



Published in final edited form as:

Biometrics. 2022 June ; 78(2): 777–788. doi:10.1111/biom.13459.

Inverse Probability Weighted Estimators of Vaccine Effects Accommodating Partial Interference and Censoring

Sujatro Chakladar¹, Samuel P. Rosin¹, Michael G. Hudgens^{1,*}, M. Elizabeth Halloran^{2,3}, John D. Clemens^{4,5}, Mohammad Ali⁶, Michael E. Emch⁷

¹Department of Biostatistics, University of North Carolina, Chapel Hill, North Carolina, U.S.A.

²Department of Biostatistics, University of Washington, Seattle, Washington, U.S.A.

³Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Center, Seattle, Washington, U.S.A.

⁴Department of Epidemiology, University of California, Los Angeles, California, U.S.A.

⁵International Centre for Diarrhoeal Disease Research, Bangladesh

⁶Department of International Health, Johns Hopkins University, Baltimore, Maryland, U.S.A.

⁷Department of Geography, University of North Carolina, Chapel Hill, North Carolina, U.S.A.

Summary:

Estimating population-level effects of a vaccine is challenging because there may be interference, i.e., the outcome of one individual may depend on the vaccination status of another individual. Partial interference occurs when individuals can be partitioned into groups such that interference occurs only within groups. In the absence of interference, inverse probability weighted (IPW) estimators are commonly used to draw inference about causal effects of an exposure or treatment. Tchetgen Tchetgen and VanderWeele (2012) proposed a modified IPW estimator for causal effects in the presence of partial interference. Motivated by a cholera vaccine study in Bangladesh, this paper considers an extension of the Tchetgen Tchetgen and VanderWeele IPW estimator to the setting where the outcome is subject to right censoring using inverse probability of censoring weights (IPCW). Censoring weights are estimated using proportional hazards frailty models. The large sample properties of the IPCW estimators are derived, and simulation studies are presented demonstrating the estimators' performance in finite samples. The methods are then used to analyze data from the cholera vaccine study.

Keywords

Causal Inference; Interference; Right censoring; Survival

* mhudgens@email.unc.edu .

Supporting Information

Web Appendix A, Tables, and Figures referenced in Sections 2-4 are available with this paper at the Biometrics website on Wiley Online Library. R code based on the simulations in Section 3 is included with the Supporting Information and is also available at <https://github.com/samrosin/interference-IPCW>.

1. Introduction

Estimating population-level effects of a vaccine is challenging because there may be interference, i.e., the outcome of one individual may depend on the vaccination status of another individual. Partial interference is a special case of interference where individuals can be partitioned into groups such that interference does not occur between individuals in different groups but may occur between individuals in the same group (Sobel 2006). Partial interference might be a reasonable assumption if groups of individuals are sufficiently separated geographically, socially, and/or temporally. In this paper we consider assessing the effects of cholera vaccination based on a study in Matlab, Bangladesh (Ali et al. 2005), assuming partial interference based on the spatial location of residences of study participants (Perez-Heydrich et al. 2014). Effects due to interference, also known as spillover effects or peer effects, are of interest in many areas, including criminology, developmental psychology, econometrics, education, political science, social media and network analysis.

Inferential methods about spillover effects have been developed for randomized experiments (Rosenbaum, 2007; Hudgens and Halloran, 2008; Baird et al., 2018). However, in some settings it may not be feasible or ethical to randomize groups or individuals to different treatment or exposure conditions. In the observational setting, Tchetgen Tchetgen and VanderWeele (henceforth TV) (2012) proposed inverse probability weighted (IPW) estimators for different causal effects when there may be partial interference. Large sample properties of these IPW estimators were considered by Perez-Heydrich et al. (2014) and Liu et al. (2016). While motivated by observational studies, these estimators may also be applied in cluster randomized trials where partial interference is assumed and there is non-compliance, i.e., not all individuals receive the treatment assigned to their cluster. These estimators are also applicable to settings such as the cholera vaccine study where all individuals in the study were randomized but only a subset chose to participate in the trial.

In settings where the outcome of interest is a time to event, the outcome may be subject to right censoring due to study completion or participant drop-out. For example, in the Bangladesh cholera vaccine trial, some study participants emigrated out of the field trial area and hence were lost to follow-up. To date most methods that permit interference do not allow for right censored outcomes. One exception is Loh et al. (2020), who allow for right censored outcomes and interference, but require treatment to be randomized. In the absence of interference, censoring is often accommodated via inverse probability of censoring weights along with inverse probability treatment weights (e.g., Robins and Finkelstein (2000)).

In this paper, an extension of the TV IPW estimator is considered for observational studies where there may be partial interference and the outcome is subject to right censoring using inverse probability of censoring weights (IPCW). The proposed methods are developed in Section 2. In Section 3 simulation results are presented demonstrating the empirical performance of the proposed methods in finite sample settings. In Section 4 the methods are used to analyze the Bangladesh cholera vaccine study. Section 5 concludes with a discussion.

2. Methods

2.1 Estimands

Suppose data are observed from m groups of individuals, with n_j individuals per group for $i = 1, \dots, m$. In the cholera vaccine study analysis presented below, participants are partitioned into $m = 700$ groups (neighborhoods) according to the geographical location of their household of residence. Let $A_{ij} = 1$ if individual j in group i receives treatment (e.g., vaccine) and $A_{ij} = 0$ otherwise. Let $\mathbf{A}_i = (A_{i1}, A_{i2}, \dots, A_{in_i})$ and $\mathbf{A}_{i,-j} = (A_{i1}, A_{i2}, \dots, A_{ij-1}, A_{ij+1}, \dots, A_{in_i})$. Let \mathbf{a}_i and $\mathbf{a}_{i,-j}$ denote possible realizations of \mathbf{A}_i and $\mathbf{A}_{i,-j}$ and let $\mathcal{A}(n)$ denote the set of all possible 2^n treatments for a group size of $n = 1, 2, \dots$. Assume partial interference and denote the potential time to event for individual j in group i if, possibly counter to fact, group i receives treatment \mathbf{a}_i by $T_{ij}(\mathbf{a}_i)$. In the motivating vaccine study, we are interested in time until incident cholera. The notation $T_{ij}(\mathbf{a}_i)$ reflects the partial interference assumption, i.e., the potential outcome of individual j in group i does not depend on the treatment of individuals outside group i . Below the notation $T_{ij}(a, \mathbf{a}_{i,-j})$ is sometimes used to make explicit the treatment for individual j and the treatment for all other individuals in group i . Let $\mathbf{T}_i(\cdot) = \{T_{ij}(\mathbf{a}_i) : \mathbf{a}_i \in \mathcal{A}(n_i), j = 1, 2, \dots, n_i\}$ denote the set of all potential event times for individuals in group i .

Suppose the event times are subject to right censoring, e.g., due to loss to follow-up or study completion. Let C_{ij} denote the potential censoring times for individual j in group i . Let $\delta_{ij} = 1$ if $T_{ij}(\mathbf{A}_i) \leq C_{ij}$ and $\delta_{ij} = 0$ otherwise, and let $X_{ij} = \min(T_{ij}(\mathbf{A}_i), C_{ij})$. Define $\mathbf{X}_i = (X_{i1}, X_{i2}, \dots, X_{in_i})$ and $\mathbf{i} = (i_1, i_2, \dots, i_{n_i})$. Denote by \mathbf{L}_{ij} the vector of baseline covariates for subject j in group i and by \mathbf{L}_i baseline covariates for all subjects in group i , i.e., $\mathbf{L}_i = (\mathbf{L}_{i1}, \mathbf{L}_{i2}, \dots, \mathbf{L}_{in_i})$. Assume that the m groups are randomly sampled from an infinite superpopulation of groups such that the observed data are m i.i.d. copies of $\mathbf{O}_i = (\mathbf{L}_i, \mathbf{A}_i, \mathbf{X}_i, \mathbf{i})$.

In the absence of interference, treatment effects are typically defined as contrasts in mean potential outcomes for different counterfactual scenarios, e.g., the average treatment effect is usually defined as the difference in the mean potential outcome had all individuals received treatment versus had no individuals received treatment. Similarly, in the setting where there is partial interference, causal effects may be defined as contrasts in mean potential outcomes for different counterfactual scenarios. Here we consider counterfactual scenarios where the marginal probability that an individual receives treatment, $\Pr_\alpha(A_{ij} = 1)$, equals α for different values of $\alpha \in (0, 1)$. The notation $\Pr_\alpha(\cdot)$ indicates that the probability corresponds to the distribution under the counterfactual scenario. Specifically, the Bernoulli treatment allocation strategy (or policy) described in TV is considered wherein individuals independently select treatment with probability α . Let $\pi(\mathbf{a}_i, \alpha)$ denote the probability that group i receives treatment \mathbf{a}_i under Bernoulli allocation strategy α . That is, $\pi(\mathbf{a}_i, \alpha) = \Pr_\alpha(\mathbf{A}_i = \mathbf{a}_i) = \prod_{k=1}^{n_i} \alpha^{a_{ik}} (1 - \alpha)^{1 - a_{ik}}$. Similarly let $\pi(\mathbf{a}_{i,-j}, \alpha) = \Pr_\alpha(\mathbf{A}_{i,-j} = \mathbf{a}_{i,-j} \mid A_{ij} = a) = \prod_{k=1, k \neq j}^{n_i} \alpha^{a_{ik}} (1 - \alpha)^{1 - a_{ik}}$.

The causal estimands of interest defined below are contrasts in the risk of having an event by time t for different combinations of treatment a and allocation strategies α . To define these estimands, let

$$\bar{F}_{ij}(t, a, \alpha) = \sum_{\mathbf{a}_i, -j \in \mathcal{A}(n_i - 1)} I\{T_{ij}(a, \mathbf{a}_i, -j) \leq t\} \pi(\mathbf{a}_i, -j, \alpha),$$

and

$$\bar{F}_i(t, \alpha) = \sum_{\mathbf{a}_i \in \mathcal{A}(n_i)} I\{T_{ij}(\mathbf{a}_i) \leq t\} \pi(\mathbf{a}_i, \alpha).$$

In words, $\bar{F}_{ij}(t, a, \alpha)$ is the probability that individual j in group i will have an event by time t when receiving treatment a and the group adopts policy α . Likewise, $\bar{F}_i(t, \alpha)$ is the probability that individual j in group i will have an event by time t when the group adopts allocation strategy α . Denote the group average risks by $\bar{F}_i(t, a, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \bar{F}_{ij}(t, a, \alpha)$ and $\bar{F}_i(t, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \bar{F}_{ij}(t, \alpha)$. Let $\mu(t, a, \alpha) = E_\alpha\{\bar{F}_i(t, a, \alpha)\}$ and $\mu(t, \alpha) = E_\alpha\{\bar{F}_i(t, \alpha)\}$ where $E_\alpha\{\cdot\}$ denotes the expected value under the counterfactual setting when policy α is adopted in the superpopulation of groups. In the cholera vaccine study described in Section 4, $\mu(t, a, \alpha)$ denotes the average risk of acquiring cholera by time t when an individual receives treatment a and other individuals receive vaccine with probability α .

Various effects of treatment can be defined by contrasts in $\mu(t, a, \alpha)$ and $\mu(t, \alpha)$. The direct effect is obtained by comparing the probability of an event when an individual receives treatment versus when not receiving treatment for a fixed allocation strategy. In particular, the direct effect at time t corresponding to policy α is defined to be $DE(t, \alpha) = \mu(t, 0, \alpha) - \mu(t, 1, \alpha)$. The indirect (or spillover) effect is the difference in the probability of an event by time t for two different policies when the individual does not receive treatment. Specifically, the indirect effect is given by $IE(t, \alpha_1, \alpha_2) = \mu(t, 0, \alpha_1) - \mu(t, 0, \alpha_2)$ for allocation strategies α_1 and α_2 . An indirect effect can analogously be defined when an individual is vaccinated. The total effect is the difference between the probability of an event by time t when an individual does not receive treatment under policy α_1 and when an individual receives treatment under policy α_2 , i.e., $TE(t, \alpha_1, \alpha_2) = \mu(t, 0, \alpha_1) - \mu(t, 1, \alpha_2)$. Finally, the overall effect is the difference between the probability of an event by time t for policy α_1 versus α_2 , i.e., $OE(t, \alpha_1, \alpha_2) = \mu(t, \alpha_1) - \mu(t, \alpha_2)$.

2.2 Assumptions

Assume the following:

- i. Conditional independent treatment: $\mathbf{A}_j \perp\!\!\!\perp \mathbf{T}_j(\cdot) \mid \mathbf{L}_j$
- ii. Treatment positivity: $\Pr(\mathbf{A}_j = \mathbf{a}_j \mid \mathbf{L}_j) > 0$ for all $\mathbf{a}_j \in \mathcal{A}(n_j)$
- iii. Conditional independent censoring: $C_{ij} \perp\!\!\!\perp T_{ij}(\mathbf{A}_j) \mid \{\mathbf{L}_j, \mathbf{A}_j\}$
- iv. Non-censoring positivity: $\Pr(ij = 1 \mid \mathbf{L}_j, \mathbf{A}_j) > 0$

Assumption I states that the potential event times for individuals within the same group are conditionally independent of the treatment received by the group given covariates; this is a group-level generalization of the usual individual-level no unmeasured confounders assumption often made in the absence of interference. In the cholera vaccine study, Assumption I would be violated if there was an unobserved common cause of one or more components of vaccination status and time to cholera. For instance, individuals who were more concerned about cholera may have been more likely to participate in the trial and also more likely to practice other measures (e.g., hand washing) to prevent cholera. Treatment positivity assumes that each group has a non-zero probability of being assigned every possible treatment combination given covariates for the group (Perez-Heydrich et al., 2014). Assumption II is reasonable in the cholera study because all women and children in Matlab were invited to participate in the vaccine trial. Assumption III supposes that conditional on baseline group covariates and group treatment, an individual's failure time is independent of their censoring time. Assumption III would be violated in the cholera study if there was some variable other than \mathbf{L}_j and \mathbf{A}_j which were prognostic of both the censoring and time to cholera. For example, this assumption might not hold if individuals living in areas of Matlab with a high burden of cholera were more likely to emigrate outside of the study area. Note censoring times between individuals within the same cluster need not be independent; indeed, the frailty model introduced below allows for such dependency via a cluster random effect. Finally Assumption IV indicates that each individual has a non-zero probability of not being censored at each observation time (Rotnitzky et al., 2007). In the next section IPW estimators are proposed and shown to be consistent (and asymptotically normal) for the direct, indirect, total, and overall effects under Assumptions I-IV.

2.3 Proposed Estimator

In the absence of censoring, the IPW estimator proposed by TV can be used to draw inference about $\mu(t, a, \alpha)$ and $\mu(t, \alpha)$, i.e., the mean potential outcomes under the counterfactual setting where policy α is adopted. In particular, letting $Y_{ij} = I(X_{ij} \leq t)$ be the indicator variable that the observation time for individual j in group i is less than or equal to t , the TV IPW estimators are $\hat{\mu}^{TV}(t, a, \alpha) = m^{-1} \sum_{i=1}^m \hat{F}_i^{TV}(t, a, \alpha)$ and $\hat{\mu}^{TV}(t, \alpha) = m^{-1} \sum_{i=1}^m \hat{F}_i^{TV}(t, \alpha)$ where

$$\hat{F}_i^{TV}(t, a, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{\pi(\mathbf{A}_i, -j; \alpha) I(A_{ij} = a) Y_{ij}}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}})}, \quad \hat{F}_i^{TV}(t, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{\pi(\mathbf{A}_i; \alpha) Y_{ij}}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}})},$$

and $\hat{\boldsymbol{\beta}}$ is an estimator of the vector of parameters for the propensity model $\Pr(\mathbf{A}_i | \mathbf{L}_i, \boldsymbol{\beta})$. Details of the propensity model are discussed in the next sections.

In the presence of censoring, the following extension of the TV IPW estimators is proposed:

$$\hat{\mu}(t, a, \alpha) = m^{-1} \sum_{i=1}^m \hat{F}_i(t, a, \alpha) \text{ and } \hat{\mu}(t, \alpha) = m^{-1} \sum_{i=1}^m \hat{F}_i(t, \alpha) \text{ where}$$

$$\hat{F}_i(t, a, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{\pi(\mathbf{A}_i, -j; \alpha) I(A_{ij} = a) I(\Delta_{ij} = 1) I(X_{ij} \leq t)}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}}) S_C(X_{ij} | \mathbf{L}_i, \mathbf{A}_i, \hat{\boldsymbol{\gamma}})},$$

$$\hat{F}_i(t, \alpha) = n_i^{-1} \sum_{j=1}^{n_i} \frac{\pi(\mathbf{A}_i; \alpha) I(\Delta_{ij} = 1) I(X_{ij} \leq t)}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}}) S_C(X_{ij} | \mathbf{L}_i, \mathbf{A}_i, \hat{\boldsymbol{\gamma}})},$$

$S_C(t | \mathbf{L}_i, \mathbf{A}_i, \hat{\boldsymbol{\gamma}}) = \Pr(C_{ij} > t | \mathbf{L}_i, \mathbf{A}_i, \hat{\boldsymbol{\gamma}})$ and $\hat{\boldsymbol{\gamma}}$ is an estimator of the vector of the parameters for the censoring model. Details of the censoring model are discussed in the next sections. Estimates of the direct, indirect, total, and overall effects are given by $\widehat{DE}(t, \alpha) = \hat{\mu}(t, 0, \alpha) - \hat{\mu}(t, 1, \alpha)$, $\widehat{IE}(t, \alpha_1, \alpha_2) = \hat{\mu}(t, 0, \alpha_1) - \hat{\mu}(t, 0, \alpha_2)$, $\widehat{TE}(t, \alpha_1, \alpha_2) = \hat{\mu}(t, 0, \alpha_1) - \hat{\mu}(t, 1, \alpha_2)$ and $\widehat{OE}(t, \alpha_1, \alpha_2) = \hat{\mu}(t, \alpha_1) - \hat{\mu}(t, \alpha_2)$.

The proposition below shows that if the group level propensity scores and the individual censoring probabilities are known, then the proposed IPCW estimators are unbiased. A proof of the proposition is given in Web Appendix A.

Proposition. *If $\Pr(\mathbf{A}_i | \mathbf{L}_i)$ and $S_C(X_{ij} | \mathbf{L}_i, \mathbf{A}_i)$ are known for $j = 1, 2, \dots, n_i$ and $i = 1, \dots, m$, then $E\{\hat{\mu}(t, a, \alpha)\} = \mu(t, a, \alpha)$ and $E\{\hat{\mu}(t, \alpha)\} = \mu(t, \alpha)$.*

In observational studies, the conditional distribution of treatment given covariates is unknown. Likewise, in both observational studies as well as randomized trials, the conditional distribution of censoring given covariates is typically not known (one exception being studies or trials without drop-out such that the only cause of censoring is the end of administrative follow-up at some fixed time point). Therefore, we consider finite dimensional parametric models to estimate the group propensity scores and conditional probability of censoring; these estimates are then plugged into the IPCW estimators defined above.

The conditional probability of censoring is estimated using a shared frailty model (Munda et al., 2012) where the conditional hazard for C_{ij} is assumed to have the proportional hazards form $g_{ij}(c | \mathbf{L}_i, \mathbf{A}_i, e_i) = g_0(c; \boldsymbol{\theta}_h) e_i \exp(\tilde{\mathbf{L}}_{ij} \boldsymbol{\theta}_c)$ where g_0 is the baseline hazard function, $\boldsymbol{\theta}_h$ is the q' -dimensional parameter vector of the baseline hazard function, e_i is a random effect with density $f_e(e_i; \boldsymbol{\theta}_r)$, $\tilde{\mathbf{L}}_{ij}$ is some user specified function of $\{\mathbf{L}_i, \mathbf{A}_i\}$, and $\boldsymbol{\theta}_c$ is the q -dimensional column vector of coefficients. The row vector $\tilde{\mathbf{L}}_{ij}$ could include, for example, covariates and treatment for individual j (i.e., \mathbf{L}_{ij} and A_{ij}) as well as the proportion of others in the group who receive treatment (i.e., $\sum_{k \neq j} A_{ik} / (n_i - 1)$). Below the dependence of g_0 on $\boldsymbol{\theta}_h$ is suppressed for notational convenience. Let $\boldsymbol{\gamma} = (\boldsymbol{\theta}_c^\top, \boldsymbol{\theta}_h^\top, \boldsymbol{\theta}_r)$ be the vector of parameters for the frailty model. Maximum likelihood theory can be used to draw inference about $\boldsymbol{\gamma}$. Under assumption III, the contribution of group i to the log-likelihood corresponding to the frailty censoring model, denoted by $l(\mathbf{X}_i, \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma})$, equals

$$\sum_{j=1}^{n_i} (1 - \Delta_{ij}) [\log\{g_0(X_{ij})\} + \tilde{L}_{ij}\boldsymbol{\theta}_c] + \log\{(-1)^{d_i} \mathcal{L}^{(d_i)}(\sum_{j=1}^{n_i} G_0(X_{ij}) \exp(\tilde{L}_{ij}\boldsymbol{\theta}_c))\},$$

where $d_i = \sum_{j=1}^{n_i} (1 - \Delta_{ij})$ is the number of censored observations in group i , $G_0(\omega) = \int_0^\omega g_0(\kappa) d\kappa$, $\mathcal{L}(s) = \int_0^\infty \exp(-e_i s) f_e(e_i; \theta_r) de_i$ and $\mathcal{L}^{(k)}$ is the k^{th} -derivative of \mathcal{L} (Munda et al., 2012). Therefore, the maximum likelihood estimator of $\boldsymbol{\gamma}$ solves the following estimating equations

$$\sum_i \psi_{ck}(\mathbf{X}_i, \boldsymbol{\Delta}_i, \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma}) = 0 \text{ for } k = 1, \dots, q + q' + 1,$$

where $\psi_{ck} = \psi_{ck}(\mathbf{X}_i, \boldsymbol{\Delta}_i, \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma}) = \ell(\mathbf{X}_i, \boldsymbol{\Delta}_i, \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma}) / \gamma_k$ and γ_k is the k -th element of $\boldsymbol{\gamma}$. Below, the baseline hazard for the censoring model is assumed to be constant and equal to θ_h , and the frailty term e_i is assumed to follow a Gamma distribution with mean 1 and variance θ_r , such that censoring weights for an uncensored individual can be computed via

$$\begin{aligned} S_C(t, | \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma}) &= \int \Pr(C_{ij} > t | \tilde{L}_{ij}, \boldsymbol{\gamma}, e_i) f_e(e_i; \theta_r) de_i \\ &= \int \exp\{-\theta_h t \exp(\tilde{L}_{ij}\boldsymbol{\theta}_c) e_i\} \frac{e_i^{1/\theta_r - 1} \exp(-e_i / \theta_r)}{\theta_r^{1/\theta_r} \Gamma(1/\theta_r)} de_i \\ &= \{\theta_r \theta_h t \exp(\tilde{L}_{ij}\boldsymbol{\theta}_c) + 1\}^{-1/\theta_r} \end{aligned}$$

Following TV (2012), a mixed effects model may be assumed for the treatment allocation, i.e., $\Pr(A_{ij} = 1 | \mathbf{L}_{ij}, b_i) = \text{logit}^{-1}(\mathbf{L}_{ij}\boldsymbol{\theta}_x + b_i)$ where b_i is a random effect following density $f_b(b_i; \boldsymbol{\theta}_s)$. (In the application below the mixed effects model has a slightly more complicated form owing to the study design.) Let $\boldsymbol{\beta} = (\boldsymbol{\theta}_x^\top, \boldsymbol{\theta}_s)$ denote the $(p + 1)$ dimensional vector of parameters for the mixed effects model. Again, maximum likelihood theory can be used to draw inference about $\boldsymbol{\beta}$. The contribution of group i to the log-likelihood for the mixed effects model is given by $\ell(\mathbf{A}_i, \mathbf{L}_i, \boldsymbol{\beta}) = \log \Pr(\mathbf{A}_i | \mathbf{L}_i, \boldsymbol{\beta})$ where

$$\Pr(\mathbf{A}_i | \mathbf{L}_i, \boldsymbol{\beta}) = \int \prod_{j=1}^{n_i} h_{ij}(b_i, \mathbf{L}_i, \boldsymbol{\theta}_x)^{A_{ij}} \{1 - h_{ij}(b_i, \mathbf{L}_i, \boldsymbol{\theta}_x)\}^{(1 - A_{ij})} f_b(b_i; \boldsymbol{\theta}_s),$$

and $h_{ij}(b_i, \mathbf{L}_i, \boldsymbol{\beta}) = \Pr(A_{ij} = 1 | \mathbf{L}_{ij}, b_i)$. The maximum likelihood estimator of $\boldsymbol{\beta}$ is the solution to the score equations

$$\sum_i \psi_{xk}(\mathbf{A}_i, \mathbf{L}_i, \boldsymbol{\beta}) = 0 \text{ for } k = 1, \dots, p + 1,$$

where $\psi_{xk} = \psi_{xk}(\mathbf{A}_i, \mathbf{L}_i, \boldsymbol{\beta}) = \ell(\mathbf{A}_i, \mathbf{L}_i, \boldsymbol{\beta}) / \beta_k$ and β_k is the k -th element of $\boldsymbol{\beta}$.

Inference about the causal effects of interest is then based on solving the vector of estimating equations

$$\sum_i \psi(\mathbf{O}_i, \boldsymbol{\theta}) = 0, \quad (1)$$

where $\boldsymbol{\theta} = (\boldsymbol{\gamma}, \boldsymbol{\beta}, \theta)$, $\psi(\mathbf{O}_i, \boldsymbol{\theta}) = (\boldsymbol{\psi}_c, \boldsymbol{\psi}_x, \psi_{aa})^\top$, $\boldsymbol{\psi}_c = (\psi_{c1}, \psi_{c2}, \dots, \psi_{cq+q'+1})$, $\boldsymbol{\psi}_x = (\psi_{x1}, \psi_{x2}, \dots, \psi_{xp+1})$,

$$\psi_{aa} = \psi_{aa}(\mathbf{O}_i, \boldsymbol{\theta}) = \frac{g^*(\mathbf{O}_i, a, \alpha, \boldsymbol{\gamma})}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \boldsymbol{\beta})} - \theta,$$

and

$$g^*(\mathbf{O}_i, a, \alpha, \boldsymbol{\gamma}) = n_i^{-1} \sum_{j=1}^n \frac{\pi(\mathbf{A}_i, -j; \alpha) I(A_{ij} = a) I(A_{ij} = 1) I(X_{ij} \leq t)}{S_C(X_{ij} | \mathbf{L}_i, \mathbf{A}_i, \boldsymbol{\gamma})}.$$

Let $\hat{\boldsymbol{\theta}} = (\hat{\boldsymbol{\gamma}}, \hat{\boldsymbol{\beta}}, \hat{\mu}(t, a, \alpha))$ denote the solution to (1). Denote the true value of $\boldsymbol{\theta}$ by $\boldsymbol{\theta}_0 = (\boldsymbol{\gamma}_0, \boldsymbol{\beta}_0, \mu(t, a, \alpha))$ and note that

$$\int \psi_{aa}(\mathbf{o}, \boldsymbol{\gamma}_0, \boldsymbol{\beta}_0, \mu(t, a, \alpha)) dF_{\mathbf{O}}(\mathbf{o}) = E \left\{ \frac{g^*(\mathbf{O}_i, a, \alpha, \boldsymbol{\gamma}_0)}{\Pr(\mathbf{A}_i | \mathbf{L}_i, \boldsymbol{\beta}_0)} - \mu(t, a, \alpha) \right\} = 0,$$

where $F_{\mathbf{O}}$ denotes the joint distribution of the complete observed random variable \mathbf{O} and the last equality follows from the Proposition above. Therefore, assuming the parametric models above are correctly specified, it follows that $\int \psi(\mathbf{o}, \boldsymbol{\theta}_0) dF_{\mathbf{O}}(\mathbf{o}) = 0$. By M-estimation theory

(Stefanski and Boos, 2002), $\hat{\boldsymbol{\theta}} \xrightarrow{P} \boldsymbol{\theta}_0$ and $\sqrt{m}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$ converges in distribution to a Normal distribution with mean 0 and covariance matrix $\boldsymbol{\Sigma}$ equal to $U(\boldsymbol{\theta}_0)^{-1} V(\boldsymbol{\theta}_0) \{U(\boldsymbol{\theta}_0)^{-1}\}^T$ where $U(\boldsymbol{\theta}_0) = E\{-\dot{\psi}(\mathbf{O}_i, \boldsymbol{\theta}_0)\}$, $V(\boldsymbol{\theta}_0) = E\{\psi(\mathbf{O}_i, \boldsymbol{\theta}_0)\psi(\mathbf{O}_i, \boldsymbol{\theta}_0)^T\}$, and $\dot{\psi}(\mathbf{O}_i, \boldsymbol{\theta}) = \partial \psi(\mathbf{O}_i, \boldsymbol{\theta}) / \partial \boldsymbol{\theta}^T$.

Consistency and asymptotic normality of the direct, indirect and total effect estimators follows from the delta method. Similar techniques can be used to show that $\hat{\mu}(t, \alpha)$ and the overall effect estimator are also consistent and asymptotically Normal. The asymptotic variance $\boldsymbol{\Sigma}$ can be consistently estimated by $\hat{\boldsymbol{\Sigma}} = \hat{U}(\hat{\boldsymbol{\theta}})^{-1} \hat{V}(\hat{\boldsymbol{\theta}}) \{\hat{U}(\hat{\boldsymbol{\theta}})^{-1}\}^T$ where $\hat{U}(\hat{\boldsymbol{\theta}}) = m^{-1} \sum_{i=1}^m \{-\dot{\psi}(\mathbf{O}_i, \hat{\boldsymbol{\theta}})\}$ and $\hat{V}(\hat{\boldsymbol{\theta}}) = m^{-1} \sum_{i=1}^m \{\psi(\mathbf{O}_i, \hat{\boldsymbol{\theta}})\psi(\mathbf{O}_i, \hat{\boldsymbol{\theta}})^T\}$. The empirical sandwich variance estimator $\hat{\boldsymbol{\Sigma}}$ can be computed using the R package geex (Saul and Hudgens, 2020) and can be used to construct Wald confidence intervals (CIs).

3. Simulation Study

A simulation study was conducted to assess the finite sample bias of the IPCW estimator and coverage of the corresponding Wald confidence intervals. The data generating model used in the simulation study was motivated by aspects of the cholera vaccine study analysis presented in the next section. Following Perez-Heydrich et al. (2014), data were simulated according to the following steps.

- i. First, two baseline covariates L_{1ij} and L_{2ij} were randomly generated. In the application presented in Section 4, conditional independence (assumption I) is assumed given an individual's age (in decades) and the distance of their residence to the nearest river. Motivated by this example, L_{1ij} (age) and L_{2ij} (distance to river) were randomly generated as follows. First, V_{ij} was randomly generated from an Exponential distribution with mean 20, and r_{1i} and r_{2ij} were independently sampled from Normal(0, 0.1), i.e., a mean-zero Gaussian distribution with standard deviation 0.1. Then L_{1ij} was set to $\min(V_{ij} + r_{1i} + r_{2ij}, 100)/10$ and L_{2ij} randomly generated such that $\log L_{2ij} \sim \text{Normal}(r_{1i} + r_{2ij}, 0.75)$.
- ii. The random effects for the treatment model b_j were randomly sampled from a Normal distribution with mean 0 and variance 1.0859.
- iii. The treatment indicators A_{ij} were randomly sampled from a Bernoulli distribution with mean $p_{ij} = \text{expit}(0.2727 - 0.0387L_{1ij} + 0.2179L_{2ij} + b_j)$.
- iv. The potential times to event $T_{ij}(\mathbf{a}_j)$ were randomly sampled from an Exponential distribution with mean $\mu_{ij} = 200 + 100a_{ij} - 0.98L_{1ij} - 0.145L_{2ij} + 50 \sum_{k \neq j} a_{ik} / n_i$.
- v. The random effects for the censoring model e_i were randomly generated from a Gamma distribution with mean 1 and variance $\theta_r = 1.25$.
- vi. Censoring times C_{ij} were randomly sampled from an Exponential distribution with mean $1/\lambda_0$ where $\lambda_0 = 0.015 \exp(0.002L_{1ij} + 0.015L_{2ij})e_i$.
- vii. Individual censoring indicators were determined, i.e., $\delta_{ij} = 0$ if $C_{ij} < T_{ij}(\mathbf{A}_j)$.

Steps i through vii were used to stochastically generate 1000 data sets, with each data set containing 500 groups with 10 individuals per group. For each simulated data set, the IPCW estimator of $\mu(100, a, \alpha)$ was evaluated for $a = 0, 1$ and $\alpha = 0.1, 0.2, \dots, 0.9$. Estimated standard errors based on the empirical sandwich variance estimator and Wald 95% confidence intervals were also calculated for each simulated data set. Empirical standard errors were calculated by taking the standard deviation of the point estimates from all simulations.

The true value of the estimand was obtained by simulating counterfactual outcomes for $m = 10^6$ groups of individuals. Note that, according to the model used to generate the data, potential survival times depend only on $\sum_{k \neq j} a_{ik}$ SO, $\mu(t, a, \alpha)$ was approximated by (Perez-Heydrich et al., 2014)

$$m^{-1} \sum_{i=1}^m n_i^{-1} \sum_{j=1}^{n_i} \sum_{k=0}^{n_i-1} \binom{n_i-1}{k} I\{T_{ij}(a, k) \leq t\} \alpha^k (1-\alpha)^{n_i-k-1}.$$

The true value of $\mu(t, \alpha)$ was determined in a similar fashion.

Results from the simulation study are presented in Table 1. Bias of the IPCW estimator was negligible for all values of a and α . Likewise, the median estimated standard error was close to the empirical standard error. Coverage of the 95% Wald CIs generally approximated

the nominal level. However, the results in Table 1 and from other simulations described below indicate coverage may differ from the nominal level when α is close to 0 or 1. For the observed data generating distribution in the simulations, the probability an individual received treatment was approximately 0.6. This suggests caution should be used when drawing inferences for values of α far from the observed proportion of individuals treated.

Additional simulation studies were conducted to assess the performance of the proposed methods for different values of m , the total number of groups, ranging from 10 to 500. The number of individuals per group was 10, as in the previous simulations. For each $m \in \{10, 50, 100, 200, 300, 400, 500\}$, 1000 data sets were simulated according to steps i through vii above. Results are depicted in Figure 1. Bias of the IPCW estimator was small and coverage of the Wald CIs was close to the nominal level, especially as the number of groups increased. Details regarding these simulation results are provided in Web Tables 1-6.

In cluster randomized trials with small numbers of clusters, Wald CIs are often constructed using a t distribution with $m - r$ degrees of freedom, where r is the number of parameters being estimated, rather than a Normal distribution; empirical coverage of Wald CIs using the t distribution is also shown in Table 1 and Web Tables 1-6, and was similar to coverage based on the Normal distribution, except when $m = 10$ in which case the t -based CI has coverage closer to the nominal level.

Additional simulations for $m = 500$ and $n_j \in \{30, 50, 75, 100, 200\}$ are provided in Web Tables 7-11. Bias of the IPCW estimator remained negligible in these scenarios, however Wald CIs tended to undercover as the size of the groups increased. Simulations were also conducted for a range of censoring rates; see Web Tables 12-14. These tables also include performance of the TV IPW estimator, which does not account for censoring; as expected the estimator is biased and the corresponding CIs fail to cover at the nominal level, with performance worsening as censoring increases.

Validity of the IPCW estimator relies on correct specification of the treatment allocation and censoring models. Therefore additional simulations were conducted to assess robustness to mis-specification of these models. Web Tables 15-22 summarize performance of the IPCW estimator to incorrect specification of the random effect distribution in the treatment model. For these simulations, the random effect b_j was generated from the distributions in Web Figure 1 but assumed to be Gaussian when calculating the IPCW estimator. These results demonstrate robustness to mild departures from the assumed random effect distribution, whereas more severe mis-specification (e.g., for skewed or bimodal distributions) may result in bias and CI undercoverage. Web Table 23 demonstrates robustness of the method when the censoring model is mis-specified due to covariate omission. Finally Web Table 24 summarizes performance of the IPCW estimator when censoring depends on treatment; as expected, the results are similar to Table 1, demonstrating the estimator performs well in this scenario.

In summary, these empirical results suggest the proposed methods: (i) appropriately account for right censoring, unlike the TV estimator; (ii) are best suited for applications where there

are many groups and the number of individuals per group is not large; and (iii) tend to be robust to mild mis-specification of the treatment allocation and censoring models.

4. Application

4.1 Cholera Vaccine Study and Analysis

In this section, the methods described in Section 2 are used to analyze a cholera vaccine study in Matlab, Bangladesh (Ali et al., 2005). Eligible study participants were children 2–15 years of age and women greater than 15 years old. All 121,975 eligible individuals in the population were randomized to one of three vaccination groups: B subunit-killed whole cell oral cholera vaccine, killed whole-cell-only cholera vaccine, and *E. coli* K12 placebo. As in Perez-Heydrich et al. (2014), no distinction is made between the two vaccines in the analysis presented here. Individuals were considered to have participated in the randomized trial component of the study if they received two or more doses of vaccine or placebo. The primary endpoint of the trial was incident cholera. Three health centers in the Matlab area served as surveillance centers and collected endpoint data on all individuals, regardless of whether they participated in the randomized trial. The analysis presented here includes data from all individuals, i.e., trial participants as well as those who chose not to participate. Thus an approach which accounts for possible confounding, such as the IPW method described in Section 2, should be utilized to assess the effects of vaccination.

Previous analyses of this study suggest the presence of interference (Ali et al., 2005; Perez-Heydrich et al., 2014). Interference is plausible in this setting because the vaccine may (i) prevent an individual from contracting cholera or (ii) decrease the infectiousness or contagiousness of an individual who does contract cholera; for either (i) or (ii), the vaccine would make it less likely that such an individual would subsequently infect other individuals. However, these previous analyses did not formally account for censoring. Here individuals are considered right censored if they were not diagnosed with cholera during the study. Individuals who emigrated from the study location or died during the follow-up period prior to cholera infection were right censored at the time of emigration or death. Individuals who did not emigrate or die and who did not develop cholera during the study were right censored at the end of the study period.

Related individuals in Matlab live in collections of houses called baris. There were a total of 6,415 baris at the time of the vaccine trial. Perez-Heydrich et al. (2014) used a clustering algorithm to form groups (neighborhoods) based on the spatial location of the baris, with the number of groups pre-specified to be 700. The analysis here is based on the same groups as in Perez-Heydrich et al. and assumes that there is no interference between individuals in different groups, i.e, vaccination of an individual in one group has no effect on whether an individual in another group contracts cholera. When fitting the propensity model $\Pr(\mathbf{A}_i | \mathbf{L}_i; \boldsymbol{\beta})$ described below, the largest 15 groups had estimated group propensity scores that were effectively equal to zero, and therefore these groups were omitted.

Individuals participating in the vaccine trial were not all vaccinated on the same calendar day, such that the level of vaccine coverage within a group varied over a relatively brief period of calendar time at the study onset. For simplicity and because the methods developed

above do not accommodate time varying treatment, the start of follow-up for all individuals in a particular group was set to the latest date of second vaccination among all individuals in that group. Observations were excluded if individuals contracted cholera (60), died (346), or emigrated (3671) prior to the start of follow-up for their group. In total, 94,234 individuals were included in the analysis. Among these individuals, 55,413 were unvaccinated, either because they received placebo or they did not participate, and 38,821 were vaccinated with one of the two vaccines. During follow-up, there were 280 incident cases of cholera among the unvaccinated individuals and 74 cholera cases among the vaccinated individuals. Most (96%) of censoring was due to end of study follow-up ($n = 89, 956$), 3.5% of censoring was due to emigration ($n = 3, 340$), and $< 1\%$ due to death ($n = 584$).

As in Perez-Heydrich et al., the group propensity score was modeled using a mixed effects model. The particular form of the model derives from the fact that for an individual to have received a vaccine, they must have (i) chosen to participate in the trial, and (ii) been randomized to receive one of the two vaccines. To account for (i), a logistic regression model for participation was assumed. Covariates in the participation component of the model were age, squared age, distance to nearest river, and squared distance to nearest river. Accommodating (ii) in the propensity model is straightforward because, due to randomization, individuals who elected to participate in the trial were known to receive one of the two vaccines with probability $2/3$. Combining these two aspects of the model, the propensity score for group i was estimated by

$$\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}}) = \int \prod_{j=1}^{n_i} \{(2/3)h_{ij}(b_i, \mathbf{L}_{ij}, \hat{\boldsymbol{\theta}}_x)\}^{A_{ij}} \{1 - (2/3)h_{ij}(b_i, \mathbf{L}_{ij}, \hat{\boldsymbol{\theta}}_x)\}^{(1 - A_{ij})} f_b(b_i; \hat{\boldsymbol{\theta}}_s),$$

where $h_{ij}(b_i, \mathbf{L}_{ij}, \boldsymbol{\theta}_x) = \Pr(B_{ij} = 1 | b_i, \mathbf{L}_{ij}, \boldsymbol{\theta}_x) = \text{expit}(\mathbf{L}_{ij}\boldsymbol{\theta}_x + b_i)$, B_{ij} is the indicator of participation, i.e., $B_{ij} = 1$ if individual j in group i participated in the randomized trial and $B_{ij} = 0$ otherwise, and $(\hat{\boldsymbol{\theta}}_x, \hat{\boldsymbol{\theta}}_s)$ is the maximum likelihood estimate of $(\boldsymbol{\theta}_x, \boldsymbol{\theta}_s)$. Censoring was modeled using the Gamma frailty model described above, and included only age as a covariate as no other variables were associated with censoring. Estimated parameters and standard errors for the treatment and censoring models are given in Web Table 25. Over 70% of individuals belonged to groups where the vaccine coverage was between 0.3 and 0.6. Therefore, the analysis was conducted for allocation strategies ranging from 0.3 to 0.6.

4.2 Results

Figure 2 shows the IPCW estimates of the cumulative probability of cholera over time for allocation strategies 0.3, 0.45, and 0.6, both when an individual receives a vaccine and when an individual is unvaccinated. The IPCW estimates can be viewed as weighted averages, with uncensored individuals weighted by $w_{ij} = \pi(\mathbf{A}_i, -j; \alpha) / \{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}}) S_C(X_{ij} | \mathbf{L}_i, \hat{\boldsymbol{\gamma}})\}$ and $\tilde{w}_{ij} = \pi(\mathbf{A}_i; \alpha) / \{\Pr(\mathbf{A}_i | \mathbf{L}_i, \hat{\boldsymbol{\beta}}) S_C(X_{ij} | \mathbf{L}_i, \hat{\boldsymbol{\gamma}})\}$; histograms of w_{ij} and \tilde{w}_{ij} for $\alpha = 0.30, 0.45, 0.60$ are presented in Web Figures 2-3. Figure 2 shows that the estimated risk of cholera when an individual is unvaccinated decreases dramatically as α increases, suggesting the presence of interference. This decrease is more modest when an individual is vaccinated, indicating a stronger indirect effect when unvaccinated. At all time points the estimated

risk of cholera is higher when an individual is unvaccinated, suggesting a beneficial, direct effect of vaccination, especially at lower coverage levels. For $\alpha = 0.3$ and $\alpha = 0.45$, the estimated risk when unvaccinated increases suddenly between 200 and 300 days, and then again between 300 and 400 days. These results might be attributable to the known bimodal seasonality of cholera in Bangladesh (Longini et al., 2002). Note that, because the study start date varied across groups, the time scale in this analysis does not exactly coincide with calendar time. Nonetheless, 95% of individuals had a start date within a two calendar month range, such that there is a strong correlation between the analysis time scale and calendar time, and thus cholera seasonality may explain these periods of marked increase in risk.

Direct, indirect, total and overall effect estimates and 95% CIs ($\times 1000$) for different allocation strategies at time $t = 1$ year are shown in Figure 3. The direct effect estimates generally decrease as α increases. For example, the direct effect estimate for $\alpha = 0.35$ is 3.6 (95% CI 1.1, 6.2) whereas for $\alpha = 0.5$ the direct effect estimate is 1.5 (95% CI -0.5 , 3.5). The indirect, total, and overall effect estimates in Figure 3 compare the risk of cholera over a range of allocation probabilities $\alpha_1 \in [0.3, 0.6]$ versus $\alpha_2 = 0.4$. Here the indirect effect contrasts risk of cholera infection when individuals are unvaccinated. For larger values of α_1 the 95% CIs for these effects exclude the null value of zero. For example, for $\alpha_1 = 0.6$ the indirect effect estimate is 2.8 (95% CI 1.1, 4.5), providing statistical evidence of the presence of interference. These results indicate that when individuals are unvaccinated, the risk of cholera infection is significantly reduced by increasing the level of vaccine coverage in their neighborhood. The total effect estimates quantify the combined direct and indirect effects of the vaccine. The overall effect estimates may be of greatest interest from a public health or policy perspective. For $\alpha_1 = 0.6$, the overall effect estimate is 2.2 (95% CI 0.9, 3.4); in words, fewer cases of cholera per 1000 individuals per year are expected if 60% of individuals are vaccinated compared to if only 40% of individuals receive vaccine.

In previous analyses of these data, Perez-Heydrich et al. also estimated the direct, indirect, total and overall effects using a binary outcome indicating whether an individual was infected with cholera during the first year of follow-up. The IPCW estimates for $t = 1$ are similar to these previous results, e.g., Perez-Heydrich et al. estimated the direct effect for $\alpha = 0.32$ to be 5.3 (95% CI 2.5, 8.1) whereas the IPCW estimate of this effect at $t = 1$ is 4.0 (95% CI 1.6, 6.5). However, the Perez-Heydrich et al. estimates may be biased because they did not account for right censoring. That said, the IPCW results should still be viewed cautiously and only have valid causal interpretation if Assumptions I - IV hold and the treatment and censoring models are correctly specified.

5. Discussion

In this paper, the TV IPW estimator for partial interference was extended to allow for right censored outcomes. The proposed estimator was obtained by weighting the original TV estimator by censoring weights calculated from a parametric frailty model of the censoring times. The estimator was shown to be consistent and asymptotically Normal (under identifiability Assumptions I - IV), and a consistent estimator of the asymptotic variance was proposed. A simulation study demonstrated that the proposed methods performed well in finite samples provided the number of groups is sufficiently large. Analysis of a cholera

vaccine study using the proposed methods suggests vaccination had both a direct and indirect effect against cholera infection. These results are in accordance with findings by Ali et al. (2005) and Perez-Heydrich et al. (2014), but are likely more accurate since these previous analyses did not formally account for right censoring.

There are several areas of possible future research related to the methods developed here. For example, further research could entail developing estimators which perform well in settings where the number of groups (m) is small. The IPCW estimator considered here also presents computational challenges when group sizes (n_i 's) are large because the corresponding estimated group propensity scores can be approximately zero. As a result, in the cholera vaccine analysis in Section 4 the largest 15 groups were excluded, inefficiently discarding data and limiting the generalizability of the results. Other approaches are needed to better accommodate larger groups. Extensions of the IPCW estimator which allow for competing risks could also be considered. In the cholera example, such extensions would permit death to be considered a competing risk rather than treating such observations as censored. Validity of the IPCW estimator requires correct specification of parametric models, such that it is important to assess model fit in application of this method. Further research could entail developing estimators which are robust to model mis-specification; e.g., in the absence of censoring, doubly robust estimators have been developed which allow for partial interference and are consistent even if the treatment model is mis-specified (Liu et al., 2019).

Extensions of the IPCW estimator could also be considered for the setting where there is general interference, i.e., where interference is not restricted to individuals within the same group. One such approach might entail extending existing methods for observational studies with arbitrary interference (e.g., Sofrygin and van der Laan (2016); Tchetgen Tchetgen et al. (2020)) to allow for right censored outcomes, perhaps by incorporating inverse probability of censoring weights. In the cholera study, interference may have been present between individuals in different groups; for instance, Root et al. (2011) found that the risk of cholera among placebo recipients was inversely associated with level of vaccine coverage in their social networks, suggesting possible interference between socially connected individuals from different groups. Methods for observational studies with general interference and right censored outcomes could also be helpful when there is partial interference but the number of groups is small.

Methods could also be developed which allow the assumed interference structure to change over time. For example, when individuals in the cholera study were censored due to emigration out of the study area, it may be considered unlikely that those individuals would continue to interfere with others in the study; in which case, the assumed interference structure could be updated at the time when the individuals were censored.

Finally, following Tchetgen Tchetgen and VanderWeele (2012) and Perez-Heydrich et al. (2014), here we consider causal estimands corresponding to counterfactual scenarios where individuals independently select treatment with probability α . In future research alternative estimands, as in Papadogeorgou et al. (2019) and Barkley et al. (2020), could be considered which target counterfactual estimands that allow for within group treatment selection

dependence. For example, if a particular vaccine is made widely available in Matlab, we might expect vaccine uptake to be positively correlated between individuals in the same bari. Thus such estimands may be more relevant to public health officials determining vaccine policy.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgment

The authors thank the Editor, Associate Editor, two reviewers, Brian Barkley, Bradley Saul, Shaina Alexandria and Kayla Kilpatrick for their comments. This work was supported by NIH grant R01 AI085073 and NSF GRFP grant DGE-1650116. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

Data Availability Statement

The data that support the findings in this paper are available on request from the International Centre for Diarrhoeal Disease Research, Bangladesh. The data are not publicly available due to privacy or ethical restrictions.

References

- Ali M, Emch M, von Seidlein L, Yunus M, Sack DA, Rao M, Holmgren J, and Clemens JD (2005). Herd immunity conferred by killed oral cholera vaccines in Bangladesh: a reanalysis. *Lancet*, 366(9479):44–49. [PubMed: 15993232]
- Baird S, Bohren J, McIntosh C, and Özler B (2018). Optimal design of experiments in the presence of interference. *Review of Economics and Statistics*, 100(5):844–860.
- Barkley BG, Hudgens MG, Clemens JD, Ali M, and Emch ME (2020). Causal inference from observational studies with clustered interference, with application to a cholera vaccine study. *Annals of Applied Statistics*, 14:1432–1448.
- Hudgens MG and Halloran ME (2008). Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842. [PubMed: 19081744]
- Liu L, Hudgens M, and Becker-Dreps S (2016). On inverse probability-weighted estimators in the presence of interference. *Biometrika*, 103(4):829–842. [PubMed: 29422692]
- Liu L, Hudgens MG, Saul B, Clemens JD, Ali M, and Emch ME (2019). Doubly robust estimation in observational studies with partial interference. *Stat*, 8(1):e214. [PubMed: 31440374]
- Loh WW, Hudgens MG, Clemens JD, Ali M, and Emch ME (2020). Randomization inference with general interference and censoring. *Biometrics*, 76(1):235–245. [PubMed: 31388990]
- Longini IM, Yunus M, Zaman K, Siddique A, Sack RB, and Nizam A (2002). Epidemic and endemic cholera trends over a 33-year period in Bangladesh. *Journal of Infectious Diseases*, 186(2):246–251. [PubMed: 12134262]
- Munda M, Rotolo F, and Legrand C (2012). parfm: Parametric frailty models in R. *Journal of Statistical Software*, 51(11):1–20. [PubMed: 23504300]
- Papadogeorgou G, Mealli F, and Zigler CM (2019). Causal inference with interfering units for cluster and population level treatment allocation programs. *Biometrics*, 75(3):778–787. [PubMed: 30859545]
- Perez-Heydrich C, Hudgens MG, Halloran ME, Clemens JD, Ali M, and Emch ME (2014). Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70(3):731–741. [PubMed: 24845800]

- Robins JM and Finkelstein DM (2000). Correcting for noncompliance and dependent censoring in an AIDS clinical trial with inverse probability of censoring weighted (IPCW) log-rank tests. *Biometrics*, 56(3):779–788. [PubMed: 10985216]
- Root ED, Giebultowicz S, Ali M, Yunus M, and Emch M (2011). The role of vaccine coverage within social networks in cholera vaccine efficacy. *PLoS One*, 6(7):e22971. [PubMed: 21829566]
- Rosenbaum PR (2007). Interference between units in randomized experiments. *Journal of the American Statistical Association*, 102(477):191–200.
- Rotnitzky A, Farall A, Bergesio A, and Scharfstein D (2007). Analysis of failure time data under competing censoring mechanisms. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(3):307–327.
- Saul BC and Hudgens MG (2020). The calculus of M-estimation in R with geex. *Journal of Statistical Software*, 92. doi: 10.18637/jss.v092.i02.
- Sobel ME (2006). 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.
- Sofrygin O and van der Laan MJ (2016). Semi-parametric estimation and inference for the mean outcome of the single time-point intervention in a causally connected population. *Journal of Causal Inference*, 5(1).
- Stefanski LA and Boos DD (2002). The calculus of M-estimation. *The American Statistician*, 56(1):29–38.
- Tchetgen Tchetgen EJ, Fulcher IR, and Shpitser I (2020). Auto-g-computation of causal effects on a network. *Journal of the American Statistical Association*, In press.
- Tchetgen Tchetgen EJ and VanderWeele TJ (2012). On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75. [PubMed: 21068053]

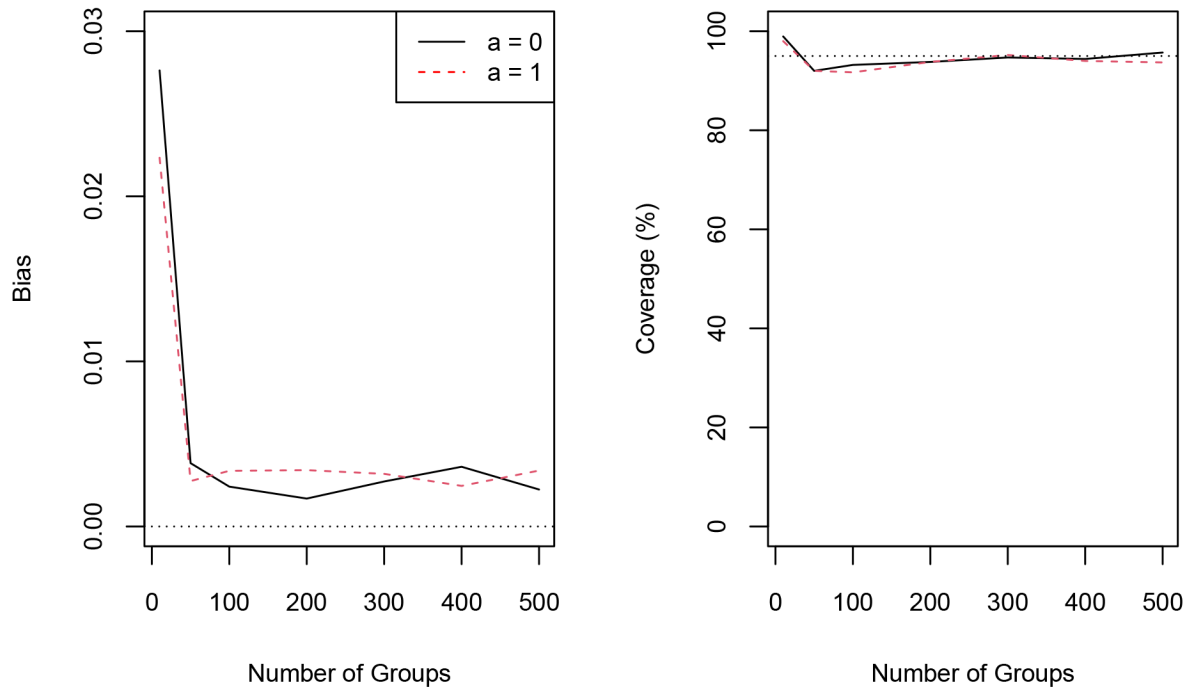


Figure 1.

Absolute bias of the IPCW estimator (left) and corresponding Wald 95% confidence interval coverage (right) for different numbers of groups for $\alpha = 0.5$. The dotted line in the right plot corresponds to 95% coverage. This figure appears in color in the electronic version of this article.




Figure 2.

Estimated cumulative probability of cholera over time when vaccinated or unvaccinated for $\alpha = 0.3$ (left), $\alpha = 0.45$ (center) and $\alpha = 0.6$ (right). This figure appears in color in the electronic version of this article.

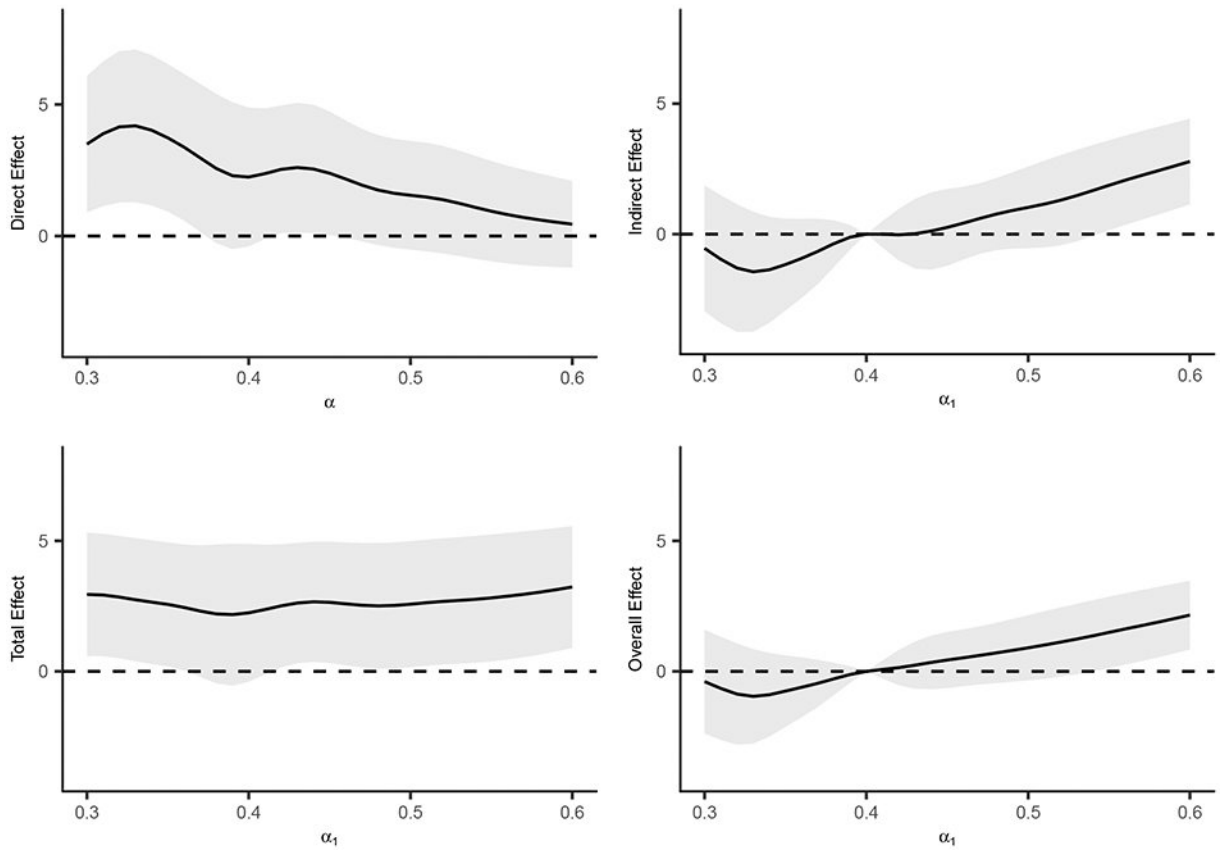


Figure 3. Direct, indirect, total and overall effect estimates ($\times 1000$) for different allocation strategies at time $t = 1$ year. Indirect, total, and overall effects are with respect to $a_2 = 0.4$. The shaded regions denote pointwise 95% confidence intervals of the estimates.

Table 1

Results from simulation study described in Section 3, where: α denotes the allocation probability; μ_a denotes the value of the target parameters $\mu(100, a, \alpha)$ for $a = 0, 1$; Bias is the average of $\mu(100, a, \alpha) - \hat{\mu}(100, a, \alpha)$ for $a = 0, 1$; ESE is the empirical standard error; MSE is the median of the sandwich variance-based standard error estimates; EC denotes the empirical coverage of the 95% Wald confidence intervals based on the Normal distribution; and EC_t denotes empirical coverage of t -distribution-based Wald confidence intervals.

| α | μ_0 | Bias | ESE | MSE | EC | EC_t | α | μ_1 | Bias | ESE | MSE | EC | EC_t |
|----------|---------|-------|------|------|-----|--------|----------|---------|-------|------|------|-----|--------|
| 0.1 | 0.39 | -0.02 | 0.07 | 0.07 | 95% | 95% | 0.1 | 0.28 | -0.01 | 0.08 | 0.08 | 92% | 92% |
| 0.2 | 0.38 | -0.01 | 0.04 | 0.04 | 95% | 95% | 0.2 | 0.28 | 0.00 | 0.04 | 0.04 | 95% | 95% |
| 0.3 | 0.38 | -0.01 | 0.03 | 0.03 | 95% | 95% | 0.3 | 0.27 | 0.00 | 0.03 | 0.03 | 94% | 94% |
| 0.4 | 0.37 | 0.00 | 0.03 | 0.02 | 95% | 95% | 0.4 | 0.27 | 0.00 | 0.02 | 0.02 | 94% | 94% |
| 0.5 | 0.37 | 0.00 | 0.03 | 0.02 | 96% | 96% | 0.5 | 0.27 | 0.00 | 0.02 | 0.02 | 94% | 94% |
| 0.6 | 0.36 | 0.00 | 0.02 | 0.02 | 94% | 95% | 0.6 | 0.27 | 0.00 | 0.02 | 0.02 | 94% | 94% |
| 0.7 | 0.35 | 0.01 | 0.02 | 0.02 | 95% | 95% | 0.7 | 0.26 | 0.00 | 0.01 | 0.01 | 96% | 96% |
| 0.8 | 0.35 | 0.00 | 0.03 | 0.03 | 96% | 96% | 0.8 | 0.26 | 0.00 | 0.02 | 0.02 | 96% | 96% |
| 0.9 | 0.34 | 0.00 | 0.05 | 0.05 | 95% | 95% | 0.9 | 0.26 | 0.00 | 0.02 | 0.02 | 96% | 96% |