

Collection of Biostatistics Research Archive
COBRA Preprint Series

Year 2011

Paper 85

Components of the indirect effect in vaccine
trials: identification of contagion and
infectiousness effects

Tyler J. VanderWeele* Eric J. Tchetgen†
M. Elizabeth Halloran‡

*Harvard University, tvanderw@hsph.harvard.edu

†Harvard University, etchetge@hsph.harvard.edu

‡University of Washington, betz@fhcrc.org

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/cobra/art85>

Copyright ©2011 by the authors.

Components of the indirect effect in vaccine trials: identification of contagion and infectiousness effects

Tyler J. VanderWeele, Eric J. Tchetgen, and M. Elizabeth Halloran

Abstract

Vaccination of one person may prevent the infection of another either because (i) the vaccine prevents the first from being infected and from infecting the second or because (ii) even if the first person is infected, the vaccine may render the infection less infectious. We might refer to the first of these mechanisms as a contagion effect and the second as an infectiousness effect. In this paper, for the simple setting of a randomized vaccine trial with households of size two, we use counterfactual theory under interference to provide formal definitions of a contagion effect and an infectiousness effect. Using ideas analogous to mediation analysis, we show that the indirect effect (the effect of one individual's vaccine on another's outcome) can be decomposed into a contagion effect and an infectiousness effect on the risk difference, risk ratio, odds ratio and vaccine efficacy scales. We provide identification assumptions for such contagion and infectiousness effects, and describe a simple statistical techniques to estimate these effects when they are identified. We also give a sensitivity analysis techniques to assess how inferences would change under violations of the identification assumptions. The concepts and results of this paper are illustrated with sample vaccine trial data.

Introduction

Administering a vaccine to one or several individuals in a population may protect not only those vaccinated individuals from infection or disease, but also other individuals as well. In the causal inference vaccine literature, the protection afforded other unvaccinated individuals has been called the indirect effects of vaccination. A number of papers have considered the methodology of estimating such indirect effects.^{1–5} With an indirect effect, we might conceive of two distinct mechanisms by which such an effect may operate. Suppose we have two persons in a household and that we vaccinate the first. Vaccinating the first person may protect the second by preventing the infection in the first and thereby preventing the infection from spreading from the first person to the second. Alternatively, vaccinating the first person may protect the second because, even if the first person is infected the vaccine may render the infection less infectious, thereby preventing the first person from infecting the second. This latter effect is sometimes referred to as an "infectiousness effect".^{3,6,7} We will refer to the former as a "contagion effect", following terminology in the social network literature^{8,9}, though we acknowledge that "infectiousness" and "contagion" are sometimes used interchangeably in the infectious disease literature.

In this paper we show that in households of size two, the indirect effect can be decomposed into a contagion effect and an infectiousness effect. We draw on theory for causal inference under interference^{4,7,10,11} and on mediation analysis^{12–15} to provide formal counterfactual decompositions for each of these effects. We show that decompositions of the indirect effect into a contagion and infectiousness effect hold for the risk difference, risk ratio, odds ratio and vaccine efficacy scales. We discuss identification assumptions that suffice to estimate these effects from vaccine trial data and propose a simple statistical modeling strategy using logistic regressions to estimate these effects. We also describe a sensitivity analysis technique for these effects that can be employed when the identification assumptions do not hold to assess the sensitivity of the estimates to violations in the assumptions being made. We illustrate the methodology with application some sample vaccine trial data.

Concepts and Definitions

We consider a setting similar to that in VanderWeele and Tchetgen Tchetgen⁷ in which there are N households indexed by $i = 1, \dots, N$ such that each household consists of two persons indexed by $j = 1, 2$. We will generalize this setting somewhat in the final section of the paper. We let A_{ij} denote the vaccine status for individual j in household i . We let $A_{ij} = 1$ denote that the

person j in household i was vaccinated and let $A_{ij} = 0$ denote that the person was not vaccinated. We let Y_{ij} denote the infection status of individual j in household i after some fixed follow up period. We let $Y_{ij}(a_{i1}, a_{i2})$ denote the counterfactual outcome for person j in household i if the two persons in that household had, possibly contrary to fact, vaccination status of (a_{i1}, a_{i2}) . For example, $Y_{i2}(1, 0)$ would denote what would have happened to person 2 if person 1 had received the vaccine and person 2 had not; and $Y_{i1}(0, 0)$ denotes what would have happened to person 1 if neither person 1 nor person 2 had received the vaccine. Note that under this counterfactual or "potential outcomes" notation, the potential outcome for individual 1, $Y_{i1}(a_{i1}, a_{i2})$ depends on the vaccine status of both person 1 and person 2, and likewise the potential outcome for individual 2, $Y_{i2}(a_{i1}, a_{i2})$ depends on the vaccine status of both persons. This allows for the possibility that the exposure status of one person affects the outcome of another. In the statistics literature, this is sometimes referred to as interference or spillover effect.^{4,10,16–19}

Most literature in causal inference makes a "no interference" assumption^{16,20} that one person's outcome does *not* depend on the exposure of others. In the current context this would imply that $Y_{i1}(a_{i1}, a_{i2}) = Y_{i1}(a_{i1})$ and $Y_{i2}(a_{i1}, a_{i2}) = Y_{i2}(a_{i2})$ so that each person's outcome depends only on his or her own exposure status. This type of no-interference assumption is implausible in the infectious disease context, and so we do not make it here. We do, however, assume that the exposure status of persons in one household in the study do not affect the outcomes of those in other study households; this is sometimes referred to as an assumption of partial inference.^{10,17} This might be plausible if the various households are sufficiently geographically separated or do not interact with one another. Throughout this paper we will assume a simple randomized experiment in which one of the two persons is randomized to receive a vaccine or control and the second person is always unvaccinated. However, we note in the Discussion section, that a similar analysis would be applicable if the second person were always vaccinated, or in a trial in which the second person were randomized to vaccination and the clusters in which the second person were and were not vaccinated were analyzed separately. In the Discussion section, we also briefly consider settings with multiple people per household. We will let $j = 1$ denote the individual who may or may not be vaccinated and $j = 2$ the individual who is always unvaccinated.

Using this counterfactual notation, the average indirect effect would then simply be

$$E[Y_{i2}(1, 0) - Y_{i2}(0, 0)]$$

i.e. the difference in infection status for person 2 if person 1 is vaccinated versus

unvaccinated.¹⁰ If vaccine status is randomized then this can be estimated simply by^{3,7}:

$$E[Y_{i2}|A_{i1} = 1, A_{i2} = 0] - E[Y_{i2}|A_{i1} = 0, A_{i2} = 0].$$

Halloran and Hudgens³ also refer to this as the "ITT (intention to treat) indirect effect."

To proceed with decomposing this indirect effect into a contagion effect and an infectiousness effect we need also to consider counterfactuals of a different form. From this point onwards, we assume that only person 1, not person 2, can be infected from outside the household; person 2 can only be infected by person 1. Thus if $Y_{i1}(a_{i1}, a_{i2}) = 0$ then $Y_{i2}(a_{i1}, a_{i2}) = 0$. This might be a plausible assumption if person 2 were an infant or a home-bound elderly person. The assumption would also be plausible in certain case-control study designs of rare or acute infections in which it is highly unlikely that both persons are infected from outside the household, but this setting would require further development. Suppose then that in addition to potentially intervening to give person 1 the vaccine we could also, at least hypothetically, consider intervening to give or prevent the infections in person 1. We could then let $Y_{i2}(a_{i1}, a_{i2}, y_{i1})$ denote the infection status of person 2 if we had intervened, possibly contrary to fact, to set the vaccine status of person 1 and person 2 to a_{i1} and a_{i2} respectively and the infection status of person 1 to y_{i1} . This in some sense formalizes, using counterfactual notation, ideas that were proposed by Halloran and Struchiner.²

Because, throughout the paper, we assume that individual 2 is always unvaccinated, we can somewhat simplify the notation above. The counterfactuals $Y_{i1}(a_{i1}, a_{i2})$ and $Y_{i2}(a_{i1}, a_{i2})$, we can write as $Y_{i1}(a_{i1}) := Y_{i1}(a_{i1}, 0)$ and $Y_{i2}(a_{i1}) := Y_{i2}(a_{i1}, 0)$. Note that we are still assuming interference/spillover in that the vaccine of individual 1 affects the outcome of person 2. Likewise the vaccine status of person 2 would affect the outcome of person 1 but in this simple randomized experiment, person 2 always remains unvaccinated. This simple setting in which person 2 always remains unvaccinated also allows us to rewrite the counterfactual $Y_{i2}(a_{i1}, a_{i2}, y_{i1})$ as $Y_{i2}(a_{i1}, y_{i1}) := Y_{i2}(a_{i1}, 0, y_{i1})$. We thus will be considering counterfactuals of the form $Y_{i1}(a_{i1})$, $Y_{i2}(a_{i1})$ and $Y_{i2}(a_{i1}, y_{i1})$. The direct effect of the person 1's vaccine on person 1's outcome is $E[Y_{i1}(1) - Y_{i1}(0)]$; the indirect effect of the person 1's vaccine on person 2's outcome is simply $E[Y_{i2}(1) - Y_{i2}(0)]$. In the next section we will use these counterfactuals to define contagion and infectiousness effects.

Contagion and Infectiousness Effects

Consider now the counterfactual contrast

$$E[Y_{i2}(0, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(0))].$$

This considers what would have happened to person 2 had person 1 been left unvaccinated but if we had set the infection status of person 1 to the level it would have been if person 1 was vaccinated; this is $Y_{i2}(0, Y_{i1}(1))$. The contrast compares this counterfactual to what would have happened to person 2 had person 1 been left unvaccinated but if we had set the infection status of person 1 to the level it would have been if person 1 was unvaccinated; this is $Y_{i2}(0, Y_{i1}(0))$. For this contrast to be non-zero, $Y_{i1}(1)$ and $Y_{i1}(0)$ would have to differ, i.e. the vaccine for person 1 would have to have an effect on the infection status of person 1, and that change in infection for person 1 would have to change the infection status for person 2, even if person 1 had been left unvaccinated. Essentially, the contrast is non-zero if the vaccine prevents person 1 from being infected and preventing person 1 from being infected in turn prevents person 2 from being infected. We thus refer to this counterfactual contrast as a contagion effect.

As noted in the introduction, vaccinating person 1 may prevent the infection of person 2 not simply by preventing person 1 from being infected but also potentially because, even if person 1 is infected, the vaccine may render the infection less infectious. Consider now the contrast

$$E[Y_{i2}(1, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(1))].$$

This compares what would have happened to person 2 if person 1 had been vaccinated versus unvaccinated and had person 1 had the infection status that would have occurred with the vaccine. This contrast will essentially only be non-zero if person 1 is infected with the vaccine (since person 1's vaccination status will not affect person 2's outcome unless person 1 is infected). If the contrast is non-zero then this will be because even when person 1 is vaccinated and infected, the vaccine itself changes whether person 2 is infected by person 1. This is thus one way to measure what in the infectious disease literature is referred to as the "infectiousness effect." Other counterfactual formalizations of the infectiousness effect have also been proposed,^{3,7} and we will discuss the relation of these measures to that proposed above at the end of the paper.

These counterfactual definitions of the contagion and infectiousness effects have the desirable feature that we can decompose an indirect effect into a contagion and an infectiousness effect essentially by taking the indirect effect

and adding and subtracting the term $E[Y_{i2}(0, Y_{i1}(1))]$:

$$\begin{aligned} E[Y_{i2}(1) - Y_{i2}(0)] &= E[Y_{i2}(1, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(0))] \\ &= E[Y_{i2}(1, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(1))] \\ &\quad + E[Y_{i2}(0, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(0))] \end{aligned}$$

where the first term in the sum is the infectiousness effect and the second term in the sum is the contagion effect. This decomposition is analogous to what in the mediation analysis literature is sometimes referred to as "natural direct and indirect effects"^{12,13}. We will in fact be exploiting this analogy in subsequent sections in our discussion of identification, estimation, and sensitivity analysis. The term "indirect effect" is used differently in mediation analysis than in causal inference vaccine literature on interference. In the mediation analysis literature, "indirect effect" is used to describe the effect of an exposure on an outcome for one individual that operates through some intermediate or mediator in that same individual. This is also referred to as a "mediated effect". In the literature on causal inference in the presence of interference, the indirect effect of say, vaccinating some persons in a population is a contrast of potential outcomes comparing the outcomes in those other persons who did not receive the vaccine to what their outcomes would have been if the vaccinated persons were not vaccinated. The latter notion of an indirect effect in the presence of interference is also called a spillover effect in the social science literature. See the Appendix for further discussion.

Thus far we have been considering measures of effect on a risk difference scale. However, risk ratio, odds ratio, or vaccine efficacy measures are more commonly employed in the vaccine literature. The effects described above and the decomposition described above also have analogues for ratio and vaccine efficacy measures. For example, the indirect effect on the risk ratio and odds ratio could be defined respectively as $\frac{E[Y_{i2}(1)]}{E[Y_{i2}(0)]}$ or $\frac{E[Y_{i2}(1)]/\{1-E[Y_{i2}(1)]\}}{E[Y_{i2}(0)]/\{1-E[Y_{i2}(0)]\}}$. Decompositions for the indirect effect into a contagion and infectiousness effect also hold for the risk ratio or odds ratio. For example, for the risk ratio we have that:

$$\frac{E[Y_{i2}(1)]}{E[Y_{i2}(0)]} = \frac{E[Y_{i2}(1, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(1))]} \times \frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]}.$$

Here the first term in the product is the infectiousness effect on the risk ratio scale and the second term is the contagion effect on the risk ratio scale; the indirect effect is now the product of the contagion and infectiousness effects on the risk ratio scale, rather than their sum. A similar decomposition holds for odds ratio measures.

Likewise, similar definitions and a somewhat analogous decomposition holds with a vaccine efficacy measure. As in Halloran and Hudgens³, the vaccine efficacy measure for the indirect effect would be defined as:

$$VE_{indirect} = 1 - \frac{E[Y_{i2}(1)]}{E[Y_{i2}(0)]}.$$

We might likewise define the vaccine efficacy for the contagion effect and infectiousness effect measures as:

$$VE_{cont} = 1 - \frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]}$$

$$VE_{inf} = 1 - \frac{E[Y_{i2}(1, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(1))]}.$$

Some algebra gives:

$$1 - \frac{E[Y_{i2}(1)]}{E[Y_{i2}(0)]} = \left(1 - \frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]} \right) + \frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]} \left(1 - \frac{E[Y_{i2}(1, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(1))]} \right)$$

and we thus have:

$$VE_{indirect} = VE_{cont} + \left(\frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]} \right) VE_{inf}.$$

In words, the vaccine efficacy measure for the indirect effect is the sum of the vaccine efficacy for the contagion effect and that of the infectiousness effect where the vaccine efficacy of the infectiousness effect is adjusted by the factor $\left(\frac{E[Y_{i2}(0, Y_{i1}(1))]}{E[Y_{i2}(0, Y_{i1}(0))]} \right)$ to account for the fact that when the infectiousness effect operates, the contagion effect has essentially already occurred (the infectiousness effect makes the infection less infectious but this infectiousness effect will not operate if the vaccine in fact prevents person 1 from being infected).

We could likewise define each of these effect measure conditional on covariates C . For example, the contagion and infectiousness effects on the risk ratio scale conditional on covariates $C = c$ would be $\frac{E[Y_{i2}(0, Y_{i1}(1))|C=c]}{E[Y_{i2}(0, Y_{i1}(0))|C=c]}$ and $\frac{E[Y_{i2}(1, Y_{i1}(1))|C=c]}{E[Y_{i2}(0, Y_{i1}(1))|C=c]}$ respectively.

Identification of Contagion and Infectiousness Effects

We have defined the contagion and infectiousness effects in terms of counterfactuals that are not immediately estimable from the data. Although these effects may be of substantive interest, we cannot estimate them without fur-

ther assumptions. In the appendix we draw on results from counterfactual theory to show that the contagion and infectiousness effects as defined above will be identified under the following four assumptions. We assume that data is available on some set of baseline covariates C that may be attributes of person 1 or of person 2 or household-level attributes. More rigorous statements of these assumptions are given in the appendix. We assume that conditional on the set of covariates C the following assumptions hold:

- (i) The effect of person 1's vaccine on person 2's infection status is unconfounded.
- (ii) The covariates C contain all of the common causes of person 1's infection status and person 2's infection status so that the effect of person 1's infection status on person 2's infection status is unconfounded.
- (iii) The effect of person 1's vaccine on person 1's infection status is unconfounded.
- (iv) There is no common cause of person 1's infection status and person 2's infection status that is itself affected by the vaccine.

Under these four assumptions the contagion and infectiousness effects are identified from the data. Empirical formulas for identification are given in the appendix. In the next section we will describe how these effects can be estimated using statistical models.

We now assess the four assumptions in a bit more detail. If the vaccine of person 1 is randomized as we have been assuming throughout then assumptions (i) and (iii) will hold by randomization. In an observational setting assumptions (i) and (iii) would only hold if a sufficiently rich set of covariates C were available so that vaccination was effectively randomized within strata of covariates C .

Assumption (ii) is a strong assumption. Assumption (ii) effectively requires that within the set of available covariates C we have all variables that are common causes of person 1's infection status and person 2's infection status. Such common causes might include for example environmental factors related to the sanitary, spatial and nutritional characteristics of the household. Assumption (ii) can perhaps be made more plausible by attempting to control for such variables, but in general it will not be possible to verify assumption (ii). Assumption (iv) by contrast is arguably somewhat weaker: it requires that of all the common causes of person 1's and person 2's infection status, none of these common causes are affected by the vaccine itself. Since most of these common causes are likely to be characteristics of the household environment, it seems reasonably plausible that such characteristics would not be changed

by the vaccine.

The key to identification of the contagion and infectiousness effects thus arguably lies with trying to ensure the validity of assumption (ii): trying to adjust for covariates that may be common causes of person 1's and person 2's infection status.

Statistical Models to Estimate Contagion and Infectiousness Effects

The previous section described the identification assumptions required for estimating the contagion and infectiousness effects and the appendix gives nonparametric empirical expressions for these effects. Here we consider the use of two logistic regression models to estimate these effects when they are in fact identified.

Suppose that the following two logistic regression models are fit to the observed data, (i) for the probability of infection for person 1 conditional on person 1's vaccine status a_1 and the covariates c and (ii) for the probability of infection for person 2, conditional on person 1's vaccine status a_1 , person 1's infection outcome and the covariates c :

$$\begin{aligned} \text{logit}\{P(Y_1 = 1|a_1, c)\} &= \beta_0 + \beta_1 a_1 + \beta_2' c. \\ \text{logit}\{P(Y_2 = 1|a_1, y_1, c)\} &= \theta_0 + \theta_1 a_1 + \theta_2 y_1 + \theta_3 a_1 y_1 + \theta_4' c. \end{aligned}$$

Note that the model for person 2's infection status allows for potential statistical interaction between the effects of the vaccine status of person 1 and the infection status of person 1. Such interaction would likely be present as the vaccine status of person 1 is unlikely to have an effect on whether person 2 is infected unless person 1 is in fact infected.

In the results that follow we will suppose that the infection outcome for person 2 is sufficiently rare so that odds ratios approximate risk ratios and the logistic link approximates a log link. If the infection outcome for person 2 is not rare then the results given below will hold if the logistic regression model for Y_2 is instead replaced by a log-linear model but the model for Y_1 is kept as a logistic model. No rare outcome assumption or log-linear model is needed for Y_1 .

If the covariates C suffice in satisfying assumptions (i)-(iv) above, and the models above are correctly specified then, as shown in the appendix, the contagion effect on the risk ratio scale conditional on the covariates $C = c$ is given by:

$$\frac{E[Y_{i2}(0, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(0))|c]} = \frac{(1 + e^{\beta_0 + \beta_2' c})(e^{\beta_0 + \beta_1 + \beta_2' c + \theta_0 + \theta_2 + \theta_4' c} + e^{\theta_0 + \theta_4' c})}{(1 + e^{\beta_0 + \beta_1 + \beta_2' c})(e^{\beta_0 + \beta_2' c + \theta_0 + \theta_2 + \theta_4' c} + e^{\theta_0 + \theta_4' c})} \quad (1)$$

and the infectiousness effect on the risk ratio scale conditional on the covariates is given by:

$$\frac{E[Y_{i2}(1, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(1))|c]} = \frac{e^{\theta_1(1 + e^{\beta_0 + \beta_1 + \beta'_2 c + \theta_2 + \theta_3})}}{(1 + e^{\beta_0 + \beta_1 + \beta'_2 c + \theta_2})}. \quad (2)$$

These expressions can be obtained directly from the estimates of the logistic regression parameters and standard errors for these could be obtained using the delta method. In the appendix we discuss adapting SAS and SPSS macros for mediation analysis¹⁵ to compute these contagion and infectiousness effects as well as standard errors and confidence intervals for these parameters.

Sensitivity Analysis for Contagion and Infectiousness Effects

The identification and estimation of the contagion and infectiousness effects depend critically on assumptions (i)-(iv) above. Unfortunately, these are fairly strong assumptions, especially assumption (ii) that the set of observed covariates C contains all common causes of the infection status of person 1 and person 2. In this section we give a relatively straightforward sensitivity analysis technique that can be employed to assess how sensitive one's estimates and conclusions are to violations of assumption (ii). The technique assumes that there is an unmeasured binary confounding variable U that is a common cause of the infection status of person 1 and person 2, and that assumptions (i)-(iv) would hold conditional on (C, U) but not on the measured covariates C alone. The investigator can then specify sensitivity parameters corresponding to (i) the effect of the unmeasured confounding U on the infection status of person 2 conditional on the vaccine status of person 1, the infection status of person 1, and the observed covariates C and (ii) the prevalence of U within each stratum defined by the vaccine status of person 1 and the infection status of person 1, conditional on the observed covariates C . The technique then uses the estimates obtained by controlling only for observed covariates C along with these sensitivity parameters to calculate the corrected estimates that would have been obtained had it been possible to control for the unmeasured confounding variable U as well. The sensitivity analysis parameters can then be varied across a range of plausible values to assess how sensitive the conclusions and estimates are to a potential unmeasured common cause of the infection status of person 1 and person 2.

The technique assumes that the effect of U on the infection status of person 2 is constant across the vaccine status of person 1 and the infection status of

person 1 and is given by

$$\gamma = \frac{P(Y_2 = 1|a_1, y_1, c, U = 1)}{P(Y_2 = 1|a_1, y_1, c, U = 0)}.$$

The sensitivity analysis parameter γ thus captures the effect of U on the infection status of person 2. The investigator also specifies the prevalence of U in each stratum defined by the vaccine status of person 1 and the infection status of person 1 conditional on the observed covariates C :

$$\pi_{rs} = P(U = 1|a_1 = r, y_1 = s, c).$$

From these sensitivity analysis parameters the following can be calculated

$$\begin{aligned} B_0 &= \frac{1 + (\gamma - 1)\pi_{10}}{1 + (\gamma - 1)\pi_{00}} \\ B_1 &= \frac{1 + (\gamma - 1)\pi_{11}}{1 + (\gamma - 1)\pi_{01}} \\ B_2 &= \frac{1 + (\gamma - 1)\pi_{01}}{1 + (\gamma - 1)\pi_{00}} \end{aligned}$$

It follows from derivations in VanderWeele²¹ that if we let

$$\begin{aligned} \theta_1^\dagger &= \theta_1 - \log(B_0) \\ \theta_2^\dagger &= \theta_2 - \log(B_2) \\ \theta_3^\dagger &= \theta_3 - \log(B_1) + \log(B_0) \end{aligned}$$

and replace $(\theta_1, \theta_2, \theta_3)$ with $(\theta_1^\dagger, \theta_2^\dagger, \theta_3^\dagger)$ in formulas (1) and (2) then this would give corrected contagion and infectiousness effect estimates corresponding to what would have been obtained had we been able to adjust for U and C rather than only the observed covariates C alone. In general we will not know the true values of the sensitivity analysis parameters; however, by varying the parameters γ and $\pi_{00}, \pi_{10}, \pi_{01}, \pi_{11}$ we will be able to have some sense as to how sensitive the results are to potential unmeasured common causes of the infection status of person 1 and person 2. The sensitivity technique is of course also limited by the assumptions made which are (i) a single unmeasured binary confounder and (ii) that the effect of U on the infection status of person 2 is constant across the vaccine status of person 1 and the infection status of person 1.

Illustration

Consider the hypothetical vaccine trial data in Table 1 in which person 1 is randomized to the vaccine, there are two persons per household, outcome data is available on both persons, and information is also available on household level socioeconomic status.

Table 1. Numbers infected, (Y_{i1}, Y_{i2}) , from a hypothetical randomized vaccine trial, by vaccination status (A_{i1}, A_{i2}) and socioeconomic status (SES)

	$Y_{i1} = 0, Y_{i2} = 0$	$Y_{i1} = 1, Y_{i2} = 0$	$Y_{i1} = 1, Y_{i2} = 1$
Low SES: $A_{i1} = 0, A_{i2} = 0$	200	120	180
Low SES: $A_{i1} = 1, A_{i2} = 0$	350	96	54
High SES: $A_{i1} = 0, A_{i2} = 0$	250	125	125
High SES: $A_{i1} = 1, A_{i2} = 0$	400	75	25

If we fit a logistic model for the probability of infection for person 1 conditional on person 1's vaccine status a_1 and the covariates c and a log-linear model for the probability of infection for person 2, conditional on person 1's vaccine status a_1 , person 1's infection outcome and the covariates c and then use the expressions (1) and (2) above for the contagion and infectiousness effects, setting the covariate to its mean value, we obtain, under assumptions (i)-(iv), an overall estimate of the indirect effect on the risk ratio scale of 0.63 (95% CI: 0.56, 0.70), an estimate of the contagion effect on the risk ratio scale of 0.80 (95% CI: 0.74, 0.85) and an estimate of the infectiousness effect on the risk ratio scale of 0.79 (95% CI: 0.71, 0.87). The indirect effect on the risk ratio scale decomposes into the product of the contagion and infectiousness effects: $0.63 = 0.80 \times 0.79$. On the vaccine efficacy scale, we would have an overall indirect effect vaccine efficacy of $1 - 0.63 = 37\%$, a contagion effect vaccine efficacy of $1 - 0.80 = 20\%$, an infectiousness effect vaccine efficacy of $1 - 0.79 = 21\%$, and vaccine efficacy component due to the infectiousness effect of $(0.80)(21\%) = 17\%$ (essentially taking into account the fact that the infectiousness effect will operate only if the contagion effect has not). We can then decompose the indirect effect vaccine efficacy into the sum of the contagion effect vaccine efficacy and the vaccine efficacy component due to infectiousness: $37\% = 20\% + 17\%$. In this hypothetical example, roughly equal portions of the indirect effect of person 1's vaccine on person 2's infection status appear to be due to the contagion effect versus the infectiousness effect.

Discussion

In this paper we have considered how an indirect effect of the vaccine of one person on the outcome of another can be decomposed into two components:

one corresponding to the vaccine preventing the infection in person 1 which then protects person 2 (the contagion effect) and another corresponding to the fact that even if person 1 is infected the vaccine may render the infection less infectious (the infectiousness effect). The infectiousness effect has been considered in other work in the vaccine literature^{6,22,23} and within causal inference.^{3,7} Halloran and Hudgens³ and VanderWeele and Tchetgen Tchetgen⁷ formalize the infectiousness effect by examining the effect of the vaccine of person 1 on the infection status of person 2 in the "principal stratum"²⁴ in which person 1 would be infected irrespective of vaccine status. Issues of inference for this "infectiousness effect" are described elsewhere.^{3,7} This infectiousness effect based on principal strata is somewhat different than that considered here: essentially the "principal stratum" infectiousness effect is a conditional effect (it conditions on the subgroup for which person 1 would be infected irrespective of vaccine status), whereas as the infectiousness effect considered here is an unconditional infectiousness effect - it averages over also those clusters for whom person 1 is uninfected (for which any potential infectiousness effect of the vaccine would not have the opportunity to operate). These issues are important in the interpretation of these effects; both types of infectiousness effects (conditional and unconditional) could potentially be reported. The advantage of the infectiousness effect given in this paper (the unconditional version) is that it can be used to decompose the overall effect into the contagion and infectiousness components.

Our work here could be extended in a number of directions and is also subject to various limitations which could likewise prompt further research on this topic. First, we have considered the setting in which there are two persons per cluster and only one person is randomized to vaccination. However, in settings in which both are randomized to vaccination, the same type of analysis as that described here could be pursued separately for households in which person 2 is or is not vaccinated. Another simple extension to the work here might involve settings in which only one person in each household is randomized to the vaccine but outcome data is collected on numerous additional individuals per household. In such cases the outcome Y_{i2} in this paper could be replaced with the proportion in the household who are infected (other than the person randomized); the logistic regression would then have to be replaced with a linear or log-linear regression but similar methods from the mediation analysis literature could potentially be adapted and applied.¹⁵ If the numbers in each household vary across households, this number could also be controlled for in the analysis. The approach described here could perhaps be further extended to other settings, possible including the use of network data. One limitation of the approach described here is that the analysis assumed that the regression

models were correctly specified. In settings with a large number of covariates this may be a difficult assumption to make plausible. Future research could consider adapting robust statistical methods from the mediation analysis literature²⁵ to help deal with this issue of model specification. Another limitation is that we have assumed that only person 1, not person 2, can be infected from outside of the household. While this may be plausible if person 2 is elderly or an infant (or perhaps in certain case-control study designs for rare infections), in many other settings the assumption will not be plausible. However, once this assumption is relaxed so that person 2 can be infected outside the household then it will be possible for person 2 to infect person 1 and then the temporal ordering between the infection status of person 1 and person 2 is no longer clear, rendering it difficult to utilize the methodology employed here. Future research could consider extending the current methodology to settings in which both persons can be infected outside the household by using data on the timing of infections.³

Halloran⁵ proposed the minicommunity design to estimate indirect effects of vaccination. In the minicommunity design, the household or other small transmission unit serves as the cluster in which to estimate indirect effects of vaccination, similar to studies in larger communities to estimate indirect, total, and overall effects. In some individually randomized, controlled vaccine trials, it may be straightforward to enroll households of trial participants for follow up; (see Trollfors et al.²⁶ for an example. A similar suggestion, called the augmented study design, was made by Longini et al.²⁷ and Datta et al.²⁸ to estimate vaccine efficacy for infectiousness in HIV vaccine trials. Such studies would be relatively cost-effective to conduct. Because establishing that vaccination can have indirect effects and estimating the effects of vaccination on reducing infectiousness for others could have important implications for global vaccine policy, it is important to consider collecting outcome data on other household members in vaccine trials. Such studies would allow estimation of the indirect, contagiousness and infectiousness effects described in this paper.

Appendix

Formalizations and Derivations

In this appendix we give a formal statement of the identification assumption (i)-(iv) in the text, provide non-parametric empirical expressions for the contagion and infectiousness effects when they are identified, derive closed form expressions for these when logistic or log-linear regression models are used to model the probabilities of infection and provide a sensitivity analysis technique

when the identification assumptions are violated. Most of this is accomplished by noting an analytic relation between the contagion and infectiousness effects defined in the text and what are sometimes called "natural direct and indirect effects" in the literature on mediation^{12-15,21}. Within the mediation analysis literature, the interest lies in assessing the extent to which the effect of an exposure A on outcome Y is mediated by some intermediate M . Essentially if within the mediation context, we take the exposure as person 1's vaccine status, the mediator as person 1's infection status, and the outcome as person 2's infection status, then the contagion and infectiousness effects defined in this paper correspond to the "total" natural direct effect and the "pure" natural indirect effect respectively in the mediation analysis literature.^{12,14,29}

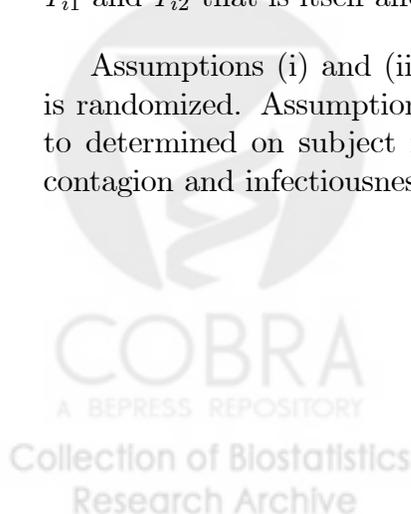
We first formalize identification assumptions (i)-(iv) above. We use $X \perp\!\!\!\perp Y|Z$ to denote that X is conditionally independent of Y given Z . In counterfactual notation, assumptions (i)-(iv) in the text can be formally stated as:

- (i) $Y_{i2}(a_{i1}, y_{i1}) \perp\!\!\!\perp A_{i1}|C$
- (ii) $Y_{i2}(a_{i1}, y_{i1}) \perp\!\!\!\perp Y_{i1}|(C, A_{i1})$
- (iii) $Y_{i1}(a_{i1}) \perp\!\!\!\perp A_{i1}|C$
- (iv) $Y_{i2}(a_{i1}, y_{i1}) \perp\!\!\!\perp Y_{i1}(a'_{i1})|C$

Drawing on the analogy with the mediation analysis literature, the interpretation of (i)-(iv) above is essentially that^{13,14}:

- (i) The effect of A_{i1} on Y_{i2} is unconfounded conditional on C
- (ii) The effect of Y_{i1} on Y_{i2} is unconfounded conditional on (C, A_{i1})
- (iii) The effect of A_{i1} on Y_{i1} is unconfounded conditional on C
- (iv) Given that (ii) holds, there is no confounder of the relationship between Y_{i1} and Y_{i2} that is itself affected by A_{i1}

Assumptions (i) and (iii) will hold if A_{i1} , the vaccine status of person 1, is randomized. Assumptions (ii) and (iv) are substantial and would have to be determined on subject matter grounds. Under assumptions (i)-(iv), the contagion and infectiousness effects are identified from the vaccine trial data.



To see this, note that:

$$\begin{aligned}
E[Y_{i2}(a_{i1}, Y_{i1}(a'_{i1}))|c] &= \sum_{y_1} E[Y_{i2}(a_{i1}, y_1)|Y_{i1}(a'_{i1}) = y_1, c]P(Y_{i1}(a'_{i1}) = y_1|c) \\
&= \sum_{y_1} E[Y_{i2}(a_{i1}, y_1)|c]P(Y_{i1}(a'_{i1}) = y_1|c) \\
&= \sum_{y_1} E[Y_{i2}(a_{i1}, y_1)|a_{i1}, c]P(Y_{i1}(a'_{i1}) = y_1|a'_{i1}, c) \\
&= \sum_{y_1} E[Y_{i2}(a_{i1}, y_1)|a_{i1}, y_1, c]P(Y_{i1}(a'_{i1}) = y_1|a'_{i1}, c) \\
&= \sum_{y_1} E[Y_{i2}|a_{i1}, y_1, c]P(Y_{i1} = y_1|a'_{i1}, c)
\end{aligned}$$

where the first equality holds by iterated expectations, the second by assumption (iv), the third by assumptions (i) and (iii), the fourth by assumption (ii) and the final equality holds by what is sometimes referred to as "consistency". The final expression is given in terms of the observed data. If we first let $a_{i1} = 0, a'_{i1} = 1$ and then $a_{i1} = 0, a'_{i1} = 0$, we obtain that the contagion effect conditional on C is given by $E[Y_{i2}(0, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(0))|c] =$

$$\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]\{P(Y_{i1} = y_1|A_{i1} = 1, c) - P(Y_{i1} = y_1|A_{i1} = 0, c)\}.$$

If we first let $a_{i1} = 1, a'_{i1} = 1$ and then $a_{i1} = 0, a'_{i1} = 1$, we obtain that the infectiousness effect is given by $E[Y_{i2}(1, Y_{i1}(1)) - Y_{i2}(0, Y_{i1}(1))|c] =$

$$\sum_{y_1} \{E[Y_{i2}|A_{i1} = 1, y_1, c] - E[Y_{i2}|A_{i1} = 0, y_1, c]\}P(Y_{i1} = y_1|A_{i1} = 1, c).$$

The contagion effect then essentially contrasts the observed expectation $E[Y_{i2}|A_{i1} = 0, y_1, c]$ as standardized by the distribution of the infection status of person 1 among the households with person 1 vaccinated versus unvaccinated. The infectiousness effect then effectively is the observed expectation contrast $E[Y_{i2}|A_{i1} = 1, y_1, c] - E[Y_{i2}|A_{i1} = 0, y_1, c]$ standardized by the distribution of the infection status of person 1 among the households with person 1 vaccinated.

Likewise on a risk ratio scale we have that the contagion effect is given by:

$$\frac{E[Y_{i2}(0, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(0))|c]} = \frac{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)}{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 0, c)}$$

and the infectiousness effect is given by:

$$\frac{E[Y_{i2}(1, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(1))|c]} = \frac{\sum_{y_1} E[Y_{i2}|A_{i1} = 1, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)}{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)}.$$

Suppose now that the following two models were fit to the data:

$$\begin{aligned} \log it\{P(Y_1 = 1|a_1, c)\} &= \beta_0 + \beta_1 a_1 + \beta'_2 c. \\ \log it\{P(Y_2 = 1|a_1, y_1, c)\} &= \theta_0 + \theta_1 a_1 + \theta_2 y_1 + \theta_3 a_1 y_1 + \theta'_4 c \end{aligned}$$

and that the infection outcome Y_2 for person 2 is sufficiently rare so that odds ratios approximated risk ratios (and the logit link approximated a log-link). Using these models for the conditional predicted probabilities for Y_1 and Y_2 gives, for the contagion effect:

$$\begin{aligned} \frac{E[Y_{i2}(0, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(0))|c]} &= \frac{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)}{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 0, c)} \\ &\approx \frac{e^{\theta_0 + \theta_2 + \theta'_4 c} \frac{e^{\beta_0 + \beta_1 + \beta'_2 c}}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}} + e^{\theta_0 + \theta'_4 c} \frac{1}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}}}{e^{\theta_0 + \theta_2 + \theta'_4 c} \frac{e^{\beta_0 + \beta'_2 c}}{1 + e^{\beta_0 + \beta'_2 c}} + e^{\theta_0 + \theta'_4 c} \frac{1}{e^{\beta_0 + \beta'_2 c}}} \\ &= \frac{(1 + e^{\beta_0 + \beta'_2 c})(e^{\beta_0 + \beta_1 + \beta'_2 c + \theta_0 + \theta_2 + \theta'_4 c} + e^{\theta_0 + \theta'_4 c})}{(1 + e^{\beta_0 + \beta_1 + \beta'_2 c})(e^{\beta_0 + \beta'_2 c + \theta_0 + \theta_2 + \theta'_4 c} + e^{\theta_0 + \theta'_4 c})} \end{aligned}$$

and for the infectiousness effect:

$$\begin{aligned} \frac{E[Y_{i2}(1, Y_{i1}(1))|c]}{E[Y_{i2}(0, Y_{i1}(1))|c]} &= \frac{\sum_{y_1} E[Y_{i2}|A_{i1} = 1, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)}{\sum_{y_1} E[Y_{i2}|A_{i1} = 0, y_1, c]P(Y_{i1} = y_1|A_{i1} = 1, c)} \\ &\approx \frac{e^{\theta_0 + \theta_1 + \theta_2 + \theta_3 + \theta'_4 c} \frac{e^{\beta_0 + \beta_1 + \beta'_2 c}}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}} + e^{\theta_0 + \theta_1 + \theta'_4 c} \frac{1}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}}}{e^{\theta_0 + \theta_2 + \theta'_4 c} \frac{e^{\beta_0 + \beta_1 + \beta'_2 c}}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}} + e^{\theta_0 + \theta'_4 c} \frac{1}{1 + e^{\beta_0 + \beta_1 + \beta'_2 c}}} \\ &= \frac{e^{\theta_1} (1 + e^{\beta_0 + \beta_1 + \beta'_2 c + \theta_2 + \theta_3})}{(1 + e^{\beta_0 + \beta_1 + \beta'_2 c + \theta_2})}. \end{aligned}$$

If the infection outcome for person 2 is not rare then the results above will

hold if the logistic regression model for Y_2 is replaced by a log-linear model but the model for Y_1 is kept as a logistic model. No rare outcome assumption or log-linear model is needed for Y_1 . Standard errors and confidence intervals for these expressions can be obtained via the delta method as in Valeri and VanderWeele.¹⁵ In fact, the SAS and SPSS macros in Valeri and VanderWeele¹⁵ can be directly adapted to estimate these effects and their standard errors and confidence intervals by: specifying the exposure as the vaccine status of person 1, the mediator as the infection status of person 1, the outcome as the infection status of person 2, the outcome model as logistic (or log-linear if the infection outcome for person 2 is not rare), the mediator model as logistic and requesting the option that the full output be given. The estimates reported for the "pure natural indirect effect" can then be taken as a measure of the contagion effect on the conditional risk ratio scale and that reported for the "total natural direct effect" can be taken as the measure of the infectiousness effect on the conditional risk ratio scale. The macro provides standard errors and confidence intervals for these estimates. The formal analytic relation between natural direct and indirect effects and the contagion and infectiousness effects also allows us to adapt sensitivity analysis techniques for natural direct and indirect effects²¹ to apply to contagion and infectiousness effects as in the text.

A few further technical comments merit attention. VanderWeele and Tchetgen Tchetgen⁷ provided an alternative definition of the infectiousness effect on a ratio scale as $E[Y_{i2}(1, Y_{i1}(1)) | Y_{i1}(1) = Y_{i1}(0) = 1] / E[Y_{i2}(0, Y_{i1}(0)) | Y_{i1}(1) = Y_{i1}(0) = 1]$ i.e. the effect of the vaccine of person 1 on the infection status of person 2 in the "principal stratum"²⁴ in which person 1 would be infected irrespective of vaccine status. This infectiousness effect is "conditional" in the sense that it is conditional on the subgroup for which person 1 would be infected irrespective of vaccine status, whereas the infectiousness effect in the text is an unconditional infectiousness effect - it averages over also those households for whom person 1 is uninfected. Yet another definition of an infectiousness effect could be given as $E[Y_{i2}(1, 1) | c] / E[Y_{i2}(0, 1) | c]$ i.e. the effect of the vaccine of person 1 on the outcome of person 2, intervening to set person 1's infection status to present. This effect is analogous to the "controlled direct effect"^{12,13} in the mediation literature. Under assumptions (i) and (ii) above it is identified by $E[Y_{i2} | A_{i1} = 1, Y_{i1} = 1, C = c] / E[Y_{i2} | A_{i1} = 0, Y_{i1} = 1, C = c]$. It is a marginal effect insofar as it is for the entire population for which $C = c$; however it is "conditional on infection" in the sense that it considers a hypothetical contrast in which, in all households, person 1 is infected. Under the logistic regression models given above (assuming rare outcome or using a log-linear model rather than logistic model for Y_{i2}), this would be $e^{\theta_1 + \theta_3}$. It should also be noted that under the exclusion restriction that the vaccine of person

1 does not affect the infection status of person 2 unless person 1 is infected we would have $E[Y_{i2}(1, 0)|c]/E[Y_{i2}(0, 0)|c] = 1$. Under assumptions (i) and (ii) and the two regression models we then have $1 = E[Y_{i2}(1, 0)|c]/E[Y_{i2}(0, 0)|c] = E[Y_{i2}|A_{i1} = 1, Y_{i1} = 0, C = c]/E[Y_{i2}|A_{i1} = 0, Y_{i1} = 0, C = c] = e^{\theta_1}$ i.e. $\theta_1 = 0$. This implication of the exclusion restriction can be tested empirically.

A Note on Terminology

In this paper we have exploited relations between what we have defined as the "contagion and infectiousness effects" on the one hand and "natural direct and indirect effects" on the other. Because of the terminological overlap, the language employed can be somewhat confusing. In the mediation analysis literature¹³, "indirect effect" is used to describe situations in which the effect of an exposure on an outcome for one person operates through some intermediate or mediator for that individual. The "contagion effect" and "infectiousness effect" in this paper are, analytically somewhat analogous to the "natural indirect effect" and "natural direct effect", respectively, in the mediation analysis literature. The "contagion effect" is essentially the effect of person 1's vaccine on person 2's infection outcome mediated by person 1's infection outcome. The "infectiousness effect" is essentially the effect of person 1's vaccine on person 2's infection outcome not mediated by person 1's infection outcome.

In the infectious disease and vaccine literature, the "indirect effect of vaccination" has one more general usage and also a more technical meaning.¹ In general, an "indirect effect of vaccination" is used to describe settings in which vaccination of one person affects the outcome of another individual. This is a specific case of the dependent happenings described by Sir Ronald Ross³⁰ wherein the number of events depends on how many others are already affected. However, in the causal inference literature for vaccine effects, there are several effects of vaccination strategies due to the interference between individuals², wherein the treatment assignment of one person affects the potential outcomes of other persons.^{16,19,20} In this literature, the indirect effect of vaccination is the effect of a vaccination strategy in a population in those individuals, or a subpopulation of those individuals, who were not vaccinated. The total effect of vaccination is the effect of a vaccination strategy in a population in those individuals, or a subpopulation of those individuals, who were vaccinated. More recent formal papers on these indirect and total effects in the presence of interference include Hudgens and Halloran¹⁰, VanderWeele and Tchetgen Tchetgen¹¹, and Tchetgen Tchetgen and VanderWeele.⁴ In other statistical and causal inference literature the effects due to interference are sometimes called "spillover effects".¹⁷

In this paper, we have decomposed the "indirect effect of a vaccination" in the literature on causal inference in the presence of interference into the "natural indirect effect" and "natural direct effect" of mediation analysis. Because these two literatures, causal inference in the presence of interference on the one hand and causal inference mediation analysis on the other hand - use the same terms for different concepts, and moreover because, as we have seen in this paper, these concepts are not entirely unrelated, it is important to clarify in each instance how specifically the various terms are being employed.

References

1. Halloran ME, Struchiner CJ. Study designs for dependent happenings. *Epidemiology* 1991; 2:331-338.
2. Halloran ME, Struchiner CJ. Causal inference for infectious diseases. *Epidemiology* 1995; 6:142-51.
3. Halloran ME, Hudgens MG. Causal inference for vaccine effects on infectiousness. *International Journal of Biostatistics*, in press.
4. Tchetgen Tchetgen EJ, VanderWeele TJ. Estimation of causal effects in the presence of interference. *Statistical Methods in Medical Research - Special Issue on Causal Inference*, in press.
5. Halloran ME. The minicommunity for assessing indirect effects of vaccination. *Epidemiologic Methods*, in press.
6. Datta S, Halloran ME, Longini IM. Efficiency of estimating vaccine efficacy for susceptibility and infectiousness: randomization by individual versus household. *Biometrics*, 1999; 55:792-798.
7. VanderWeele TJ, Tchetgen Tchetgen EJ. Bounding the infectiousness effect in vaccine trials. *Epidemiology* 2011;22:686-693.
8. Aral S, Muchnik L, Sundararajan A (2009). Distinguishing influence based contagion from homophily driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences (USA)*, 106:21544-21549.
9. Christakis NA, Fowler JH. Social contagion theory: examining dynamic social networks and human behavior. *Statistics in Medicine*, to appear.
10. Hudgens MG, Halloran ME. Towards causal inference with interference. *Journal of the American Statistical Association*, 103:832-842, 2008.

11. VanderWeele TJ, Tchetgen Tchetgen EJ. Effect partitioning under interference for two-stage randomized vaccine trials. *Statistics and Probability Letters* 2011;81:861-869.
12. Robins JM, Greenland S. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 1992;3:143-155.
13. Pearl J. Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty and Artificial Intelligence*. San Francisco: Morgan Kaufmann, 2001:411-420.
14. VanderWeele TJ, Vansteelandt S. Odds ratios for mediation analysis with a dichotomous outcome. *American Journal of Epidemiology*, 2010;172:1339-1348.
15. Valeri L, VanderWeele TJ. Mediation analysis allowing for exposure-mediator interactions and causal interpretation: theoretical assumptions and implementation with SAS and SPSS macros. Technical Report.
16. Rubin DB. Comment on: Neyman (1923) and Causal Inference in Experiments and Observational Studies. *Statistical Science* 1990;5:472-480.
17. Sobel ME. 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.
18. Hong G, Raudenbush SW. Evaluating kindergarten retention policy: A case study of causal inference for multilevel observational data. *Journal of the American Statistical Association*, 2006; 101:901-910.
19. Rosenbaum PR. Interference between units in randomized experiments. *Journal of the American Statistical Association*, 2007; 102:191-200.
20. Cox DR. *The Planning of Experiments*. New York: Wiley, 1958.
21. VanderWeele TJ. Bias formulas for sensitivity analysis for direct and indirect effects. *Epidemiology*, 2010;21:540-551.
22. Halloran ME, Préziosi M-P, Chu H. Estimating vaccine efficacy from secondary attack rates. *Journal of the American Statistical Association* 2003; 98:38-46.

23. Becker NG, Britton T, O'Neill PD, Estimating vaccine effects from studies of outbreaks in household pairs. *Statistics in Medicine*, 2006; 25:1079-1093.
24. Frangakis CE, Rubin DB. Principal stratification in causal inference. *Biometrics*, 2002;58:21-29.
25. Tchetgen Tchetgen EJ, Spitzer I. Semiparametric estimation of models for natural direct and indirect effects. Harvard University Biostatistics Working Paper Series. Working Paper 129. <http://www.bepress.com/harvardbiostat/paper129>, 2011.
26. Trollfors B, Taranger J, Lagergard T, Sundh V, Bryla DA, Schneerson R, Robbins JB. Immunization of children with pertussis toxoid decreases spread of pertussis within the family. *Pediatr Infect Dis J* 1998;17:196-199.
27. Longini IM, Datta S, Halloran, ME. Measuring vaccine efficacy for both susceptibility to infection and reduction in infectiousness for prophylactic HIV-1 vaccines. *J Acq Immun Def Synd* 1996;13:440-447.
28. Datta S, Halloran ME, Longini IM. Augmented HIV vaccine trial designs for estimating reduction in infectiousness and protective efficacy. *Stat Med* 1998;17:185-200.
29. Robins JM. Semantics of causal DAG models and the identification of direct and indirect effects. In *Highly Structured Stochastic Systems*, Eds. P. Green, N.L. Hjort, and S. Richardson, 70-81. Oxford University Press, New York, 2003.
30. Ross R. An application of the theory of probabilities to the study of *a priori* pathometry, Part 1. *Proc R Soc Series A* 1916;92:204-230.

