability times the rate of contacts of infecteds in i). We assume no differential activity by infection status. If we assume, in addition, no differential mixing by infection status the m_{ij} can be estimated from the full population. Note that, since this is a closed population, $m_{ii} = 1 - \sum_{j \neq i} m_{ij}$.

Let the counts of individuals in the infected states, I_i, be noted by Y_i. This gives us an incidence rate in cluster i of $\delta_i = \sum_{j=1}^{C} \alpha_{ij} Y_j / N_j = \sum_{j=1}^{C} m_{ij}\eta_j Y_j / N_j = \sum_{j=1}^{C} \nu_{ij}(t)\eta_j$. Note that this is a function of time, as Y_i changes as the epidemic progresses. If we can assume that prevalence within clusters (Y_i/N_i) is roughly constant over time, we can ignore the dependence of ν on t. In what follows, we make this assumption, and note the prevalence in cluster i as π_i.

The probability that a susceptible in cluster i is infected in the interval $[a, b]$ is $\Pr(T_i \in [a, b]) = \int_a^b \delta_i \cdot \exp(-\delta_i t)\, dt = \exp(\delta_i a) - \exp(\delta_i b)$.

While this model does not perfectly describe most epidemic processes, it often serves as a reasonable approximation. Although this model is presented to describe our framework, the following results hold for other varieties of epidemic models, such as an SIR model, SI model, or a staged model where infectivity is constant across stages.

3 Proportional hazards and the epidemic framework

3.1 The proportional hazards model in the absence of mixing

The proportional hazards model for survival data states that the instantaneous rate of failure, or hazard, takes the form $\lambda(t, X) = \lambda_0(t)\exp(\beta X)$, where X is a set of covariates that explain variability in hazard

rates. This form is a natural parallel to the probability of infection taken from the epidemic model in Section 2, when the covariates X consist of only a treatment indicator, Z. Under the assumed epidemic process, with $m_{ij} = 0 \ \ \forall i \neq j$, we derive the proportional hazards estimate of the effect of treatment, Z:

$$\begin{aligned}
\Pr\left(T_i > t | Z = 0\right) &= \exp\left(-t \frac{1}{c_1} \sum_{i=1}^{c_1} \delta_i\right) \\
\Pr\left(T_i > t | Z = 1\right) &= \exp\left(-t \frac{1}{C - c_1} \sum_{i=c_1+1}^{C} \delta_i\right) \\
\Pr\left(T_i > t | Z\right) &= \Pr\left(T_i > t | Z = 0\right) \cdot \exp\left(-t \left(\frac{1}{C - c_1} \sum_{i=c_1+1}^{C} \delta_i - \frac{1}{c_1} \sum_{j=1}^{c_1} \delta_j\right) Z\right) \\
&= \Pr\left(T_i > t | Z = 0\right) \cdot \exp\left(-t \left(\frac{1}{C - c_1} \sum_{i=c_1+1}^{C} \eta_i \pi_i - \frac{1}{c_1} \sum_{j=1}^{c_1} \eta_j \pi_j\right) Z\right) \qquad (3.1)
\end{aligned}$$

This means that we would like to recover $\frac{1}{C-c_1} \sum_{i=c_1+1}^{C} \eta_i \pi_i - \frac{1}{c_1} \sum_{j=1}^{c_1} \eta_j \pi_j$ to estimate the counterfactual treatment effect in the absence of interference when mixing is present.

Inclusion of other covariates that inform the hazard rate is also consistent with the epidemic model framework; we focus on the treatment alone for clarity of presentation. In the presence of other covariates, the compartments or sub-populations of the epidemic model would be further subdivided by covariates. For example, one could divide the groups presented in Section 2 (S_i, E_i, I_i and R_i) by sex s, yielding S_{is}, E_{is}, I_{is} and R_{is}. The summation then would be over community as well as covariates; the m_{ij} could either be assumed to be constant over covariates or could be estimated within subgroups.

3.2 The proportional hazards model in the presence of mixing

In this section, we build upon the result from Section 3.1 to show what the coefficient on the treatment indicator in the proportional hazards model represents when the m_{ij} are not zero between clusters. In this case,

$$\begin{aligned}
\Pr\left(T_i > t | Z = 0\right) &= \exp\left(-t \left(\frac{1}{c_1} \sum_{i=1}^{c_1} \delta_i\right)\right) \\
\Pr\left(T_i > t | Z = 1\right) &= \exp\left(-t \left(\frac{1}{C - c_1} \sum_{i=c_1+1}^{C} \delta_i\right)\right)
\end{aligned}$$

Thus, we can rewrite the model in terms of a baseline hazard when treatment is absent and a proportional adjustment when treatment is applied:

$$\Pr(T_i > t|Z) = \Pr(T_i > t|Z=0) \cdot \exp\left(-tZ\left(\frac{1}{C-c_1}\sum_{i=c_1+1}^{C}\sum_{j=1}^{C}\nu_{ij}\eta_j - \frac{1}{c_1}\sum_{i=1}^{c_1}\sum_{j=1}^{C}\nu_{ij}\eta_j\right)\right)$$

This can be shown to be equivalent to

$$\begin{aligned}\Pr(T_i > t|Z) = \Pr(T_i > t|Z=0) \cdot \exp\Bigg(-tZ\Bigg(&\frac{1}{C-c_1}\sum_{j=c_1+1}^{C}\eta_j\pi_j(m_j^{in}-m_j^{out}) \\ &-\frac{1}{c_1}\sum_{j=1}^{c_1}\eta_j\pi_j(m_j^{in}-m_j^{out}) + \sum_{j=1}^{C}\frac{2c_1-C}{(C-c_1)c_1}\eta_j\pi_j m_j^{out}\Bigg)\Bigg)\end{aligned} \tag{3.2}$$

where m_j^{in} and m_j^{out} are the total mixing with cluster j of clusters of the same and opposite treatment status, respectively. Note that if the number of treatment and control groups are equal, the last term is zero; for nearly balanced group sizes it will be very small. This suggests that weighting the treatment variable Z by the function $(m_j^{in} - m_j^{out})$ of cross-cluster mixing rates will allow us to estimate the overall treatment effect, $\sum_{i=c_1+1}^{C}\eta_i\pi_i - \sum_{j=1}^{c_1}\eta_j\pi_j$.

The estimator of the overall treatment effect can thus be expressed as a scaling of the randomized treatment effect estimator. This method will alter point estimates of the treatment effect, but not the results of tests of null hypotheses of no treatment effect. Standard errors scale with the point estimates, so a confidence interval that includes 0 for the randomized effect will also include 0 when adjusted for interference. In the case where $(m_j^{in} - m_j^{out})$ is the same across clusters, the resulting estimate will be $\frac{\hat{\tau}^*}{(m_j^{in}-m_j^{out})}$, where $\hat{\tau}^*$ is the estimate of the randomized treatment effect.

This result can be extended directly to a frailty model with cluster-specific random effects. Suppose that each cluster has a random deviation, b_i, on the force of infection, $\delta_i = \xi + b_i + \tau \cdot I(Z=1)$, where τ is the treatment effect and ξ is the average force of infection in the absence of treatment. In the absence of mixing, this gives $\sum_{i=c_1+1}^{C}\delta_i - \sum_{j=1}^{c_1}\delta_j = \tau - \left(\sum_{i=c_1+1}^{C}b_i - \sum_{j=1}^{c_1}b_j\right)$. Assuming that the b are identically distributed, the expected value of $\left(\sum_{i=c_1+1}^{C}b_i - \sum_{j=1}^{c_1}b_j\right) = 0$, implying that the overall treatment effect τ is still equal to $\frac{1}{C-c_1}\sum_{i=c_1+1}^{C}\eta_i\pi_i - \frac{1}{c_1}\sum_{j=1}^{c_1}\eta_j\pi_j$ and the randomized treatment effect τ^* is still equal to $\frac{1}{C-c_1}\sum_{i=c_1+1}^{C}\sum_{j=1}^{C}\nu_{ij}\eta_j - \frac{1}{c_1}\sum_{i=1}^{c_1}\sum_{j=1}^{C}\nu_{ij}\eta_j$.

Note that for pairs of clusters forming closed units $(m_j^{in} - m_j^{out}) = (m_{11} + m_{00} - 1)$. The model described here can be fit either to a single system of multiple clusters, or multiple sets of closed systems

(i.e., pairs of clusters).

4 Simulation study

We test this theory by simulating paired clusters that form closed populations with disease spread modeled by a stochastic SEIR model with varying population size over time. We begin with 400 individuals in cluster 1 and 600 in cluster 0; total population size varies stochastically over time as individuals enter and leave the population (birth/death rate $\mu = 1/10000$). The intensity of contacts is set so that individuals have an average of two contacts on any given day. The R_0 value for each cluster is set to be 2.5 initially, and falls to 1.25 in cluster 1 when the treatment is rolled out at the start of data collection. Mixing is established by forcing 0 to 50% of contacts of members of cluster 1 to be with members of cluster 0. This yields an asymmetry in m_{10} and m_{01}. Disease spread starts with 5% initial infecteds, and data are "collected" over a three-year period of roughly constant prevalence after the initial ramp-up of infections. We then record the rates of mixing between clusters, infection times and censoring status for individuals. This is repeated for 20 simulated pairs of clusters. For each pair, the infectivity parameter implied by R_0 is perturbed to generate random effects by cluster. We fit an unadjusted frailty model and an adjusted frailty model.

The results of the simulation are given in Figure 1: each point represents the average point estimate over 500 simulated data sets, and the 95% confidence interval is given by the vertical lines. Note the rapid attrition of the estimated treatment effect as the mixing rate increases for the unadjusted model. The adjusted estimates for all levels of mixing are well within the confidence bounds for the no-mixing case. The standard error of the adjusted estimates increases proportionately to the increase in the point estimate, as expected; this method does not add information to increase power in estimation, but rather allows us to calculate an alternate estimand to the randomized treatment effect.

5 Errors-in-variables

The weighting approach to estimating the overall treatment effect creates a new cluster-level treatment variable, Z, that is 0 for all untreated clusters, but a function of mixing rates for treated clusters. The mixing parameters are likely to be measured with error, yielding an observed variable W that is a noisy version of the true Z. In some cases, it may be possible to characterize this error; for example, estimating mixing rates as proportions of reported partnerships for which both partners reside in the same cluster from behavioral surveys leads to a natural binomial form for the variability of the estimates. As a consequence,

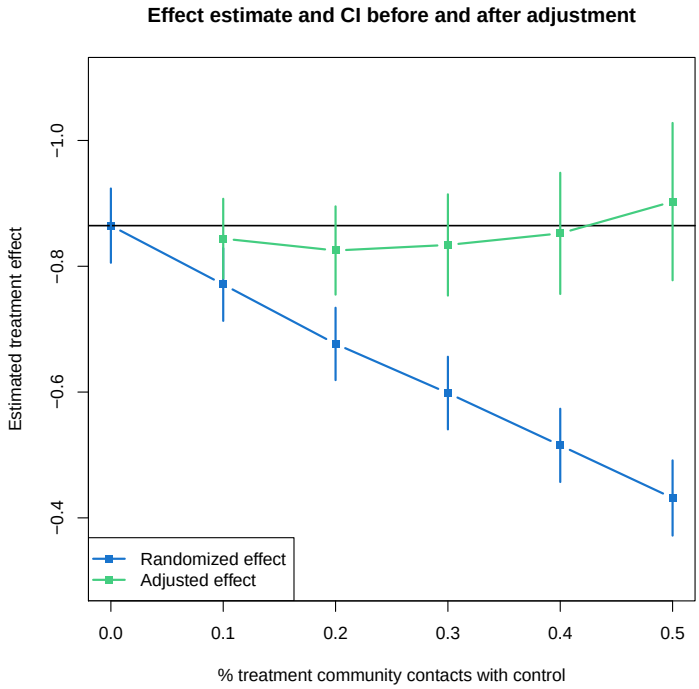

Figure 1: **Simulation results from true SEIR epidemic.** Average point estimate and 95% confidence interval for randomized and overall treatment effects over 500 simulated data sets with SEIR model. The horizontal black line indicated the expected value of the treatment effect in the absence of interference.

we recommend consideration of a correction for errors-in-variables when implementing this method.

Measurement error in proportional hazards models has been well studied [11–15], but less so in the setting of clustered survival data and frailty models. Li and Lin have proposed two approaches for measurement error in frailty models: a structural model that is fit using non-parametric maximum likelihood [16], and a functional model using simulation and extrapolation (SIMEX) [17]. The functional approach is appealing as it does not require assumptions about the distribution of the unobserved error-prone covariate and is easy to implement using standard software.

The SIMEX approach is implemented in two stages: simulation and extrapolation. We suppose that the measurement error variance σ_u^2 is known, and let W be the observed noisy covariate. In the simulation step we generate data sets with gradually increasing error relative to the observed W. For a given $\zeta > 0$, we generate C noisy versions of W by adding $U_c \sim \text{iid}N(0, \zeta\sigma_u^2)$, yielding a new noisy covariate $W_c^* = W + U_c$ with error variance $(\zeta + 1)\sigma_u^2$. We then fit a frailty model without measurement error for each W_c^* and average the resulting coefficient (or frailty variance) estimates to obtain $\hat{\beta}_\zeta$ (or $\hat{\theta}_\zeta$). In the extrapolation step, a regression model of $\hat{\beta}_\zeta$ ($\hat{\theta}_\zeta$) as a function of ζ is fit to the simulated means. Using the fitted model, we extrapolate back to $\zeta = -1$, corresponding to no measurement error, to obtain $\hat{\beta}$ ($\hat{\theta}_\zeta$). Similar methods

for obtaining estimates of the standard error of $\hat{\beta}$ are also described in [17].

The development in Li and Lin (2003) considers an individual-level noisy covariate, but holds also for group-level covariates, as in our setting. In this case, the random errors in the simulation step are iid at the group level, but the algorithm otherwise remains the same. Li and Lin (2003) show that this estimator is consistent and normally distributed *assuming that one has the correct extrapolation function.* They suggest using a quadratic function, but provide no justification for this choice.

In the case where $Z \sim N(\mu_x, \sigma_z^2)$, it is easy to show that $\hat{\beta}_\zeta = \frac{\sigma_z^2}{\sigma_z^2+(\zeta+1)\sigma_u^2}\beta$ and $\hat{\theta}_\zeta = \theta + \frac{\sigma_z^2\cdot(\zeta+1)\sigma_u^2}{\sigma_z^2+(\zeta+1)\sigma_u^2}\beta^2$. This indicates that the correct extrapolation function for β is to model $\hat{\beta}_\zeta^{-1}$ as a linear function of ζ; however, this imposes a distributional assumption that we are trying to avoid by using the functional model. An alternative would be to approximate the true extrapolation function via Taylor series by fitting a polynomial function of ζ. We compare via simulation the differences in results that could be obtained using different functional forms.

For this study, we use a setup similar to [17], with some notational differences. Survival times T_{ij} are generated from the conditional hazard $\lambda_{ij}(t) = \lambda_0(t)\exp\left(\beta_z Z_i + \beta_x X_{ij} + b_i\right), j = 1, \cdots, n; i = 1, \cdots, m$. The Z_i are the cluster-level treatment variable of interest, while the X_{ij} (representing other measured covariates) are standard normal, and the cluster-level random effects $b_i \sim N(0, \theta)$. We compare results for $Z_i \sim N(\mu_z, \sigma_z^2)$ and $Z_i \sim \text{Gamma}(\alpha = 2, \beta = \sqrt{2})$. Censoring times are drawn uniformly on $[0, r]$, with r chosen to yield 0% or 30% censoring. The baseline hazard $\lambda_0(t) = 2t$.

In the simulated experiments, $m = 40$ or 100, $n = 3$, $\beta_z = 2$, $\beta_x = 1$, and $\theta = 0.5$. We generate a noisy $W_i = Z_i + U_i$, where $U_i \sim N(0, \sigma_u^2 = 0.5)$. In the SIMEX procedure, C = 200 and $\zeta \in \{0.25, 0.5, \cdots, 3\}$. The extrapolation functions compared are: two- through six-degree polynomials in ζ; second and third degree polynomials with e^ζ; and the inverse model linear in ζ suggested by the theory for normal Z.

Table 1 gives the results for experiments with 0% censoring; results for 30% censoring (and for individual-level Z) are similar and available upon request. Each experiment averages over 200 simulated data sets. The first two columns provide the estimates of β_z and θ when fitting a frailty model with the true Z and the noisy W, respectively. As expected, the estimate of β_z is attenuated in the presence of measurement error and the frailty variance θ is over-estimated. When comparing the extrapolation models in the SIMEX procedure, we see that the accuracy improves steadily with polynomial degree up to degree 5 or 6. Adding exponential terms does not lead to a significant improvement over the polynomial. When Z is normally distributed, the inverse model for β_z performs best, as expected given that it is the correct function. Even with gamma-distributed Z, however, the inverse is among the best models for β_z.

Normally-distributed Z

m		Model for hazard $X+Z$	Model for hazard $X+W$	Extrapolation model ζ^2	ζ^3	ζ^4	ζ^5	ζ^6	ζ^2+e^ζ	ζ^3+e^ζ	Inverse
40	β_z	1.99	1.58	1.89	1.95	1.97	1.98	1.95	1.92	1.96	1.99
	θ	0.48	1.27	0.63	0.52	0.51	0.50	0.66	0.56	0.52	–
100	β_z	1.99	1.19	1.59	1.74	1.82	1.88	1.96	1.68	1.80	2.03
	θ	0.47	0.38	0.43	0.43	0.43	0.45	0.45	0.43	0.43	–

Gamma-distributed Z

m		Model for hazard $X+Z$	Model for hazard $X+W$	Extrapolation model ζ^2	ζ^3	ζ^4	ζ^5	ζ^6	ζ^2+e^ζ	ζ^3+e^ζ	Inverse
40	β_z	1.99	1.55	1.86	1.92	1.94	1.97	2.03	1.89	1.93	1.97
	θ	0.52	1.30	0.70	0.62	0.56	0.51	0.21	0.65	0.58	–
100	β_z	1.96	1.52	1.85	1.92	1.95	1.96	2.02	1.89	1.94	1.96
	θ	0.44	1.13	0.60	0.50	0.44	0.38	0.25	0.54	0.46	–

Table 1: **Bias due to measurement error and SIMEX errors-in-variables adjustment results for a variety of extrapolation models.** The first two columns give the average parameter estimates from models with the true simulated Z and the noisy W. An extrapolation model described as ζ^i represents a polynomial function of ζ with maximum power i. "Inverse" indicates a model of the inverse of the parameter β_z as a linear function of ζ (note that this is not theoretically supported as a model for θ). Results shown are for 0% censoring.

6 Application to design issues in Botswana CRT

To demonstrate the performance of the adjustment in a more realistic setting, we apply the mixing adjustment to the results of an agent-based simulation of a cluster randomized trial of HIV prevention originally reported in Wang et al. [8]. In that model, the epidemic spreads over dynamic networks in paired clusters of equal size. The proportion of relationships that are between clusters range from 0 (no mixing) to 50% (clusters are effectively indistinguishable). For each mixing level there are 1000 simulated pairs of clusters. See the original article for details of the epidemic model and intervention.

We divide the 1000 pairs into simulated CRTs of 10 pairs (consistent with the required number of clusters per arm to achieve 90% power found in [8]), yielding 100 simulated data sets. For each data set we compute the randomized treatment effect estimate and standard error, and the overall treatment effect estimate and standard error (where appropriate).

Figure 2 shows the decline in the randomized treatment effect as mixing increases from 0 to 50% of relationships and the resulting estimates of the overall treatment effect when adjusting for mixing. Note that we cannot adjust the estimate for 50% mixing, because the difference between within and between cluster mixing rates is precisely 0 in this case, and the randomized treatment effect estimate is very close to zero. The dashed black line gives the treatment effect in the absence of mixing, while the dashed red line at 0 is included for reference when assessing the significance of the result.

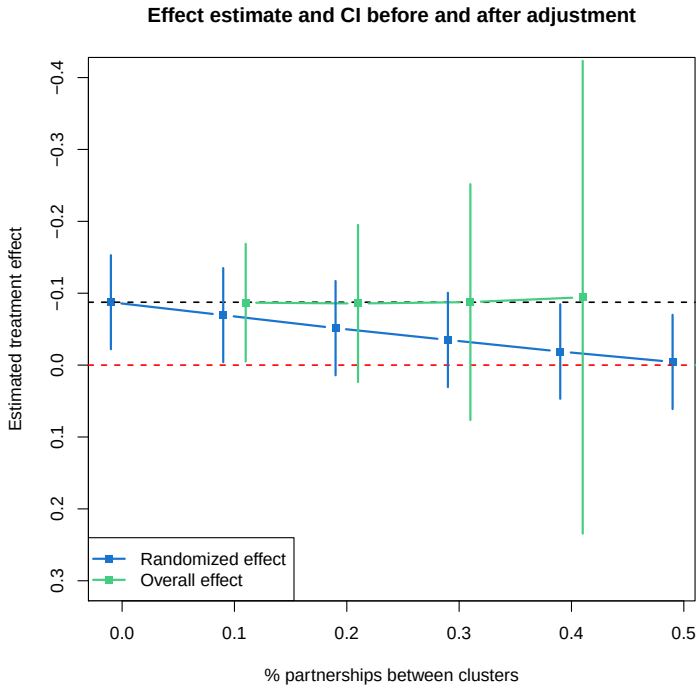

Figure 2: **Community–randomized trial model study results.** Randomized and overall treatment effect estimates over a range of mixing rates for Botswana CRT simulation. Vertical bars give the 95% CI for each estimate (note that mixing values are slightly jittered to prevent overlap of CIs).

In this application, the mass action assumption of the compartmental models described in Section 2 is violated, as contacts are contained within relationships that are sustained over time. In addition, the treatment effect is not a population-wide additive change in the transmissibility (η) parameter, but rather arises from different prevention modalities with differing efficacy being adopted by members of the population. Nonetheless, we see that the weighting method accurately recovers the overall treatment effect, even for relatively large mixing rates. The uncertainty increases substantially and, as expected, in all cases where the CI for the randomized treatment effect includes zero, the CI for the overall treatment effect does as well.

7 Discussion

Interference between units is a challenging problem in causal inference. While trials using two-stage randomization can identify causal effects in the presence of interference, this approach may be infeasible for cluster- or community-level outcomes for two reasons. First, it requires randomization of super-clusters and smaller clusters within them, which may require more experimental units than are available or feasible to include. Second, it would still require an assumption of no interference between super-clusters, which

may not be any more reasonable than assuming no interference between clusters. This paper presents an alternative approach to extract the overall treatment effect from a traditional cluster-randomized trial design with interference between clusters when the interactions between clusters causing the interference can be measured. The methods weight the treatment variable in a frailty or proportional hazards model, and so are simple to implement. Intuitively, the approach assigns each community a fractional treatment exposure and rescales the overall effect to reflect that a treatment cluster may have less than 100% of exposure from treated contacts, while a control cluster may actually be exposed some treated contacts.

This adjustment is applicable in most settings, but as mixing approaches 50%, the adjustment will become unstable as the effect will be nearly zero as will the inverse of the adjustment factor. When mixing is precisely 50%, no adjustment is possible. This situation is unlikely to arise, however, if clusters are defined with sufficient consideration of the issue of mixing. A mixing rate of 50% would effectively imply that two clusters are arbitrary subsets of the same group; clusters need to be selected with a goal of minimizing mixing through separation whether geographically, socially, or in some other way.

The development presented here requires a fairly strict set of assumptions, including that of mass-action assumption in the epidemic model and of constant prevalence, but the results for the design study for the BCPP suggest that the results are not sensitive to violations of the epidemic model assumptions. Further work is required to relax the assumption of constant prevalence. In the particular case of HIV in a mature epidemic, it may in fact be quite reasonable. Because HIV infection is of such long duration, even if new infections were sharply reduced, it may take years for the prevalence to decline as the infected population died of causes related to the disease or otherwise.

The fraction of social contacts sufficient for transmission that happen within treatment clusters can be difficult to measure, and is likely to be subject to measurement error. This error will further attenuate the estimate of the treatment effect, so we recommend consideration of the use of an errors-in-variables (EIV) adjustment in the implementation of the model. Section 5 presented a functional approach to EIV correction, SIMEX, using simulation and extrapolation to adjust for mismeasurement without making assumptions about the functional form of the mismeasured covariate. We also provided guidelines for selecting an extrapolation function; evidence from simulation studies suggests that a high degree (≥ 5) polynomial function will yield better results.

The methods discussed here allow estimation of the overall treatment effect in cases where contacts sufficient for transmission occur between treatment and control clusters. Another potential source of treatment effect attenuation that we do not address is the adoption of the treatment in control clusters.

This may be particularly likely to occur in educational interventions, where social contacts between clusters can spread the intended message into the control clusters.

References

[1] Michael E. Sobel. What do randomized studies of housing mobility demonstrate?: causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476):1398–1407, 2006.

[2] Guanglei Hong and Stephen W. Raudenbush. Evaluating kindergarten retention policy: a case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 101(475):901–910, 2006.

[3] Michael G. Hudgens and M. Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008. doi: 10.1198/016214508000000292.

[4] A Kamali, M Quigley, J Nakiyingi, J Kinsman, J Kengeya-Kayondo, R Gopal, A Ojwiya, P Hughes, L M Carpenter, and J Whitworth. Syndromic management of sexually-transmitted infections and behaviour change interventions on transmission of HIV-1 in rural Uganda: a community randomised trial. *Lancet*, 361:645–652, 2003.

[5] Rachel Jewkes, M Nduna, J Levin, N Jama, K Dunkle, A Puran, and N Duwury. Impact of Stepping Stones on incidence of HIV and HSV-2 and sexual behaviour in rural South Africa: cluster randomised controlled trial. *BMJ*, 337:a606, 2008. doi: 10.1136/bmj.a506.

[6] Dylan S. Small, Thomas R. Ten Have, and Paul R. Rosenbaum. Randomization inference in a group-randomized trial of treatments for depression: Covariate adjustment, noncompliance, and quantile effects. *Journal of the American Statistical Association*, 103(481):271–279, 2008. doi: 10.1198/016214507000000897.

[7] ALfred B. Tiono, Alphonse Ouédraogo, Bernhards Ogutu, Amidou Diarra, Sam Coulibaly, Adama Gansané, Sodiomon B. Sirima, Gregory O'Neil, Amitava Mukhopadhyay, and Kamal Hamed. A controlled, parallel, cluster-randomized trial of community-wide screening and treatment of asymptomatic carriers of *Plasmodium falciparum* in Burkina Faso. *Malaria Journal*, 12:79, 2013. doi: 10.1186/1475-2875-12-79.

[8] Rui Wang, Ravi Goyal, Quanhong Lei, M Essex, and Victor De Gruttola. Sample size considerations in the design of cluster randomized trials of combination HIV prevention. *Clinical Trials*, 11(309–318), 2014.

[9] Eric J. Tchetgen Tchetgen and Tyler J. Vander Weele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75, 2012. doi: 10.1177/0962280210386779.

[10] Todd Datz. Harvard School of Public Health awarded $20 million CDC grant to study HIV prevention in Botswana. *HSPH News*, September 20, 2011, 2011. URL `http://www.hsph.harvard.edu/news/press-releases/20-million-grant-hiv-botswana [Last accessed: 2 May 2013]`.

[11] R. L. Prentice. Covariate measurement errors and parameter estimation in a failure time regression model. *Biometrika*, 69(2):331–42, 1982.

[12] Tsuyoshi Nakamura. Proportional hazards model with covariates subject to measurement error. *Biometrics*, 48:829–838, 1992.

[13] Michael D. Hughes. Regression dilution in the proportional hazards model. *Biometrics*, 49:1056–1066, 1993.

[14] Yijian Huang and C. Y. Wang. Cox regression with accurate covariates unascertainable: a nonparametric-correction approach. *Journal of the American Statistical Association*, 95(452):1209–1219, 2000.

[15] Chengcheng Hu and D. Y. Lin. Cox regression with covariate measurement error. *Scandinavian Journal of Statistics*, 29:637–655, 2002.

[16] Yi Li and Xihong Lin. Covariate measurement errors in frailty models for clustered survival data. *Biometrika*, 87(4):849–866, 2000.

[17] Yi Li and Xihong Lin. Functional inference in frailty measurement error models for clustered survival data using the SIMEX approach. *Journal of the American Statistical Association*, 98(461):191–203, 2003.