Toward Causal Inference With Interference

Michael G. Hudgens and **M. Elizabeth Halloran**
Michael G. Hudgens is Research Associate Professor, Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599 (E-mail: mhudgens@bios.unc.edu). M. Elizabeth Halloran is Professor, Program in Biostatistics and Biomathematics, Fred Hutchinson Cancer Research Center, and Department of Biostatistics, University of Washington, Seattle, WA 98185 (E-mail: betz@u.washington.edu)

Abstract

A fundamental assumption usually made in causal inference is that of no interference between individuals (or units); that is, the potential outcomes of one individual are assumed to be unaffected by the treatment assignment of other individuals. However, in many settings, this assumption obviously does not hold. For example, in the dependent happenings of infectious diseases, whether one person becomes infected depends on who else in the population is vaccinated. In this article, we consider a population of groups of individuals where interference is possible between individuals within the same group. We propose estimands for direct, indirect, total, and overall causal effects of treatment strategies in this setting. Relations among the estimands are established; for example, the total causal effect is shown to equal the sum of direct and indirect causal effects. Using an experimental design with a two-stage randomization procedure (first at the group level, then at the individual level within groups), unbiased estimators of the proposed estimands are presented. Variances of the estimators are also developed. The methodology is illustrated in two different settings where interference is likely: assessing causal effects of housing vouchers and of vaccines.

Keywords

Group-randomized trials; Potential outcomes; Stable unit treatment value assumption; SUTVA; Vaccine

1. INTRODUCTION

1.1 Background and Outline

A fundamental assumption usually made in the potential outcomes approach to causal inference is that of no interference between individuals (Cox 1958), a critical component of the stable unit treatment value assumption (SUTVA) (Rubin 1980). Under the no-interference assumption, the potential outcomes of any individual are assumed to be unaffected by the treatment assignment of every other individual. However, in many settings, this assumption obviously does not hold. A classical example is given by the dependent happenings of infectious diseases (Ross 1916, p. 211), where whether one person becomes infected depends on who else in the population is vaccinated. In econometrics, a household's decision whether to move may be affected by whether their neighbors receive a housing voucher to move (Sobel 2006). In education, interventions given to certain students may affect other students in the same class (Rubin 1990;Rosenbaum 2007). Sobel (2006) and Rosenbaum (2007) gave several other examples where interference is likely. In some settings, interference is a nuisance while in other settings it creates effects of interest. An example of the former includes agricultural experiments, where fallow rows between treatment plots can sometimes eliminate interference

between plots. An example of the latter includes vaccinating against infectious diseases, where interference is an inherent result of the biology of transmission and is intrinsically of interest.

The assumption of no interference between individuals is often made without critical examination. Models not requiring this assumption have been considered in the context of plant variety evaluation (Kempton 1997) and cross-over trials (Senn 1993; Bailey and Kunert 2006). However, these methods typically assume a specific interference structure that is local in either space or time. Without making any such assumptions about the nature of interference, Struchiner, Halloran, Robins, and Spielman (1990) and Halloran and Struchiner (1991) conceptually defined several different types of causal effects of interventions that are possible in the presence of interference, namely, direct, indirect, total, and overall effects. To estimate the latter three effects, they noted one needs a population of groups as in group-randomized studies (Murray 1998). Several vaccination studies have been conducted or analyzed with the intent to estimate certain of these effects (Moulton et al. 2001; Longini, Halloran, and Nizam 2002; Ali et al. 2005; King et al. 2006).

Halloran and Struchiner (1995) delineated many of the complications of using potential outcomes to define causal estimands for the different types of effects possible in the presence of interference. They used Rubin's (1978, 1990) suggestion for a general notation in the presence of interference to define individual direct, indirect, total, and overall effects by letting the potential outcomes for any individual depend on the vector of treatment assignments to the other individuals in the population. However, they found this approach impracticable because the number of possible potential outcomes becomes unwieldy for any reasonably sized population. More recently, Sobel (2006) proposed causal estimands for assessing housing voucher effects defined by averaging causal effects over all possible treatment assignments for a particular voucher allocation strategy compared to a benchmark allocation wherein all households receive no voucher. Rosenbaum (2007) developed nonparametric tests and confidence intervals for assessing treatment effect in the presence of interference.

In this article, we consider a population of groups of individuals where interference is possible between individuals within the same group. We propose causal estimands for direct, indirect, total, and overall causal effects of treatment assignment strategies based on Sobel's approach of averaging over all possible treatment assignments (Sec. 3). Relations among the estimands are established and inference concerning the estimands is considered (Sec. 4). Using an experimental design with a two-stage randomization procedure (the first at the group level, the second at the individual level within groups), unbiased estimators of the proposed estimands are presented. Estimating the variance of the estimators is also considered. The methodology is illustrated in two different settings where interference is likely: assessing causal effects of housing vouchers and of vaccines (Sec. 5). Proofs are given in the Appendix. We begin with an example to motivate the development of the rest of the article.

1.2 Motivating Example

In this section, we consider data from an individually randomized, placebo-controlled trial of killed oral cholera vaccines to illustrate the direct, indirect, total, and overall effects as defined by Halloran and Struchiner (1991). Table 1 presents data from a reanalysis of this trial where the interest was in determining whether the level of vaccine coverage in a residential area, called a bari, was related to the incidence of cholera in individual vaccine recipients or placebo recipients residing in the bari (Ali et al. 2005). The target population was divided into groups by level of vaccine coverage. For illustration, we consider the groups with more than 50% and less than 28% coverage, which we denote as groups A and B.

The effects of vaccination can be estimated based on differences in the incidence of cholera during the first year of follow-up of the trial. The direct effects are estimated by comparing the

incidence (risk per 1,000 population) between vaccinated individuals and unvaccinated individuals within each group. For example, the estimated direct effect in group B is 7.01-2.66 = 4.35, suggesting vaccination results in 4.35 fewer cases of cholera per 1,000 individuals per year. The estimated direct effect in group A is 1.47-1.27 = .20, considerably lower than in group B. The difference in the two estimates illustrates one of the challenges in making comparisons directly within groups when interference is present. If an analysis were limited to group A only, the evidence would suggest that the vaccine has little effect.

The indirect effects of vaccination are those effects due to the level of coverage. They can be estimated by comparing the outcomes in the unvaccinated in the two groups or the outcomes in the vaccinated in the two groups. For instance, the estimated indirect effect in the unvaccinated is 7.01 - 1.47 = 5.54. Note this estimate is greater than the estimated direct effect in either of the groups, highlighting the importance of looking beyond direct effects in the presence of interference. Based on similar analyses, Ali et al. concluded that the vaccines provide significant indirect protection to nonvaccinated individuals.

Total and overall effects provide summary measures that combine direct and indirect effects. The total effect of vaccination is the effect of being vaccinated in the group with higher coverage (A) compared to not being vaccinated in the group with lower coverage (B). The estimated total effect (B - A) is 7.01-1.27 = 5.74. Note the total effect (B - A) estimate equals the direct effect estimate in group A plus the indirect effect estimate in the unvaccinated (B - A). The overall effect is the average effect of being in the group with higher coverage compared to being in the group with lower coverage. The overall effect can be estimated by the difference in incidence between the two groups, that is, 35/8,479 - 25/18,623 = 2.79/1,000.

2. PRELIMINARIES

2.1 Potential Outcomes

Suppose there are $N > 1$ groups of individuals [or blocks of units using Rosenbaum's (2007) terminology]. For $i = 1, \ldots, N$, let n_i denote the number of individuals in group i and let $\mathbf{Z}_i \equiv (Z_{i1}, \ldots, Z_{in_i})$ denote the treatments those n_i individuals receive. We assume throughout that assignment of an individual to a particular treatment is equivalent to receipt of that treatment that is, there is perfect compliance. Assume Z_{ij} is a dichotomous random variable having values 0 or 1 such that $\mathbf{Z}_i$ can take on 2^{n_i} possible values. Let $\mathbf{Z}_{i(j)}$ denote the n_i - 1 subvector of $\mathbf{Z}_i$ with the j th entry deleted. The vector $\mathbf{Z}_i$ will be referred to as an intervention or treatment *program*, to distinguish it from the individual treatment Z_{ij}. Let $\mathbf{z}_i$ and z_{ij} denote possible values of $\mathbf{Z}_i$ and $\mathbf{Z}_i$ and Z_{ij}. Define R^j to be the set of vectors of possible treatment programs of length j for $j = 1, 2, \ldots$; for example, $R^2 \equiv \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. Possible values $\mathbf{z}_i$ of $\mathbf{Z}_i$ are elements of R^{n_i}. For positive integer n and $k \in \{0, \ldots, n\}$, define R_k^n to be the subset of R^n

wherein exactly k individuals receive treatment 1; for example, $\sum_{j=1}^{n_i} z_{ij} = k$ for all $\mathbf{z}_i \in R_k^{n_i}$.

Denote the potential outcome of individual j in group i under treatment $\mathbf{z}_i$ as $Y_{ij}(\mathbf{z}_i)$. Following the usual approach to causal inference (see, e.g., Rosenbaum 2007), we assume the $Y_{ij}(\mathbf{z}_i)$ potential responses are fixed because they do not depend on the realized random assignment of treatments $\mathbf{Z}_i$, whereas the observed responses $Y_{ij}(\mathbf{Z}_i)$ do depend on $\mathbf{Z}_i$ and, thus, are random variables. The notation $Y_{ij}(\mathbf{z}_i)$ allows for the possibility that the potential outcome for individual j may depend on another individual's treatment assignment in group i; that is, there may be interference between individuals within a group. Implicit in this notation is the assumption that the potential outcomes for individuals in group i do not depend on treatment assignments of individuals in group i' for i' $\neq i$. In other words, we assume no interference between individuals in different groups but allow for interference between individuals within the same group (Halloran and Struchiner 1991, 1995). This will be a reasonable assumption provided the

groups are sufficiently separate (e.g., in space or time). Sobel (2006) called this a partial interference assumption. In the literature of group-randomized studies, violation of no interference across groups is called contamination.

2.2 Treatment Assignment Mechanisms

Let ψ and ϕ denote parameterizations that govern the distribution of $\mathbf{Z}_i$ for $i = 1, \ldots, N$. For example, ψ might correspond to randomly assigning half of individuals in a group to treatment 1 and the other half to treatment 0, while ϕ might correspond to assigning all individuals in a group to treatment 0. We refer to ψ and ϕ as individual treatment assignment strategies. Our goal is to assess the causal effects of assigning groups to ψ compared to ϕ.

As is typical of causal inference articles, we use randomization inference whereby the randomization distribution induced by the experimental design forms the basis for statistical inference. For the experimental design, we consider a two-stage randomization procedure. In the first stage, each of the N groups is randomly assigned to either ϕ or ψ. In the second stage, individuals are randomly assigned treatment conditional on their group's assignment in the first stage. For example, in the first stage, half of the N groups might be assigned to an allocation strategy ϕ and the other half ψ; in the second stage, two-thirds of the individuals within a group are randomly assigned treatment 1 for groups assigned ϕ, while one-third of the individuals within a group are randomly assigned treatment 1 for groups assigned ψ. Such a design has been referred to as splitplot (Hayes, Alexander, Bennett, and Cousens 2000) or pseudo-cluster (Borm, Melis, Teerenstra, and Peer 2005) randomization and has been proposed for evaluation of intervention programs in the elderly (Melis et al. 2005) and vaccine efficacy (see Sec. 5.2). This design can be employed to answer questions such as: How many infections will be averted by vaccinating two-thirds of the population compared to only vaccinating one-third of the population? What proportion of households will move if two-thirds receive vouchers compared to only one-third receiving vouchers?

Corresponding to the first stage of randomization, let $\mathbf{S} \equiv (S_1, \ldots, S_N)$ denote the group assignments with $S_i = 1$ if the ith group is assigned to ψ and 0 otherwise. Let v denote the parameterization that governs the distribution of $\mathbf{S}$ and let $C \equiv \Sigma_i S_i$ denote the number of groups assigned ψ. Define v to be a *mixed* (Sobel 2006) or *permutation* (Friedman, Furberg, and DeMets 1998) group assignment strategy if $0 < C < N$ and $\Pr_v (\mathbf{S} = \mathbf{s}) = C!(N - C)!/N!$ if $\mathbf{s} \in R_C^N$, 0 otherwise. In other words, under a mixed group assignment strategy, a fixed number C of N groups are assigned ψ, with each of the $\binom{N}{C}$ possible group assignments receiving equal probability. Similarly, corresponding to the second stage of randomization, let $K_i \equiv \Sigma_j Z_{ij}$ and define ϕ and ψ to be mixed individual group assignment strategies if K_i is fixed given S_i, with $0 < K_i < n_i$ and each of the $\binom{n_i}{K_i}$ possible individual treatment assignments receiving equal probability.

3. CAUSAL ESTIMANDS

3.1 Average Potential Outcomes

A fundamental problem in causal inference is that, in general, it is not possible to observe more than one potential outcome for an individual. Faced with this problem, causal estimands are typically defined in terms of averages of potential outcomes that are identifiable from observable random variables. Following this approach, we begin by writing the potential outcomes for individual j in group i under $z_{ij} = z$ as

$$Y_{ij}\left(\mathbf{z}_{i(j)}, z_{ij}=z\right) \tag{1}$$

for $z = 0, 1$. Because (1) depends on $\mathbf{z}_{i(j)}$, define the *individual average potential outcome* under treatment assignment z by

$$\bar{Y}_{ij}(z;\psi) \equiv \sum_{\omega \in R^{n_i-1}} Y_{ij}\left(\mathbf{z}_{i(j)}=\omega, z_{ij}=z\right) \times \Pr_\psi\left(\mathbf{Z}_{i(j)}=\omega | Z_{ij}=z\right).$$

In other words, the individual average potential outcome is the conditional expectation of Y_{ij} ($\mathbf{Z}_i$) given $Z_{ij} = z$ under assignment strategy ψ. Averaging over individuals, define the *group average potential outcome* under treatment assignment z as $\bar{Y}_i(z;\psi) \equiv \sum_{j=1}^{n_i} \bar{Y}_{ij}(z;\psi)/n_i$. Finally, averaging over groups, define the *population average potential outcome* under treatment assignment z as $\bar{Y}(z;\psi) \equiv \sum_{i=1}^{N} \bar{Y}_i(z;\psi)/N$.

The average potential outcomes discussed previously are defined as functions of both the group assignment ψ (or ϕ) and the individual treatment assignment z. We can also define average potential outcomes solely as a function of ψ. For example, define the *marginal* individual average potential outcome by $\bar{Y}_{ij}(\psi) \equiv \sum_{z \in R^{n_i}} Y_{ij}(\mathbf{z}) \Pr_\psi(\mathbf{Z}_i = \mathbf{z})$, that is, the average potential outcome for individual j in group i when group i is assigned ψ. Similarly, define the marginal group and population average potential outcomes by $\bar{Y}_i(\psi) \equiv \sum_{j=1}^{n_i} \bar{Y}_{ij}(\psi)/n_i$ and $\bar{Y}(\psi) \equiv \sum_{i=1}^{N} \bar{Y}_i(\psi)/N$.

In the following sections, causal estimands are defined in terms of these various average potential outcomes.

3.2 Direct Causal Effects

Halloran and Struchiner (1991) defined the direct effect of a treatment on an individual as the difference between the potential outcome for that individual given treatment compared to the potential outcome for that individual without treatment, all other things being equal. Formally, following Halloran and Struchiner (1995), we define the *individual direct causal effect* of treatment 0 compared to treatment 1 for individual j in group i by

$$CE_{ij}^{D}\left(\mathbf{z}_{i(j)}\right) \equiv Y_{ij}\left(\mathbf{z}_{i(j)}, z_{ij}=0\right) - Y_{ij}\left(\mathbf{z}_{i(j)}, z_{ij}=1\right). \tag{2}$$

Next, define the *individual average direct causal effect* for individual j in group i by

$$\bar{CE}_{ij}^{D}(\psi) \equiv \bar{Y}_{ij}(0;\psi) - \bar{Y}_{ij}(1;\psi), \tag{3}$$

that is, the difference in individual average potential outcomes when $z_{ij} = 0$ and when $z_{ij} = 1$ under ψ. Using Rubin's (2005) terminology, (3) is a marginal causal effect in that a comparison is being made between expected values of the marginal distributions of $Y_{ij}(\mathbf{Z}_{i(j)}, Z_{ij} = 0)$ and of $Y_{ij}(\mathbf{Z}_{i(j)}, Z_{ij} = 1)$. Finally, define the *group average direct causal effect* by

$\bar{CE}_i^{D}(\psi) \equiv \bar{Y}_i(0;\psi) - \bar{Y}_i(1;\psi) = \sum_{j=1}^{n_i} \bar{CE}_{ij}^{D}(\psi)/n_i$ and the *population average direct causal effect* by $\bar{CE}^{D}(\psi) \equiv \bar{Y}(0;\psi) - \bar{Y}(1;\psi) = \sum_{i=1}^{N} \bar{CE}_i^{D}(\psi)/N$.

3.3 Indirect Causal Effects

In contrast to direct effects, an indirect effect describes the effect on an individual of the treatment received by others in the group. In particular, Halloran and Struchiner (1991) defined the indirect effect of a treatment on an individual as the difference between the potential outcomes for that individual without treatment when the group (i) receives an intervention

program and (ii) receives the benchmark program of no intervention. Similar to Halloran and Struchiner (1995), we define the *individual indirect causal effect* of treatment program $\mathbf{z}_i$ compared with $\mathbf{z}'_i$ on individual j in group i by

$$CE_{ij}^I\left(\mathbf{z}_{i(j)}, \mathbf{z}'_{i(j)}\right) \equiv Y_i\left(\mathbf{z}_{i(j)}, z_{ij}{=}0\right) - Y_i\left(\mathbf{z}'_{i(j)}, z'_{ij}{=}0\right),$$

(4)

where $\mathbf{z}'_i$ is another n_i-dimensional vector of individual treatment assignments. (Note $\mathbf{z}'_i$ does not denote the transpose of $\mathbf{z}_i$.)

Remark—Definition (4) does not restrict either $\mathbf{z}_i$ or $\mathbf{z}'_i$ to be the benchmark program of no intervention; that is, individual indirect causal effects may exist between two different intervention programs. The same is true for the definitions of individual total and overall causal effects.

Remark—The individual indirect causal effect could be defined analogously for individuals with $z_{ij}{=}z'_{ij}{=}1$; that is, individuals under either treatment may experience indirect effects. This yields two individual indirect causal effects, which need not be equal. For simplicity, only indirect effects based on (4) are considered in the rest of this article.

Similar to direct effects, define the *individual average indirect causal effect* by $\overline{CE}_{ij}^I(\phi,\psi) \equiv \overline{Y}_{ij}(0;\phi) - \overline{Y}_{ij}(0;\psi)$. Clearly, if $\psi = \phi$, then $\overline{CE}_{ij}^I(\phi,\psi){=}0$; that is, there are no individual average indirect causal effects. Finally, define the *group average indirect causal effect* as $\overline{CE}_i^I(\phi,\psi) \equiv \overline{Y}_i(0;\phi) - \overline{Y}_i(0;\psi) = \sum_{j=1}^{n_i} \overline{CE}_{ij}^I(\phi,\psi)/n_i$ and the *population average indirect causal effect* as $\overline{CE}^I(\phi,\psi) \equiv \overline{Y}(0;\phi) - \overline{Y}(0;\psi) = \sum_{i=1}^{N} \overline{CE}_i^I(\phi,\psi)/N$.

3.4 Total Causal Effects

Total effects describe both the direct and the indirect effects of a particular treatment assignment on an individual. Halloran and Struchiner (1991) defined the total effect of a treatment on an individual as the difference between the potential outcomes for that individual (i) with treatment when the group receives an intervention program and (ii) without treatment when the group receives no intervention. Following Halloran and Struchiner (1995), we define the *individual total causal effects* for individual j in group i as

$$CE_{ij}^T\left(\mathbf{z}_{i(j)}, \mathbf{z}'_{i(j)}\right) \equiv Y_{ij}\left(\mathbf{z}_{i(j)}, z_{ij}{=}0\right) - Y_{ij}\left(\mathbf{z}'_{i(j)}, z'_{ij}{=}1\right).$$

(5)

Define the *individual average total causal effect* by $\overline{CE}_{ij}^T(\phi,\psi) \equiv \overline{Y}_{ij}(0;\phi) - \overline{Y}_{ij}(1;\psi)$, the *group average total causal effect* by $\overline{CE}_i^T(\phi,\psi) \equiv \overline{Y}_i(0;\phi) - \overline{Y}_i(1;\psi) = \sum_{j=1}^{n_i} \overline{CE}_{ij}^T(\phi,\psi)/n_i$, and the *population average total causal effect* by $\overline{CE}^T(\phi,\psi) \equiv \overline{Y}(0;\phi) - \overline{Y}(1;\psi) = \sum_{i=1}^{N} \overline{CE}_i^T(\phi,\psi)/N$.

Remark—It follows from (2), (4), and (5) that the individual total causal effect is the sum of individual direct and indirect causal effects, that is, $CE_{ij}^T\left(\mathbf{z}_{i(j)}, \mathbf{z}'_{i(j)}\right){=}CE_{ij}^D\left(\mathbf{z}'_{i(j)}\right){+}CE_{ij}^I\left(\mathbf{z}_{i(j)}, \mathbf{z}'_{i(j)}\right)$. Likewise, the total causal effects can be decomposed as the sum of direct and indirect causal effects at the individual average, group average, and population average levels, for example, $\overline{CE}^T(\phi,\psi){=}\overline{CE}^D(\psi){+}\overline{CE}^I(\phi,\psi)$. This result formalizes, using a causal framework, models from the vaccine and plant variety

evaluation literature, which assume the total effect is the sum of direct and indirect effects (Halloran and Struchiner 1991,1995;Kempton 1997;Moulton et al. 2006).

Remark—A few other characteristics of the algebra of causal effects bear mentioning. First, total causal effects are not commutative; for example, $\bar{CE}^T(\phi,\psi)$ will not necessarily equal $\bar{CE}^T(\psi,\phi)$ for $\phi \neq \psi$. However, indirect effects have the property $\bar{CE}^I(\psi,\phi) = -\bar{CE}^I(\phi,\psi)$, implying $\bar{CE}^D(\psi) + \bar{CE}^D(\phi) = \bar{CE}^T(\phi,\psi) + \bar{CE}^T(\psi,\phi)$. Thus, the total causal effects, while not necessarily equal, are constrained in sum to equal the sum of the direct effects. Also note that if $\bar{CE}^I(\psi,\phi) = \bar{CE}^I(\phi,\psi) = 0$, then $\bar{CE}^T(\phi,\psi) = \bar{CE}^T(\psi,\phi)$ if and only if $\bar{CE}^D(\phi) = \bar{CE}^D(\psi)$; that is, in the absence of indirect effects, the total effects are commutative if and only if the direct effects are equal.

3.5 Overall Causal Effect

Halloran and Struchiner (1991) defined the overall causal effect to be the average effect of an intervention program relative to no intervention. We define the *individual overall causal effect* of treatment $\mathbf{z}_i$ compared to treatment $\mathbf{z}'_i$ for individual j in group i by $CE^O_{ij}(\mathbf{z}_i, \mathbf{z}'_i) \equiv Y_{ij}(\mathbf{z}_i) - Y_{ij}(\mathbf{z}'_i)$. Similarly, for the comparison of ϕ to ψ, define the *individual average overall causal effect* by $\bar{CE}^O_{ij}(\phi,\psi) \equiv \bar{Y}_{ij}(\phi) - \bar{Y}_{ij}(\psi)$, the *group average overall causal effect* by $\bar{CE}^O_i(\phi,\psi) \equiv \bar{Y}_i(\phi) - \bar{Y}_i(\psi)$, and the *population average overall causal effect* by $\bar{CE}^O(\phi,\psi) \equiv \bar{\bar{Y}}(\phi) - \bar{\bar{Y}}(\psi)$.

3.6 No Interference

The estimands defined previously simplify under the assumption of no interference between individuals within a group, that is, under the assumption $Y_{ij}(\mathbf{z}_i) = Y_{ij}(\mathbf{z}'_i)$ for any two treatment programs $\mathbf{z}_i = (z_{i1}, \ldots, z_{in_i})$ and $\mathbf{z}'_i = (z'_{i1}, \ldots, z'_{in_i})$ such that $z_{ij} = z'_{ij}$ (Rubin 1980; Angrist, Imbens, and Rubin 1996). Assuming no interference, the potential outcomes for individual j in group i can be written simply as $Y_{ij}(0)$ and $Y_{ij}(1)$. In turn, the individual direct causal effect equals $Y_{ij}(0) - Y_{ij}(1)$. The corresponding group average direct causal effect becomes $\sum_{j=1}^{n_i} \{Y_{ij}(0) - Y_{ij}(1)\}/n_i$, that is, the usual average causal effect (ACE) estimand. By (4), the individual indirect causal effect equals 0 for all individuals assuming no-interference. Similarly, by (5), the individual total causal effect equals the individual direct causal effect. Likewise, at the group and population average levels, under the no-interference assumption the indirect causal effect is 0 and the direct causal effect equals the total causal effect. Assuming no interference also implies the direct, indirect, and total effects do not depend on the treatment assignment strategies ϕ and ψ, whereas in the presence of interference within a group, they do in general.

4. INFERENCE

In this section, we consider drawing inference about the estimands defined previously. Throughout this section, we assume:

Assumption 1. ν, ϕ, and ψ are mixed assignment strategies.

In Section 4.1, we present estimators for the estimands defined previously and show they are unbiased under Assumption 1. In Section 4.2, we consider the variances of these estimators.

4.1 Estimators

Theorem 1—Suppose $S_i = 1$ and let

$$\widehat{Y}_i(z;\psi) \equiv \frac{\sum_{j=1}^{n_i} Y_{ij}(\mathbf{Z}_i) I\left[Z_{ij}=z\right]}{\sum_{j=1}^{n_i} I\left[Z_{ij}=z\right]} \quad \text{for} \quad z=0,1;$$

(6)

that is, $\widehat{Y}_i(z;\psi)$ is the average of observed outcomes for individuals in group i receiving treatment z under treatment program $\mathbf{Z}_i$. Under Assumption 1, $E\left\{\widehat{Y}_i(z;\psi)|S_i=1\right\}=\bar{Y}_i(z;\psi)$ for $z = 0, 1$.

Corollary: Under Assumption 1, $\widehat{CE}_i^D(\psi) \equiv \widehat{Y}_i(0;\psi) - \widehat{Y}_i(1;\psi)$ is a conditionally unbiased estimator of $\bar{CE}_i^D(\psi)$ given $S_i = 1$.

Remark: Unbiased estimators of the group average indirect, total, and overall causal effects do not exist without further assumptions because the same group is not observed under ϕ and ψ.

Theorem 2—For $z = 0, 1$, let $\widehat{Y}(z;\psi) \equiv \sum_{i=1}^{N}\widehat{Y}_i(z;\psi) \times I[S_i=1]/\sum_{i=1}^{N} I[S_i=1]$ Under Assumption 1, $E\left\{\widehat{Y}(z;\psi)\right\}=\bar{Y}(z;\psi)$ for $z = 0.1$.

Corollary: Under Assumption 1, unbiased estimators for the population average direct, indirect, and total causal effects are given by $\widehat{CE}^D(\psi) \equiv \widehat{Y}(0;\psi) - \widehat{Y}(1;\psi)$, $\widehat{CE}^I(\phi,\psi) \equiv \widehat{Y}(0;\phi) - \widehat{Y}(0;\psi)$, and $\widehat{CE}^T(\phi,\psi) \equiv \widehat{Y}(0;\phi) - \widehat{Y}(1;\psi)$, where $\widehat{Y}(z;\phi)$ is defined analogously to $\widehat{Y}(z;\psi)$ for $z = 0, 1$.

Theorem 3—Let $\widehat{Y}_i(\psi) \equiv \sum_{j=1}^{n_i} Y_{ij}(\mathbf{Z}_i)/n_i$ and $\widehat{Y}(\psi) \equiv \sum_{i=1}^{N}\widehat{Y}_i(\psi) I[S_i=1]/\sum_{i=1}^{N} I[S_i=1]$ Under Assumption 1, $E\left\{\widehat{Y}_i(\psi)|S_i=1\right\}=\bar{Y}_i(\psi)$ and $E\left\{\widehat{Y}(\psi)\right\}=\bar{Y}(\psi)$.

Corollary: Under Assumption 1, an unbiased estimator of $\bar{CE}^O(\phi,\psi)$ is given by $\widehat{CE}^O(\phi,\psi) \equiv \widehat{Y}(\phi) - \widehat{Y}(\psi)$, where $\widehat{Y}(\phi)$ is defined analogously to $\widehat{Y}(\psi)$.

4.2 Variance Estimators

In general, unbiased estimators of the variances of the estimators discussed previously do not exist without making further assumptions. For example, consider estimating $\text{Var}\left(\widehat{Y}_i(z;\psi)|S_i=1\right)$ under Assumption 1. The estimator $\widehat{Y}_i(1;\psi)$ is based on sampling from the set of potential outcomes $\left\{Y_{ij}(\mathbf{z}_i):\mathbf{z}_i \in R_{K_i}^{n_i}, z_{ij}=1\right\}$ for some fixed value of K_i. This set can be partitioned into $\binom{n_i}{K_i}$ clusters of size K_i, where each cluster corresponds to a particular $\mathbf{z}_i \in R_{K_i}^{n_i}$. Moreover, given $S_i = 1$ from the first stage of randomization, the second randomization

stage entails selecting exactly one of these clusters according to $\mathbf{Z}_i$. Thus, $\widehat{Y}_i(1;\psi)$ can be viewed as the sample mean from a single systematic sample. It is known that, in general, unbiased estimators of the variance of the sample mean from a single systematic sample do not exist without making further assumptions about the underlying population (Som 1973, sec. 4.4; Thompson 1992, chap. 12.4).

Therefore, to make progress in deriving variance estimators, in Section 4.2.1 an additional assumption is introduced about the structure of interference (*stratified interference*), which may be plausible in a broad range of settings. In Section 4.2.2, variance estimators of the direct, indirect, total, and overall causal effect estimators are proposed. Under the additional assumption of stratified interference, these variance estimators are shown to be unbiased if the causal effects are additive and positively biased otherwise.

4.2.1 Stratified Interference—Suppose that R^{n_i}, that is, the set of possible treatment programs for group i, can be partitioned into strata such that within strata there is no interference. In particular, we assume:

Assumption 2 (Stratified interference). For $k = 1, \ldots, n_i - 1$, $Y_{ij}(\mathbf{z}_i) = Y_{ij}(\mathbf{z}_i')$ for all $\mathbf{z}_i, \mathbf{z}_i' \in R_k^{n_i}$ such that $z_{ij} = z_{ij}'$.

To illustrate the meaning of Assumption 2, consider a study of the effects of an intervention on children in a school. The stratified interference assumption states that the outcome for a child receiving the intervention will be the same when $k - 1$ schoolmates also receive the intervention, regardless of which particular $k - 1$ schoolmates receive the intervention. This assumption can be viewed as an intermediate assumption between (i) assuming no interference within a group and (ii) making no assumptions about the nature of interference within a group. Moreover, because there are n_i possible values of K_i given $z_{ij} = z$, it follows that $Y_{ij}(z_{i(j)}, z_{ij} = z)$ can take on n_i values. Thus, for a given $z_{ij} = z$, an individual has n_i potential outcomes under Assumption 2 compared to only one potential outcome under (i) and 2^{n_i-1} potential outcomes under (ii).

To illustrate the utility of Assumption 2, again consider estimating $\mathrm{Var}\left(\widehat{Y}_i(1;\psi)|S_i=1\right)$.

Suppose, by way of contradiction, there exists an unbiased estimator $\mathrm{Var}\left(\widehat{Y}_i(1;\psi)|S_i=1\right)$ in general, that is, under Assumption 1 only. Denote this estimator by $g(O_{i1}(\mathbf{Z}_i))$, where $O_{i1}(\mathbf{Z}_i) \equiv \{Y_{ij}(\mathbf{Z}_i) : Z_{ij} = 1\}$ is the set of observed outcomes for individuals in group i assigned treatment $Z_{ij} = 1$ and g is some real-valued function of $O_i(\mathbf{Z}_i)$. By our supposition, $E\{g(O_{i1}(\mathbf{Z}_i))|S_i=1\} = \mathrm{Var}\left(\widehat{Y}_i(1;\psi)|S_i=1\right)$ under Assumption 1. Consider the following two different scenarios where $n_i = 3$ and $K_i = 2$ given $S_i = 1$.

First, suppose $Y_{ij}(\mathbf{z}_i) = k$ for $j = 1, 2, 3$, and $\mathbf{z}_i \in R_2^3 = \{(011),(110),(101)\}$, where κ is some constant. In other words, the potential outcomes in group i are constant for all individuals and all treatment programs given $S_i = 1$. Then $\mathrm{Var}\left(\widehat{Y}_i(1;\psi)|S_i=1\right) = 0$, implying $g(\{\kappa, \kappa\}) = 0$ for any constant κ.

Second, suppose $Y_{ij}(\mathbf{z}_i) = f(\mathbf{z}_i)$ for $j = 1, 2, 3$ and $\mathbf{z}_i \in R_2^3$, where f is some real-valued function of $\mathbf{z}_i$. In other words, for any treatment program $\mathbf{z}_i \in R_2^3$, all individuals in group i have the same response. Suppose also that $f(011) \neq f(110) \neq f(101)$. Now $\widehat{Y}_i(1;\psi) = f(\mathbf{Z}_i)$, implying $\mathrm{Var}\left(\widehat{Y}_i(1;\psi)|S_i=1\right)$ equals the sample variance of the set $\{f(\mathbf{z}_i):\mathbf{z}_i \in R_2^3\}$. Thus,

$\text{Var}\left(\widehat{Y}_i(1;\psi)|S_i{=}1\right)>0$. However, because $O_{i1}(\mathbf{Z}_i)=\{f(\mathbf{Z}_i),f(\mathbf{Z}_i)\}$ in this scenario, it follows that $E\{g(O_{i1}(\mathbf{Z}_i))|\ S_i=1\}=0$. Thus, g is a biased estimator of $\text{Var}\left(\widehat{Y}_i(1;\psi)|S_i{=}1\right)$, a contradiction.

Intuitively, an unbiased estimator of $\text{Var}\left(\widehat{Y}_i(1;\psi)|S_i{=}1\right)$ does not exist in general because the observed data no way to distinguish between these two scenarios under either scenario, the observed outcomes are all equal, that is, $Y_{i1}(\mathbf{Z}_i)=Y_{i2}(\mathbf{Z}_i)=Y_{i3}(\mathbf{Z}_i)$. However, with the addition of Assumption 2, one can rule out the possibility of the second scenario. Namely, under Assumption 2, $Y_{i1}(110)=Y_{i1}(101)$ and $Y_{i2}(110)=Y_{i2}(011)$, implying $f(011)=f(110)f(101)$.

More generally, unbiased variance estimators do not exist without further assumptions (such as stratified interference) because observing $Y_{ij}(\mathbf{z}_i)$ provides no information about $Y_{ij}(\boldsymbol{\omega}_i)$ for $\boldsymbol{\omega}_i\neq\mathbf{z}_i$. Under Assumption 2, each individual now has only two potential outcomes, one for $z_{ij}=0$ and one for $z_{ij}=1$, within a particular stratum $R_k^{n_i}$. Therefore, given S_i, the observed data under one treatment program will provide information about the potential outcomes under other treatment programs. For example, suppose ψ is a mixed strategy such that K_i is fixed. Then, under Assumption 2, the outcomes for individual j are constant for all $\mathbf{z}_i\in R_{K_i}^{n_i}$ such that $z_{ij}=1$. Denote this value by $Y_{ij}(1;\psi)$, that is, $Y_{ij}(1;\psi)\equiv Y_{ij}(\boldsymbol{\omega},z_{ij}=1)$ for any $\boldsymbol{\omega}\in R_{K_i-1}^{n_i-1}$. Define $Y_{ij}(0;\psi)$ similarly.

4.2.2 Variance Estimators Assuming Stratified Interference

Theorem 4: Let

$$\widehat{\text{Var}}\left(\widehat{Y}_i(1;\psi)|S_i{=}1\right)\equiv\left(1-\frac{K_i}{n_i}\right)\frac{\widehat{\sigma}_{i1}^2(\psi)}{K_i},$$

where $\widehat{\sigma}_{i1}^2(\psi)\equiv\sum_{j=1}^{n_i}\left\{Y_{ij}(1;\psi)-\widehat{Y}_i(1;\psi)\right\}^2 Z_{ij}/(K_i-1)$ is the within-group sample variance, and

$$\widehat{\text{Var}}\left(\widehat{Y}_i(1;\psi)\right)\equiv\left(1-\frac{C}{N}\right)\frac{\widehat{\sigma}_{g1}^2(\psi)}{C}+\frac{1}{CN}\sum_{i=1}^{N}\left(1-\frac{K_i}{n_i}\right)\frac{\widehat{\sigma}_{i1}^2(\psi)}{K_i}S_i,$$

where $\widehat{\sigma}_{g1}^2(\psi)\equiv\sum_{i=1}^{N}\left\{\widehat{Y}_i(1;\psi)-\widehat{Y}(1;\psi)\right\}^2 S_i/(C-1)$. Define $\widehat{\text{Var}}\left(\widehat{Y}_i(0;\psi)|S_i{=}1\right)$, $\widehat{\sigma}_{i0}^2(\psi)$, $\widehat{\text{Var}}\left(\widehat{Y}(0;\psi)\right)$, and $\widehat{\sigma}_{g0}^2(\psi)$ analogously. Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{Y}_i(z;\psi)|S_i{=}1\right)|S_i{=}1\right\}=\text{Var}\left(\widehat{Y}_i(z;\psi)|S_i{=}1\right) \tag{7}$$

and $E\left\{\widehat{\text{Var}}\left(\widehat{Y}(z;\psi)\right)\right\}=\text{Var}\left(\widehat{Y}(z;\psi)\right)$ for $z=0,1$ for $z=0,1$.

Theorem 5: Let

$$\widehat{\text{Var}}\left(\widehat{CE}_i^D(\psi)|S_i{=}1\right)\equiv\frac{\widehat{\sigma}_{i1}^2(\psi)}{K_i}+\frac{\widehat{\sigma}_{i0}^2(\psi)}{n_i-K_i}. \tag{8}$$

Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}_i^D(\psi)|S_i{=}1\right)|S_i{=}1\right\}=\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i{=}1\right)+\sigma_{i(0-1)}^2(\psi)/n_i,$$

where $\sigma^2_{i(0-1)}(\psi) \equiv \sum_{j=1}^{n_i} \left[\left\{Y_{ij}(0;\psi) - Y_{ij}(1;\psi)\right\} - \left\{\bar{Y}_i(0;\psi) - \bar{Y}_i(1;\psi)\right\}\right]^2 / (n_i - 1)$ is the variance of the n_i differences $Y_{ij}(0;\psi) - Y_{ij}(1;\psi)$.

Corollary: Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)|S_i=1\right\} \geq \text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right),$$

with equality holding if and only if

$$Y_{ij}(0;\psi) = Y_{ij}(1;\psi) + \eta_{Di} \tag{9}$$

for fixed constant η_{Di} and $j = 1, \ldots, n_i$.

Remark: The corollary to Theorem 5 says (8) is a conditionally unbiased estimator of $\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)$ if and only if the individual direct effect is additive. If (9) does not hold, (8) will be a positively biased estimator of $\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)$. This could occur, for instance, if the potential outcomes are binary, taking on values 0 and 1 only. In this case, (9) will only be true if either (i) $\eta_{Di} = 0$ or (ii) $|\eta_{Di}| = 1$, with (ii) corresponding to the scenario that either $(Y_{ij}(0;\psi), Y_{ij}(1;\psi)) = (0, 1)$ for all j or $(Y_{ij}(0;\psi), Y_{ij}(1;\psi)) = (1, 0)$ for all j.

Theorem 6: Let

$$\widehat{\text{Var}}\left(\widehat{CE}^D(\psi)\right) \equiv \left(1 - \frac{C}{N}\right)\frac{\widehat{\sigma}_D^2(\psi)}{C} + \frac{1}{CN}\sum_{i=1}^N \widehat{\text{Var}}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)S_i, \tag{10}$$

where $\widehat{\sigma}_D^2(\psi) \equiv \sum_{i=1}^N \left\{\widehat{CE}_i^D(\psi) - \widehat{CE}^D(\psi)\right\}^2 S_i / (C - 1)$. Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}^D(\psi)\right)\right\} = \text{Var}\left(\widehat{CE}^D(\psi)\right) + \frac{1}{N^2}\sum_{i=1}^N \sigma^2_{i(0-1)}(\psi)/n_i.$$

Corollary: Under Assumptions 1 and 2, $E\left\{\widehat{\text{Var}}\left(\widehat{CE}^D(\psi)\right)\right\} \geq \text{Var}\left(\widehat{CE}^D(\psi)\right)$ with equality holding if and only if (9) holds for all $i = 1, \ldots, N$.

Remark: The corollary to Theorem 6 is similar to the corollary to Theorem 5 in that (10) is an unbiased estimator of $\text{Var}\left(\widehat{CE}^D(\psi)\right)$ if and only if the individual direct effects are additive. If direct additivity does not hold for all individuals, (10) will be positively biased. Analogous results for the group average indirect, total, and overall effects follow from Theorems 7-9.

Theorem 7: Let $\widehat{\text{Var}}\left(\widehat{CE}^I(\phi,\psi)\right) \equiv \widehat{\sigma}_{g0}^2(\phi)/(N - C) + \widehat{\sigma}_{g0}^2(\psi)/C$. Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}^I(\phi,\psi)\right)\right\} = \text{Var}\left(\widehat{CE}^I(\phi,\psi)\right) + \sigma^2_{g(0-0)}(\phi,\psi)/N,$$

where $\sigma^2_{g(0-0)}(\phi,\psi) \equiv \sum_{i=1}^N \left[\left\{\bar{Y}_i(0;\phi) - \bar{Y}_i(0;\psi)\right\} - \left\{\bar{Y}(0;\phi) - \bar{Y}(0;\psi)\right\}\right]^2 / (N - 1)$ is the variance of the N differences $\bar{Y}_i(0;\phi) - \bar{Y}_i(0;\psi)$.

Corollary: Under Assumptions 1 and 2, $E\left\{\widehat{\text{Var}}\left(\widehat{CE}^I(\phi,\psi)\right)\right\} \geq \text{Var}\left(\widehat{CE}^I(\phi,\psi)\right)$ with equality holding if and only if $\bar{Y}_i(0;\phi) = \bar{Y}_i(0;\psi) + \eta_I$ for fixed constant η_I and $i = 1, \ldots, N$.

Theorem 8: Let $\widehat{\text{Var}}\left(\widehat{CE}^T(\phi,\psi)\right) \equiv \widehat{\sigma}_{g0}^2(\phi)/(N-C) + \widehat{\sigma}_{gI}^2(\psi)/C$. Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}^T(\phi,\psi)\right)\right\} = \text{Var}\left(\widehat{CE}^T(\phi,\psi)\right) + \sigma_{g(0-1)}^2(\phi,\psi)/N,$$

where $\sigma_{g(0-1)}^2(\phi,\psi) \equiv \sum_{i=1}^N \left[\left\{\bar{Y}_i(0;\phi) - \bar{Y}_i(1;\psi)\right\} - \left\{\bar{Y}(0;\phi) - \bar{Y}(1;\psi)\right\}\right]^2/(N-1)$.

Corollary: Under Assumptions 1 and 2, $E\left\{\widehat{\text{Var}}\left(\widehat{CE}^T(\phi,\psi)\right)\right\} \geq \text{Var}\left(\widehat{CE}^T(\phi,\psi)\right)$ with equality holding if and only if $\bar{Y}_i(0;\phi) = \bar{Y}_i(1;\psi) + \eta_T$ for fixed constant η_T and $i = 1, \ldots, N$.

Theorem 9: Let $\widehat{\text{Var}}\left(\widehat{CE}^O(\phi,\psi)\right) \equiv \widehat{\sigma}_M^2(\phi)/(N-C) + \widehat{\sigma}_M^2(\psi)/C$, where $\widehat{\sigma}_M^2(\psi) \equiv \sum_{i=1}^N \left\{\widehat{Y}_i(\psi) - \widehat{Y}(\psi)\right\}^2 S_i/(C-1)$ and $\widehat{\sigma}_M^2(\phi)$ is defined analogously. Under Assumptions 1 and 2,

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}^O(\phi,\psi)\right)\right\} = \text{Var}\left(\widehat{CE}^O(\phi,\psi)\right) + \sigma_M^2(\phi,\psi)/N,$$

where $\sigma_M^2(\phi,\psi) \equiv \sum_{i=1}^N \left[\left\{\bar{Y}_i(\phi) - \bar{Y}_i(\psi)\right\} - \left\{\bar{Y}(\phi) - \bar{Y}(\psi)\right\}\right]^2/(N-1)$.

Corollary: Under Assumptions 1 and 2, $E\left\{\widehat{\text{Var}}\left(\widehat{CE}^O(\phi,\psi)\right)\right\} \geq \text{Var}\left(\widehat{CE}^O(\phi,\psi)\right)$ with equality holding if and only if $\bar{Y}_i(\phi) = \bar{Y}_i(\psi) + \eta_O$ for fixed constant η_O and $i = 1, \ldots, N$.

5. EXAMPLES

5.1 Housing Vouchers

Motivated by randomized studies designed to assess the effect of vouchers on housing mobility, Sobel (2006) proposed causal estimands and estimators when interference between units is present. At the first level, Sobel considered the effect of housing vouchers on the lease-up rate, that is, whether a household moves. At the second level, he considered voucher effects on other outcomes such as parents' perceptions of safety, welfare receipt, and child health. In this section, some of these estimands and estimators are shown to be special cases of those defined in Sections 3 and 4.1. To begin, we demonstrate that Sobel's causal estimand and estimator of the voucher effect on the lease-up rate are examples of the group average total causal effect estimand and estimator. Because Sobel considered just one group, we drop the subscript *i* for group in the rest of this section.

Consider a study where *n* households within a neighborhood are randomized to receive a housing voucher. In our terminology, households correspond to individuals and the neighborhood corresponds to a single group. Let $Z_j = 1$ if the *j*th household receives a voucher, $Z_j = 0$ otherwise for $j = 1, \ldots, n$. Let $Y_j(\mathbf{z}) = 1$ if the *j*th household moves using a voucher, $Y_j(\mathbf{z}) = 0$ otherwise. Because moving using a voucher is clearly not possible without a voucher, it follows immediately from the definition of $Y_j(\mathbf{z})$ that

$$Y_j\left(\mathbf{z}_{(j)}, z_j = 0\right) = 0 \quad \text{for } j = 1, \ldots, n. \tag{11}$$

Suppose $\Pr_\phi(\mathbf{Z} = \mathbf{0}) = 1$; that is, ϕ corresponds to the benchmark allocation strategy where no household receives a voucher. Then (11) implies

$$\overline{CE}^{\,T}(\phi,\psi) = -\frac{1}{n}\sum_{j=1}^{n}\sum_{\mathbf{z}\in R^{n-1}} Y_j\left(\mathbf{Z}_{(j)}=\mathbf{z},Z_j=1\right) \times \Pr_\phi\left(\mathbf{Z}_{(j)}=\mathbf{z}|Z_j=1\right).$$

for any other household assignment strategy ψ. In particular, suppose ψ corresponds to the mixed assignment strategy, where exactly k of n households receive a voucher. Then

$$\begin{aligned}
\overline{CE}^{\,T}&(\phi,\psi)\\
&= -\frac{(k-1)!(n-k)!}{n!}\sum_{\mathbf{z}\in R_{k-1}^{n-1}}\sum_{j=1}^{n} Y_j\left(\mathbf{Z}_{(j)}=\mathbf{z},Z_j=1\right)\\
&= -\frac{k!(n-k)!}{n!}\sum_{\mathbf{z}\in R_{k-1}^{n-1}}\sum_{\zeta=0}^{1}\frac{1}{k}\sum_{j=1}^{n} Y_j\left(\mathbf{Z}_{(j)}=\mathbf{z},Z_j=\zeta\right)\zeta\\
&= -\frac{k!(n-k)!}{n!}\sum_{\mathbf{z}\in R_{k}^{n}}\frac{1}{k}\sum_{j=1}^{n} Y_j\left(\mathbf{Z}=\mathbf{z}\right) z_j,
\end{aligned}$$

which is equivalent (up to a minus sign) to equation (2) of Sobel in the setting where there are two levels of treatment. Sobel actually considered the more general situation of three treatment levels, which is not considered here. Sobel's corresponding estimator, the observed lease-up rate among voucher recipients, is equivalent to $\widehat{Y}(1;\psi)$ as given by (6). Under exclusion restriction (11), $\widehat{\overline{CE}}^{\,T}(\phi,\psi) = -\widehat{Y}(1;\psi)$.

Although interference is possible among those receiving housing vouchers, the exclusion restriction (11) precludes interference when a household does not receive a voucher. Thus, the indirect effect of housing vouchers on mobility is 0, $\overline{CE}^{\,I}(\phi,\psi)=0$, and the total effects equal the direct effects. If households could move without the aid of a voucher, an analysis based on (11) could potentially overestimate the magnitude of the total effect of vouchers. By instead defining $Y_j(\mathbf{z}) = 1$ if the j th household moves and 0 otherwise, a voucher allocation strategy may have indirect effects in those not receiving vouchers as well as direct and total causal effects. For example, this would allow for the possibility that if several neighbors move because they receive vouchers, household j might also move even though they did not obtain a voucher. Estimating such indirect causal effects of voucher distribution within a neighborhood on those households that do not receive a voucher would likely be of interest to policy makers. Of course, in this case, observation of potential outcomes under the benchmark allocation program $\mathbf{Z} = \mathbf{0}$ would also be necessary to estimate the indirect and total effects.

Sobel's estimand of the effects of vouchers on outcomes such as welfare receipt can be viewed as an example of the group average overall effect estimand. For instance, assume all households that receive a voucher subsequently move and ϕ is the benchmark allocation strategy. Then taking the expected value of Sobel's "average effect" [his eq. (3)] over the distribution of possible intervention programs under ψ yields the group average overall effect $\overline{CE}^{\,O}(\phi,\psi)$. Sobel noted that the average overall effect is a weighted average of the indirect effects on those not receiving a voucher, which he called spillover effects, and the effects in those receiving a voucher. More precisely, for ϕ the benchmark allocation strategy, the overall effect equals the following weighted sum of the indirect and total effects:

$$\overline{CE}^{\,O}(\phi,\psi) = \Pr_\psi\left(Z_j=0\right)\overline{CE}^{\,I}(\phi,\psi) + \Pr_\psi\left(Z_j=1\right)\overline{CE}^{\,T}(\phi,\psi).$$

Because the design of the housing voucher study does not include randomizing some neighborhoods to the benchmark allocation, outcomes such as welfare receipt and parents'

perception of safety are not observed under this allocation. Thus, without further assumptions akin to (11), voucher effects on these outcomes are not identifiable from the data. Alternatively, these effects are identifiable by considering a population of neighborhoods and a two-stage randomization design. For example, neighborhoods within a city or set of cities could be identified that were sufficiently separated geographically to ensure that the assumption of no interference between neighborhoods is plausible. Then, in the first stage of randomization, some neighborhoods could be randomly assigned the benchmark allocation and other neighborhoods to an allocation strategy where in the second stage, a specified proportion of randomly selected households would receive a voucher. Such a design would permit estimation of the direct, indirect, total, and overall effects of housing vouchers on the outcomes described previously without making exclusion restrictions such as (11). Estimation of the variances of the causal effect estimators would also be possible with this design under the additional assumption of stratified interference.

5.2 Vaccines

Direct application of the proposed methods to the data given in Table 1 is not appropriate because baris were not randomly assigned to particular levels of vaccine coverage in the actual trial. Therefore, for illustrative purposes, we consider a hypothetical two-stage randomized placebo-controlled trial of cholera vaccines in a setting similar to that of Ali et al. (2005). Suppose in the first stage of this hypothetical trial that five geographically separate groups were randomized using a mixed allocation strategy ν such that three groups were assigned ϕ and the remaining two were assigned ψ. Then, in the second stage, suppose 30% of individuals were randomly chosen to receive vaccine within groups assigned ϕ and 50% of individuals were randomly chosen to receive vaccine within groups assigned ψ. Individuals were then followed for one year for detection of cholera. Results from this hypothetical trial are given in Table 2.

Estimates of the population average direct, indirect, total, and overall effects are given in Table 3. The estimated variances are also presented. Note the direct effect estimate is nearly three times greater under ϕ (30% coverage) compared to ψ (50% coverage). Ali et al. (2005) noted a similar phenomenon and cautioned that high levels of vaccine coverage can bias estimates of vaccine efficacy (i.e., the direct effect of vaccination). In fact, the issue here is not one of bias, but rather that the parameter being estimated can depend on the level of coverage due to interference between individuals. Moreover, the totality of effects of a vaccination strategy must be viewed by considering indirect, total, and overall effects in addition to direct effects. Estimates of these other effects can easily be interpreted by investigators. For example, the indirect effect estimate in Table 3 suggests 50% vaccine coverage results in 2.8 fewer cholera cases per 1,000 unvaccinated individuals per year compared to only 30% vaccine coverage. Note the estimated total effect of being vaccinated under ψ is over three times the corresponding estimated direct effect, demonstrating the importance of randomizing groups to different levels of vaccine coverage. Had all groups received 50% coverage such that only the direct effect could be estimated, the utility of vaccination would have been substantially underestimated. The estimated overall effect provides a simple summary comparison of the two strategies, indicating that, on average, 50% vaccine coverage results in 2.4 fewer cases of cholera per 1,000 individuals per year compared to 30% vaccine coverage.

6. DISCUSSION

In this article, estimands for direct, indirect, total, and overall causal effects of different treatment strategies are proposed in the setting where interference between individuals is possible. Relations between the estimands are established, and estimators of the proposed estimands are presented. These estimators are shown to be unbiased assuming a two-stage

randomization procedure with a mixed assignment mechanism at each stage (Assumption 1). Under an additional assumption of stratified interference (Assumption 2), variance estimators of the causal effect estimators are derived that are unbiased under additivity and positively biased otherwise.

This article builds on previous work in several significant ways. First, causal inference in the presence of interference is considered in a general framework not specific to any one subject area, unifying previous work on housing mobility studies (Sobel 2006) and infectious diseases (Halloran and Struchiner 1991, 1995). The definitions of individual direct, indirect, total, and overall causal effects in the presence of interference (Halloran and Struchiner 1991, 1995) are formally extended to groups and populations of groups by averaging over all possible treatment assignments for particular allocation strategies (Sobel 2006). By considering a population of groups, rather than just one group as in Sobel (2006), unbiased estimators of the causal estimands of interest are derived without requiring exclusion restriction assumptions. The variance of causal effect estimators in the presence of interference is also considered, which had not been done previously. The utility of the proposed variance estimators will depend on whether Assumptions 1 and 2 are reasonable. Assumption 1 is determined by the experimental design and, thus, should be under control of the investigator. Assumption 2 may be reasonable in many settings, such as in the evaluation of the effects of vaccines, educational interventions, or housing vouchers. However, in other contexts, such as in cross-over trials or plant variety studies, assuming different forms of interference may be more appropriate.

The methods developed here could be extended to settings with more than two treatment levels or noncompliance as in Sobel (2006). Consideration of population subgroups may be of interest in the presence of interference. As in Halloran, Longini, Cowart, and Nizam (2002), one could define the indirect, total, and overall effects for different subgroups of the population. For example, a strategy of vaccinating 70% of children against influenza and another strategy of not vaccinating children could be compared by the indirect effects on the incidence of influenza in adults from the same population. Similarly, Moulton et al. (2001) and Sobel (2006) considered populations consisting of participants and nonparticipants; that is, a subset of the population does not receive either treatment.

In addition to these extensions, other areas of research remain to be explored in the setting where interference between individuals is present. Different randomization strategies and interference structures might be considered. An anonymous referee suggested more efficient variance estimators might be derived if one is willing to make certain additional additivity assumptions. For example, under (9), $\widehat{\sigma}_{i1}^2(\psi)$ and $\widehat{\sigma}_{i0}^2(\psi)$ are estimating the same parameter, and, thus, more efficient variance estimators might be obtained by combining data from individuals assigned $Z_{ij} = 0$ and individuals assigned $Z_{ij} = 1$. Stochastic causal models could also be considered, wherein the potential outcomes $Y_{ij}(\mathbf{z}_i)$ are treated as random rather than fixed (e.g., as in Robins and Greenland 1989). Methods are needed to construct confidence intervals for the different causal effects, perhaps by building on the work of Rosenbaum (2007).

Rubin (2005) recently noted that causal inference can be conducted by making only two assumptions: a probabilistic model about the treatment assignment mechanism and SUTVA. By providing a framework where the latter assumption is not completely necessary, our work here will hopefully contribute to a foundation for causal inference in the presence of interference upon which others can build.

Acknowledgements

APPENDIX: PROOFS OF THEOREMS 1-9

A.1 Proof of Theorem 1

Without loss of generality, let $z = 1$. Under Assumption 1, K_i is fixed, so that

$$E\left\{\widehat{Y}_i(1;\psi)|S_i=1\right\} = \frac{1}{K_i}\sum_{j=1}^{n_i}\sum_{\mathbf{z}\in R_{K_i}^{n_i}}\mathrm{Pr}_\psi\left(\mathbf{Z}_i=\mathbf{z}\right)Y_{ij}(\mathbf{z})I\left[z_{ij}=1\right].$$

Now any $\mathbf{z}$ such that $z_{ij} = 0$ does not contribute to the summation, so that we can equivalently write

$$E\left\{\widehat{Y}_i(1;\psi)|S_i=1\right\} = \frac{1}{K_i}\sum_{j=1}^{n_i}\sum_{\omega\in R_{K_i-1}^{n_i-1}}\mathrm{Pr}_\psi\left(\mathbf{Z}_{i(j)}=\omega,Z_{ij}=1\right)\times Y_{ij}\left(\mathbf{z}_{i(j)}=\omega,z_{ij}=1\right)$$

$$= \frac{1}{K_i}\sum_{j=1}^{n_i}\sum_{\omega\in R_{K_i-1}^{n_i-1}}\mathrm{Pr}_\psi\left(\mathbf{Z}_{i(j)}=\omega|Z_{ij}=1\right)\times\mathrm{Pr}_\psi\left(Z_{ij}=1\right)Y_{ij}\left(\mathbf{z}_{i(j)}=\omega,z_{ij}=1\right).$$

Under Assumption 1, $\mathrm{Pr}_\psi(Z_{ij}=1) = K_i/n_i$, implying

$$E\left\{\widehat{Y}_i(1;\psi)|S_i=1\right\} = \frac{1}{n_i}\sum_{j=1}^{n_i}\sum_{\omega\in R_{K_i-1}^{n_i-1}}\mathrm{Pr}_\psi\left(\mathbf{Z}_{i(j)}=\omega|Z_{ij}=1\right)\times Y_{ij}\left(\mathbf{z}_{i(j)}=\omega,z_{ij}=1\right)$$

$$= \bar{Y}_i(1;\psi).$$

A.2 Proof of Theorem 2

Without loss of generality, let $z = 1$. Using the fact that $E\left\{\widehat{Y}(1;\psi)\right\}=E\left[E\left\{\widehat{Y}(1;\psi)|\mathbf{S}\right\}\right]$, from Theorem 1 it follows that $E\left\{\widehat{Y}(1;\psi)\right\}=E\left\{\sum_{i=1}^{N}\bar{Y}_i(1;\psi)S_i/C\right\}=\bar{Y}(1;\psi)$.

A.3 Proof of Theorem 3

The conditional expectation result follows from

$$E\left\{\widehat{Y}_i(\psi)|S_i=1\right\} = \sum_{j=1}^{n_i}\sum_{\mathbf{z}\in R^{n_i}}Y_{ij}(\mathbf{z})\mathrm{Pr}_\psi\left(\mathbf{Z}_i=\mathbf{z}\right)/n_i$$

$$= \sum_{j=1}^{n_i}\bar{Y}_{ij}(\psi)/n_i$$

$$= \bar{Y}_i(\psi).$$

The remainder of the proof parallels that of Theorem 2.

A.4 Proof of Theorem 4

The proof follows directly from known properties of estimators of population means using simple random sampling (SRS) and two-stage cluster sampling (see, e.g., Kish 1965, chap. 2; Splawa-Neyman 1990; Thompson 1992, chaps. 2 and 3). For example, given $S_i = 1$ and Assumptions 1 and 2, $\widehat{Y}_i(1;\psi)$ can be viewed as the sample mean from a simple random sample drawn without replacement from $\{Y_{i1}(1;\psi), \ldots, Y_{in_i}(1;\psi)\}$. Thus,

$$\text{Var}\left(\widehat{\bar{Y}}_i\left(1;\psi\right)|S_i=1\right) = \left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i},$$

(A.1)

where $\sigma_{i1}^2\left(\psi\right) \equiv \sum_{j=1}^{n_i}\left\{Y_{ij}\left(1;\psi\right) - \bar{Y}_i\left(1;\psi\right)\right\}^2 / (n_i - 1)$ is the within-group variance. It is also well known that $E\left(\widehat{\sigma}_{i1}^2\left(\psi\right)|S_i=1\right) = \sigma_{i1}^2\left(\psi\right)$, implying (7) holds.

Similarly, that $E\left\{\widehat{\text{Var}}\left(\bar{Y}\left(z;\psi\right)\right)\right\} = \text{Var}\left(\bar{Y}\left(z;\psi\right)\right)$ for $z = 0, 1$ follows from known results on two-stage cluster sampling. A sketch of a proof of this follows. First, one can show

$$\text{Var}\left\{\bar{Y}\left(1;\psi\right)\right\} = \left(1 - \frac{C}{N}\right)\frac{\sigma_{g1}^2\left(\psi\right)}{C} + \frac{1}{CN}\sum_{i=1}^N\left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i},$$

(A.2)

where $\sigma_{g1}^2\left(\psi\right) \equiv \sum_{i=1}^N\left\{\bar{Y}_i\left(1;\psi\right) - \bar{Y}\left(1;\psi\right)\right\}^2 / (N - 1)$. Next, note $E(S_i) = C/N$ such that

$$E\left\{\widehat{\text{Var}}\left(\bar{Y}\left(1;\psi\right)\right)\right\} = \left(1 - \frac{C}{N}\right)\frac{E\left\{\widehat{\sigma}_{g1}^2\left(\psi\right)\right\}}{C} + \frac{1}{N^2}\sum_{i=1}^N\left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i}.$$

(A.3)

So the remaining task at hand becomes finding $E\left\{\widehat{\sigma}_{g1}^2\left(\psi\right)\right\}$, which can be shown to equal

$$\frac{1}{N}\sum_{i=1}^N\left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i} + \sigma_{g1}^2\left(\psi\right).$$

Substituting this into (A.3) implies $E\left\{\widehat{\text{Var}}\left(\bar{Y}\left(1;\psi\right)\right)\right\}$ equals

$$\left(1 - \frac{C}{N}\right)\frac{\sigma_{g1}^2\left(\psi\right)}{C} + \left\{\left(1 - \frac{C}{N}\right)\frac{1}{C}\frac{1}{N} + \frac{1}{N^2}\right\}\sum_{i=1}^N\left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i},$$

which simplifies to (A.2).

A.5 Proof of Theorem 5

The proof follows from Splawa-Neyman (1990) and Rubin (1990); a sketch is given here. First, we derive $\text{Var}\left(\widehat{CE}_i^D\left(\psi\right)|S_i=1\right)$, which, of course, equals $\text{Var}\left\{\widehat{\bar{Y}}_i\left(0;\psi\right)|S_i=1\right\} + \text{Var}\left\{\widehat{\bar{Y}}_i\left(1;\psi\right)|S_i=1\right\} - 2\text{Cov}\left\{\widehat{\bar{Y}}_i\left(0;\psi\right),\widehat{\bar{Y}}_i\left(1;\psi\right)|S_i=1\right\}$. We know the form of $\text{Var}\left\{\widehat{\bar{Y}}_i\left(z;\psi\right)|S_i=1\right\}$ for $z = 0, 1$ from the proof of Theorem 4. Additionally, one can show $\text{Cov}\left(\widehat{\bar{Y}}_i\left(0;\psi\right),\widehat{\bar{Y}}_i\left(1;\psi\right)|S_i=1\right) = \left\{\sigma_{i(0-1)}^2\left(\psi\right) - \sigma_{i0}^2\left(\psi\right) - \sigma_{i1}^2\left(\psi\right)\right\} / (2n_i)$. Therefore,

$$\text{Var}\left(\widehat{CE}_i^D\left(\psi\right)|S_i=1\right) = \frac{K_i}{n_i}\frac{\sigma_{i0}^2\left(\psi\right)}{n_i - K_i} + \left(1 - \frac{K_i}{n_i}\right)\frac{\sigma_{i1}^2\left(\psi\right)}{K_i} - \frac{1}{n_i}\left\{\sigma_{i(0-1)}^2\left(\psi\right) - \sigma_{i0}^2\left(\psi\right) - \sigma_{i1}^2\left(\psi\right)\right\},$$

which simplifies to

$$\text{Var}\left(\widehat{CE}_i^D\left(\psi\right)|S_i=1\right) = \frac{\sigma_{i0}^2\left(\psi\right)}{n_i - K_i} + \frac{\sigma_{i1}^2\left(\psi\right)}{K_i} - \frac{\sigma_{i(0-1)}^2\left(\psi\right)}{n_i}.$$

The proof is then completed by noting that

$$E\left\{\widehat{\text{Var}}\left(\widehat{CE}_i^D\left(\psi\right)|S_i=1\right)|S_i=1\right\} = \frac{\sigma_{i0}^2\left(\psi\right)}{n_i - K_i} + \frac{\sigma_{i1}^2\left(\psi\right)}{K_i}.$$

A.6 Proof of Theorem 6

The proof follows along similar lines as the derivation of $E\left\{\widehat{\text{Var}}\left(\bar{Y}(1;\psi)\right)\right\}$ in the proof of Theorem 4. In particular, one can first show

$$\text{Var}\left(\widehat{CE}^D(\psi)\right) = \left(1 - \frac{C}{N}\right)\frac{\sigma_D^2(\psi)}{C} + \frac{1}{CN}\sum_{i=1}^{N}\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right),$$

(A.4)

where $\sigma_D^2(\psi) \equiv \sum_{i=1}^{N}\left\{\bar{CE}_i^D(\psi) - \bar{CE}^D(\psi)\right\}^2 / (N-1)$. Next, similar to $E\left\{\widehat{\sigma}_{g1}^2(\psi)\right\}$ in Theorem 4, one can show $E\left\{\widehat{\sigma}_D^2(\psi)\right\} = \sum_{i=1}^{N}\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)/N + \sigma_D^2(\psi)$. Taking the expected value of (10), we have

$$E\left[\widehat{\text{Var}}\left\{\widehat{CE}^D(\psi)\right\}\right] = \left(1 - \frac{C}{N}\right)\frac{\sigma_D^2(\psi)}{C} + \left(1 - \frac{C}{N}\right)\frac{1}{C}\frac{1}{N}\sum_{i=1}^{N}\text{Var}\left(\widehat{CE}_i^D(\psi)|S_i=1\right) + \frac{1}{CN}\sum_{i=1}^{N}E\left\{\widehat{\text{Var}}\left(\widehat{CE}_i^D(\psi)|S_i=1\right)S_i\right\}.$$

Combining this result with Theorem 5 proves the theorem.

A.7 Proof of Theorem 7

The proof follows along the same lines as the proof of Theorem 5. Namely, one can show $\text{Cov}\left(\widehat{\bar{Y}}(0;\phi),\widehat{\bar{Y}}(0;\psi)\right) = \left\{\sigma_{g(0-0)}^2(\phi,\psi) - \sigma_{g0}^2(\phi) - \sigma_{g0}^2(\psi)\right\}/(2N)$, where $\sigma_{g0}^2(\phi)$ and $\sigma_{g0}^2(\psi)$ are defined analogously to $\sigma_{g1}^2(\psi)$ in Theorem 4, implying $\text{Var}\left(\widehat{CE}^I(\phi,\psi)\right)$ equals

$$\text{Var}\left(\widehat{\bar{Y}}(0;\phi)\right) + \text{Var}\left(\widehat{\bar{Y}}(0;\psi)\right) + \frac{1}{N}\left\{\sigma_{g0}^2(\phi) + \sigma_{g0}^2(\psi) - \sigma_{g(0-0)}^2(\phi,\psi)\right\}.$$

(A.5)

From the proof of Theorem 4, we have

$$\begin{aligned}E\left\{\frac{\widehat{\sigma}_{g0}^2(\psi)}{C}\right\} &= \frac{\sigma_{g0}^2(\psi)}{C} + \frac{1}{CN}\sum_{i=1}^{N}\text{Var}\left(\bar{Y}_i(0;\psi)|S_i=1\right)\\ &= \text{Var}\left\{\widehat{\bar{Y}}(0;\psi)\right\} + \frac{\sigma_{g0}^2(\psi)}{N},\end{aligned}$$

and similarly $E\left\{\widehat{\sigma}_{g0}^2(\phi)/(N-C)\right\} = \text{Var}\left\{\widehat{\bar{Y}}(0;\phi)\right\} + \sigma_{g0}^2(\phi)/N$, which together with (A.5) prove the theorem.

A.8 Proof of Theorem 8

The proof is analogous to the proof of Theorem 7.

A.9 Proof of Theorem 9

As in the proof of Theorem 4, one can show

$$\text{Var}\left(\widehat{\bar{Y}}(\psi)\right) = \left(1 - \frac{C}{N}\right)\frac{\sigma_M^2(\psi)}{C} + \frac{1}{CN}\sum_{i=1}^{N}\text{Var}\left(\widehat{\bar{Y}}_i(\psi)|S_i=1\right),$$

where $\sigma_M^2(\psi) \equiv \sum_{i=1}^{N}\left\{\bar{Y}_i(\psi) - \bar{Y}(\psi)\right\}^2/(N-1)$. Following the same lines as the proofs of Theorems 5 and 7, one can also show

$$\mathrm{Cov}\left\{\widehat{Y}(\phi),\widehat{Y}(\psi)\right\}=\left\{\sigma_M^2(\phi,\psi)-\sigma_M^2(\phi)-\sigma_M^2(\psi)\right\}/(2N),\text{ where }\sigma_M^2(\phi)\text{ is defined analogously}$$

to $\sigma_M^2(\psi)$, implying

$$\mathrm{Var}\left\{CE^0(\phi,\psi)\right\}=\mathrm{Var}\left\{\widehat{Y}(\phi)\right\}+\mathrm{Var}\left\{\widehat{Y}(\psi)\right\}+\frac{1}{N}\left\{\sigma_M^2(\phi)+\sigma_M^2(\psi)-\sigma_M^2(\phi,\psi)\right\}.$$

Additionally, one can show $E\left\{\widehat{\sigma}_M^2(\psi)\right\}=\sigma_M^2(\psi)+\sum_{i=1}^{N}\mathrm{Var}\left\{\widehat{Y}_i(\psi)\,|S_i{=}1\right\}/N$, which implies $E\left\{\widehat{\sigma}_M^2(\psi)/C\right\}=\mathrm{Var}\left\{\widehat{Y}(\psi)\right\}+\sigma_M^2(\psi)/N$. Using an analogous result for $\widehat{\sigma}_M^2(\phi)$, the theorem follows.

REFERENCES

Ali M, Emch M, von Seidlein L, Yunus M, Sack DA, Rao M, Holmgren J, Clemens JD. Herd Immunity Conferred by Killed Oral Cholera Vaccines in Bangladesh: A Reanalysis. Lancet 2005;366:44–49. [PubMed: 15993232]

Angrist JD, Imbens GW, Rubin DB. Identification of Causal Effects Using Instrumental Variables. Journal of the American Statistical Association 1996;91:444–455.

Bailey RA, Kunert J. On Optimal Crossover Designs When Carryover Effects Are Proportional to Direct Effects. Biometrika 2006;93:613–625.

Borm GF, Melis RJF, Teerenstra S, Peer PG. Pseudo Cluster Randomization: A Treatment Allocation Method to Minimize Contamination and Selection Bias. Statistics in Medicine 2005;24:3535–3547. [PubMed: 16007575]

Cox, DR. Planning of Experiments. Wiley; New York: 1958.

Friedman, LM.; Furberg, C.; DeMets, DL. Fundamentals of Clinical Trials. Springer-Verlag; New York: 1998.

Halloran ME, Struchiner CJ. Study Designs for Dependent Happenings. Epidemiology 1991;2:331–338. [PubMed: 1742381]

Halloran ME, Struchiner CJ. Causal Inference in Infectious Diseases. Epidemiology 1995;6:142–151. [PubMed: 7742400]

Halloran ME, Longini IM, Cowart DM, Nizam A. Community interventions and the Epidemic Prevention Potential. Vaccine 2002;20:3254–3262. [PubMed: 12213394]

Hayes RJ, Alexander NDE, Bennett S, Cousens SN. Design and Analysis Issues in Cluster-Randomized Trials of Interventions Against Infectious Diseases. Statistical Methods in Medical Research 2000;9:95–116. [PubMed: 10946429]

Kempton, RA. Interference Between Plots. In: Kempton, RA.; Fox, PN., editors. Statistical Methods for Plant Variety Evaluation. Chapman & Hall; London: 1997. p. 101-116.

King JC, Stoddard JJ, Gaglani MB, Moore KA, Magder L, McClure E, Rubin JD, Englund JA, Neuzil K. Effectiveness of School-Based Influenza Vaccination. New England Journal of Medicine 2006;355:2523–2532. [PubMed: 17167135]

Kish, L. Survey Sampling. Wiley; New York: 1965.

Longini IM, Halloran ME, Nizam A. Model-Based Estimation of Vaccine Effects From Community Vaccine Trials. Statistics in Medicine 2002;21:481–495. [PubMed: 11836731]

Melis RJF, van Eijken MIJ, Borm GF, Wensing M, Adang E, van de Lisdonk EH, van Achterberg T, Olde Rikkert MGM. The Design of the Dutch EASYcare Study: A Randomised Controlled Trial on the Effectiveness of a Problem-Based Community Intervention Model for Frail Elderly People. BMC Health Services Research 2005;5:65. [PubMed: 16207382]

Moulton LH, O'Brien KL, Kohberger R, Chang I, Reid R, Weatherholtz R, Hackell JG, Siber GR, Santosham M. Design of a Group-Randomized *Streptococcus pneumoniae* Vaccine Trial. Controlled Clinical Trials 2001;22:438–452. [PubMed: 11514043]

Moulton LH, O'Brien KL, Reid R, Weatherholtz R, Santosham M, Siber GR. Evaluation of the Indirect Effects of a Pneumococcal Vaccine in a Community-Randomized Study. Journal of Biopharmaceutical Statistics 2006;16:453–462. [PubMed: 16892907]

Murray, DM. Design and Analysis of Group-Randomized Trials. Oxford University Press; New York: 1998.

Robins J, Greenland S. The Probability of Causation Under a Stochastic Model for Individual Risk. Biometrics 1989;45:1125–1138. [PubMed: 2611320]

Rosenbaum PR. Interference Between Units in Randomized Experiments. Journal of the American Statistical Association 2007;102:191–200.

Ross R. An Application of the Theory of Probabilities to the Study of a priori Pathometry. I. Proceedings of the Royal Society 1916;92:204–230.Ser. A

Rubin DB. Bayesian Inference for Causal Effects: The Role of Randomization. The Annals of Statistics 1978;6:34–58.

Rubin DB. Discussion of "Randomization Analysis of Experimental Data in the Fisher Randomization Test. Journal of the American Statistical Association 1980;75:591–593.by D. Basu

Rubin DB. Comment: Neyman (1923) and Causal Inference in Experiments and Observations Studies. Statistical Science 1990;5:472–480.

Rubin DB. Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. Journal of the American Statistical Association 2005;100:322–331.

Senn, S. Cross-Over Trials in Clinical Research. Wiley; New York: 1993.

Sobel M. What Do Randomized Studies of Housing Mobility Demonstrate? Causal Inference in the Face of Interference. Journal of the American Statistical Association 2006;101:1398–1407.

Som, RK. A Manual of Sampling Techniques. Heinemann; London: 1973.

Splawa-Neyman J, Dabrowska DM, Speed TP. On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9. Statistical Science 1990;5:465–472.

Struchiner CJ, Halloran ME, Robins JM, Spielman A. The Behavior of Common Measures of Association Used to Assess a Vaccination Program Under Complex Disease Transmission Patterns—A Computer Simulation Study of Malaria Vaccines. International Journal of Epidemiology 1990;19:187–196. [PubMed: 2351514]

Thompson, SK. Sampling. Wiley; New York: 1992.

Table 1

Risk of cholera in recipients of killed oral cholera vaccines or placebo, by level of coverage of the bari during one year of follow-up, based on data from Ali et al. (2005)

Level of vaccine coverage	Vaccine recipients				Placebo recipients		
	Target population	Total	Cases	Risk per 1,000 population	Total	Cases	Risk per 1,000 population
>50%	22,394	12,541	16	1.27	6,082	9	1.47
41–50%	24,159	11,513	26	2.26	5,801	27	4.65
36–40%	24,583	10,772	17	1.58	5,503	26	4.72
28–35%	25,059	8,883	22	2.48	4,429	26	5.87
<28%	24,954	5,627	15	2.66	2,852	20	7.01

Table 2

Illustrative example of a two-stage randomized placebo-controlled vaccine trial based on data from Ali et al. (2005)

| Group i | Group assignment S_i | Vaccine recipients ($Z_{ij} = 1$) | | Placebo recipients ($Z_{ij} = 0$) | |
		Total $\Sigma_j Z_{ij}$	Cases $\Sigma_j Z_{ij} Y_{ij}$ (Z_i)	Total $\Sigma_j (1 - Z_{ij})$	Cases $\Sigma_j (1 - Z_{ij}) Y_{ij}$ (Z_i)
1	1	12,541	16	12,541	18
2	1	11,513	26	11,513	54
3	0	10,772	17	25,134	119
4	0	8,883	22	20,727	122
5	0	5,627	15	13,130	92

NOTE: Group assignment $S_i = 1$ (0) corresponds to 50% (30%) vaccine coverage.

Table 3

Estimates of population average direct, indirect, total, and overall effects per 1,000 individuals per year for data in Table 2

Effect	Parameter	Estimate	Estimated variance
Direct	$\overline{CE}^D(\psi)$	1.30	.856
Direct	$\overline{CE}^D(\phi)$	3.64	.178
Indirect	$\overline{CE}^I(\phi, \psi)$	2.81	3.079
Total	$\overline{CE}^T(\phi, \psi)$	4.11	.672
Overall	$\overline{CE}^O(\phi, \psi)$	2.37	1.430