

Causal Vaccine Effects on Binary Postinfection Outcomes

Michael G. HUDGENS and M. Elizabeth HALLORAN

The effects of vaccine on postinfection outcomes, such as disease, death, and secondary transmission to others, are important scientific and public health aspects of prophylactic vaccination. As a result, evaluation of many vaccine effects condition on being infected. Conditioning on an event that occurs posttreatment (in our case, infection subsequent to assignment to vaccine or control) can result in selection bias. Moreover, because the set of individuals who would become infected if vaccinated is likely not identical to the set of those who would become infected if given control, comparisons that condition on infection do not have a causal interpretation. In this article we consider identifiability and estimation of causal vaccine effects on binary postinfection outcomes. Using the principal stratification framework, we define a postinfection causal vaccine efficacy estimand in individuals who would be infected regardless of treatment assignment. The estimand is shown to be not identifiable under the standard assumptions of the stable unit treatment value, monotonicity, and independence of treatment assignment. Thus selection models are proposed that identify the causal estimand. Closed-form maximum likelihood estimators (MLEs) are then derived under these models, including those assuming maximum possible levels of positive and negative selection bias. These results show the relations between the MLE of the causal estimand and two commonly used estimators for vaccine effects on postinfection outcomes. For example, the usual intent-to-treat estimator is shown to be an upper bound on the postinfection causal vaccine effect provided that the magnitude of protection against infection is not too large. The methods are used to evaluate postinfection vaccine effects in a clinical trial of a rotavirus vaccine candidate and in a field study of a pertussis vaccine. Our results show that pertussis vaccination has a significant causal effect in reducing disease severity.

KEY WORDS: Causal inference; Infectious disease; Maximum likelihood; Principal stratification; Sensitivity analysis.

1. INTRODUCTION

The success of vaccines in reducing the burden of infectious diseases is one of the great achievements in public health. The goal of vaccination is often not to prevent infection, but rather to prevent or to ameliorate disease (Clements-Mann 1998). Therefore, many questions of interest in understanding the effects of prophylactic vaccines concern postinfection outcomes, such as severe disease, death, or transmission to others. As early as 1939, Kendrick and Eldering described less severe disease in children inoculated with pertussis vaccine. More recently, Préziosi and Halloran (2003a) estimated the beneficial effects of pertussis vaccination on reducing transmission to others (see also Halloran, Préziosi, and Chu 2003) and preventing severe disease (Préziosi and Halloran 2003b) in breakthrough cases. Individuals with varicella (chickenpox) vaccination also develop less severe disease (Vazquez et al. 2001). Assessing candidate human immunodeficiency virus (HIV) vaccine effects on progression to AIDS through both clinical and surrogate postinfection endpoints (Gilbert, Bosch, and Hudgens 2003a) as well as on transmission (Longini, Datta, and Halloran 1996) is crucial for evaluating their potential public health benefit.

Two main approaches are currently used for estimating postinfection effects of vaccines. The first approach includes all individuals under study, whether they become infected or not; the second includes only the individuals who become infected. The first approach enjoys the statistical validity associated with an intent-to-treat (ITT) analysis and provides an assessment of the overall benefit of vaccination. However, such an approach does not distinguish vaccine effects on susceptibility to infection from effects on the postinfection endpoints of interest. Also, it may be less sensitive in detecting postinfection

vaccine effects than an analysis that conditions on infection. The second approach conditions on infection to estimate a net effect of the vaccine on the postinfection endpoint in just those people who become infected. However, conditioning on an event, such as infection, that occurs subsequent to receipt of vaccine or control may result in selection bias (Halloran and Struchiner 1995). Because the set of individuals who would become infected if vaccinated is likely not identical to the set of those who would become infected if given control, comparisons that condition on infection do not have a causal interpretation (Frangakis and Rubin 2002). For example, a vaccine might protect only people with strong immune systems, so that infected vaccinees tend to have weaker immune systems on average compared with infected controls. As a result, infected vaccinees could have worse postinfection outcomes on average than the infected controls due solely to selection bias. The resulting false conclusion that vaccination is harmful with respect to postinfection outcomes could result in the discarding of a potentially useful vaccine candidate (Hudgens, Hoering, and Self 2003; Gilbert et al. 2003a).

As a first motivating example, Préziosi and Halloran (2003b) analyzed data from a pertussis vaccine field study in Niakhar, Senegal from January 1 to December 31, 1993, in children age 6 months through 8 years. During that one calendar year, there were 3,845 and 1,020 person-years at risk in the vaccinated and unvaccinated children (M.-P. Préziosi, personal communication). Of 548 cases in the vaccinated group, 176 were severe, and of 206 cases in the unvaccinated group, 129 were severe. Préziosi and Halloran (2003b) defined an estimator of vaccine efficacy against severe disease conditional on being infected and having disease as

$$\widehat{VE}_P = 1 - \frac{\frac{\text{no. severe vaccinated cases}}{\text{no. vaccinated cases}}}{\frac{\text{no. severe unvaccinated cases}}{\text{no. unvaccinated cases}}}. \quad (1)$$

Michael G. Hudgens is Research Assistant Professor, Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599 (E-mail: mhudgens@bios.unc.edu). M. Elizabeth Halloran is Professor, Department of Biostatistics, Emory University, Atlanta, GA 30322 (E-mail: mehallo@sph.emory.edu). This work was supported in part by National Institutes of Health grants R37-AI29168, R01-AI054165, and R01-AI32042. Halloran was also a visiting faculty member in the Statistics Division, Los Alamos National Laboratories. The authors thank Peter Gilbert, Andrew Hill, Marie-Pierre Préziosi, and Bryan Shepherd for helpful discussions.

The analysis yields $\widehat{VE}_p = .49$, 95% confidence interval (CI) = (.40, .56).

Vesikari, Rautanen, Varis, Beards, and Kapikian (1990) provided a second motivating example in their analysis of a randomized, double-blinded, placebo-controlled trial of a *Rhesus* rotavirus candidate vaccine. The trial was conducted in children age 2–5 months in 1985–1987 in Finland, with 100 children randomized to each arm. The effect of the vaccine on the clinical course of infection was considered by comparing severity (mild, moderate, or severe) between vaccinees and placebo-treated individuals with confirmed *Rotavirus* diarrhea using Fisher's exact test. If severe and moderately severe cases were combined, then the foregoing approach of Préziosi and Halloran could also be used. In particular, of 10 cases in the vaccinated group, 5 were severe or moderately severe, and of 16 cases in the placebo group, 13 were severe or moderately severe, such that, using (1), we have $\widehat{VE}_p = .38$, 95% CI = (–.11, .74).

At first glance, both vaccines seem to be protective against severe disease. However, due to the possibility of selection bias, we cannot be sure of this without further investigation. Different methods are used to adjust analyses for posttreatment variables, such as infection (Robins and Greenland 1992, 1994; Rosenbaum 1984). Our approach is based on an extension of the potential outcomes framework for estimating causal effects (Rubin 1978; Holland 1986). In particular, Frangakis and Rubin (2002) and Rubin (2000) proposed a method to adjust for posttreatment variables, called *principal stratification*, that stratifies on the joint potential posttreatment variables (e.g., infection) under each treatment being considered. The causal effects of one treatment compared with the other on a main outcome of interest (e.g., disease severity) are defined within these principal strata and are called *principal effects*. In studying HIV vaccines, Hudgens et al. (2003) and Gilbert et al. (2003b) adopted the principal stratification approach to assess HIV vaccine effects on the continuous postinfection outcome viral load in the principal stratum of individuals who would be infected under vaccine and placebo. However, they focused on testing, and did not consider binary postinfection endpoints.

In this article we use a similar approach, but in contrast to the previous work, here we use principal stratification to define causal vaccine effects on binary postinfection endpoints and focus on estimation and identifiability. We also analyze data from two actual vaccine studies. In Section 2 we formally define the estimands, particularly the causal estimand of interest and two currently used estimands for vaccine effects on postinfection outcomes, and state our three main assumptions of stable unit treatment value (SUTVA), independence, and monotonicity. In Section 3 we show the causal estimand to be not identifiable without further assumptions regarding the selective mechanism of the vaccine. In Section 4 we developed a variety of selection models that identify the causal estimand. We then derive closed-form maximum likelihood estimators (MLEs) under several of these models, including the models that assume maximum possible levels of positive and negative selection bias. Using these approaches, we analyze the rotavirus vaccine candidate and the pertussis vaccine studies.

2. DEFINING VACCINE EFFECTS

Consider a group of individuals, $i = 1, \dots, n$, each of whom potentially receives either vaccine or placebo. Following the notation of Rubin (1990a) and similar to the development of Gilbert et al. (2003b), let the vector $\mathbf{Z} = (Z_1, \dots, Z_n)$ denote the vector of treatment assignments for n individuals with $Z_i = v$ if the i th individual is assigned vaccine and $Z_i = p$ if the i th individual is assigned placebo or control. In this article assignment is assumed to be equivalent to receipt. Denote the potential infection outcome of the i th individual if the study group were assigned \mathbf{Z} as $S_i(\mathbf{Z})$, where $S_i(\mathbf{Z}) = 0$ if uninfected and $S_i(\mathbf{Z}) = 1$ if infected. Then $\mathbf{S}(\mathbf{Z})$ is the n -vector of potential infection outcomes if the assignment vector were \mathbf{Z} . For simplicity, we suppress the dependence on \mathbf{Z} by letting $\mathbf{S} = \mathbf{S}(\mathbf{Z})$. The focus of this work is on evaluating the causal effect of vaccine on a binary outcome Y that occurs after an individual becomes infected, such as transmission to another individual, severe disease, or death. Formally, if $S_i(\mathbf{Z}) = 1$, then we define the potential postinfection outcome $Y_i(\mathbf{Z}, \mathbf{S}) = 1$ if the i th individual would have the worse (i.e., more severe) postinfection outcome of interest given (\mathbf{Z}, \mathbf{S}) , and $Y_i(\mathbf{Z}, \mathbf{S}) = 0$ otherwise. If an individual's potential infection outcome for an assignment is uninfected [i.e., $S_i(\mathbf{Z}) = 0$], then $Y_i(\mathbf{Z}, \mathbf{S})$ is undefined and denoted by $*$.

We assume that the potential outcomes for each individual i are independent of the treatment assignment of other individuals. This assumption, known also as the assumption of no interference between units (Cox 1958), is obviously violated in many infectious disease settings (Rubin 1978, 1990a; Halloran and Struchiner 1995); however, we make it to clarify the problem in its simplest form. We also assume there are only two possible infection outcomes and, at most, two possible postinfection outcomes. These two assumptions are expressed formally as follows.

Assumption 1: Stable unit treatment value assumption (SUTVA) (Rubin 1978; Angrist, Imbens, and Rubin 1996). For any two treatment assignments, $\mathbf{Z} = (Z_1, \dots, Z_n)$ and $\mathbf{Z}' = (Z'_1, \dots, Z'_n)$, the following hold:

- a. If $Z_i = Z'_i$, then $S_i(\mathbf{Z}) = S_i(\mathbf{Z}')$.
- b. If $Z_i = Z'_i$ and $S_i = S'_i$, then $Y_i(\mathbf{Z}, \mathbf{S}) = Y_i(\mathbf{Z}', \mathbf{S}')$.

The SUTVA assumption allows us to write $S_i(Z_i)$ and $Y_i(Z_i, S_i(Z_i))$ rather than $S_i(\mathbf{Z})$ and $Y_i(\mathbf{Z}, \mathbf{S})$. Under SUTVA, $S_i(Z_i) = S_i(Z'_i)$ whenever $Z_i = Z'_i$, so we can write $Y_i(Z_i)$ instead of $Y_i(Z_i, S_i(Z_i))$. Finally, we let S_i^{obs} denote the observed infection outcome $S_i(v)$ or $S_i(p)$, depending on treatment assignment, and analogously, Y_i^{obs} for the observed postinfection outcome.

We define a *basic principal stratification*, P_0 , according to the joint potential infection outcomes $S^{P_0} = (S(v), S(p))$ (Frangakis and Rubin 2002). Table 1 summarizes the four basic principal strata defined by the joint potential infection outcomes, $(S(v), S(p))$, and the strata defined by the joint potential postinfection outcomes, $(Y(v), Y(p))$, within each principal stratum. The four basic principal strata are composed of immune (not infected under both vaccine and placebo), harmed (infected under vaccine but not placebo), protected (infected under placebo but not vaccine), and doomed individuals

Table 1. Basic Principal Stratification P_0 Based on the Potential Infection Outcomes ($S(v), S(p)$) With Potential Postinfection Strata Based on ($Y(v), Y(p)$)

Potential infection strata		Potential postinfection strata	
Basic principal stratum, S^{P_0}	Potential infection outcomes ($S(v), S(p)$)	Potential postinfection outcomes ($Y(v), Y(p)$)	Postinfection interpretation
Immune	(0, 0)	(* , *)	Always undefined
Harmed	(1, 0)	(0, *) (1, *)	Not severe vaccine, undefined placebo Severe vaccine, undefined placebo
Protected	(0, 1)	(* , 0) (* , 1)	Undefined vaccine, not severe placebo Undefined vaccine, severe placebo
Doomed	(1, 1)	(0, 0) (1, 0) (0, 1) (1, 1)	Never severe Harmed by vaccine Helped by vaccine Always severe

(infected under both vaccine and placebo). Because membership in a basic principal stratum is not affected by whether an individual is actually assigned to vaccine or placebo, the strata can be used in the same way as pretreatment covariates, with causal postinfection vaccine effects defined within a basic principal stratum S^{P_0} .

In general, causal effects are defined in terms of potential outcomes. Table 1 shows that the doomed basic principal stratum, $S^{P_0} = (1, 1)$, is the only stratum in which both potential postinfection endpoints, and thus their joint distribution, are defined. For this reason, in our setting, defining individual postinfection causal vaccine effects makes sense only in the doomed basic principal stratum, $S^{P_0} = (1, 1)$. In other words, we can speak of a vaccine causing an improvement or worsening of a postinfection outcome only for an individual who would become infected whether vaccinated or not. Thus, following Rubin (2000), two population-level causal estimands can be validly defined: (1) the effect of vaccine on infection (S) for all participants, and (2) the effect of vaccine on the binary postinfection outcome (Y) for those participants who would be infected under both treatment assignments.

To define the first causal estimand, the individual causal effect of vaccination on infection in individual i might be expressed as a difference or ratio (Rubin 1990b),

$$S_i(p) - S_i(v) \quad \text{or} \quad \frac{S_i(v)}{S_i(p)}. \quad (2)$$

Vaccine efficacy is typically expressed in terms of a relative risk difference, so that, similar to Struchiner and Halloran (1996), we define the individual causal effect of vaccination on infection as

$$\text{VE}_{Si} = \frac{S_i(p) - S_i(v)}{S_i(p)} = 1 - \frac{S_i(v)}{S_i(p)}, \quad (3)$$

with the convention $\text{VE}_{Si} = 0$ when $S_i(v) = S_i(p) = 0$. Recall that we assume there is perfect compliance; that is, assignment to vaccine (control) is equivalent to receipt of vaccine (control). Without this assumption, (2) and (3) would be causal effects of *assignment* to vaccine. The ramifications of noncompliance on the assessment of causal vaccine effects are not considered here.

Because for any individual at least one potential outcome is always unobserved, the individual causal effect of vaccine

on infection cannot be observed. Therefore, one must rely on contrasts between different individuals to estimate average or population-level causal effects (Holland 1986). Given the individual causal effect (3), similar to Struchiner and Halloran (1996) and Gilbert et al. (2003b), we define the corresponding population causal vaccine effect by

$$\text{VE}_S = \frac{E\{S(p) - S(v)\}}{E\{S(p)\}} = \frac{E\{S(p)\} - E\{S(v)\}}{E\{S(p)\}} = 1 - \frac{E\{S(v)\}}{E\{S(p)\}}.$$

We call this the *relative average causal effect* (RACE) of vaccination on infection. Because $S(v)$ and $S(p)$ are binary, the expectations can be replaced by probabilities, and VE_S can be interpreted as the relative reduction in the probability of infection given vaccine compared with control.

Estimation of population average causal effects requires specification of the assignment mechanism (Rubin 1990b), that is, how the individuals were assigned to vaccine or placebo. We assume that treatment assignment is independent of the potential outcomes.

Assumption 2. Z is independent of $\{Y(v), Y(p), S(v), S(p)\}$.

Here and in the sequel, $Y(z) = *$ if $S(z) = 0$ for $z = v, p$. Randomization is one assignment mechanism where the treatment assignment is independent of the potential outcomes. It now follows

$$\text{VE}_S = 1 - \frac{E\{S(v)|Z=v\}}{E\{S(p)|Z=p\}} = 1 - \frac{E\{S^{\text{obs}}|Z=v\}}{E\{S^{\text{obs}}|Z=p\}}.$$

Next we define three estimands regarding the effect of vaccination on the postinfection outcome Y . Usually, the notation VE_I is used specifically for vaccine efficacy for infectiousness, that is, the effect of vaccination on transmission to others (Halloran, Struchiner, and Longini 1997), and VE_P is used for vaccine efficacy for progression or severity as in (1). Here we use VE_I throughout the development to refer to vaccine efficacy for any binary postinfection outcome. This enables us to use standard notation for the secondary attack rate (SAR), the probability of transmission from an infected to a susceptible individual. In the development that follows, SAR is defined more generally as the proportion of infected individuals with the worse postinfection outcome.

One current approach to assessing vaccine effects on postinfection endpoints defines a *net* vaccine effect estimand that conditions on infection, that is,

$$\begin{aligned} \text{VE}_I^{\text{net}} &= 1 - \frac{E\{Y^{\text{obs}}|S^{\text{obs}} = 1, Z = v\}}{E\{Y^{\text{obs}}|S^{\text{obs}} = 1, Z = p\}} \\ &= 1 - \frac{E\{Y(v)|S(v) = 1\}}{E\{Y(p)|S(p) = 1\}}, \end{aligned}$$

with the second equality following from Assumption 2. However, in general VE_I^{net} does not have a causal interpretation, because the set of individuals with $S(v) = 1$ is not necessarily identical to the set of individuals with $S(p) = 1$ (Rosenbaum 1984; Frangakis and Rubin 2002).

Another current approach defines an estimand for the effect of vaccination on different levels of severity of disease or mortality without conditioning on being infected. Such an estimand might be considered ITT, because it does not condition on the posttreatment variable S^{obs} ; that is, it incorporates all individuals according to their treatment assignment. An example of an ITT estimand is

$$\text{VE}_I^{\text{ITT}} = 1 - \frac{E\{Y(v) \times S(v)\}}{E\{Y(p) \times S(p)\}},$$

where we adopt the convention $Y(z) \times S(z) = 0$ if $S(z) = 0$, $z = v, p$. This is a general form for what Préziosi and Halloran (2003b) called “ VE_S for severity.” Although VE_I^{ITT} does have a causal interpretation, it combines vaccine effects on susceptibility and postinfection outcome as shown by the following relation:

$$\text{VE}_I^{\text{ITT}} = 1 - (1 - \text{VE}_S)(1 - \text{VE}_I^{\text{net}}).$$

We propose a third estimand for the effect of vaccination on a binary postinfection outcome that has a causal interpretation and is separate from vaccine effects on susceptibility. To define the causal VE_I , we use the basic principal stratification shown in Table 1. In particular, the individual causal vaccine effect on the postinfection outcomes is defined as

$$\text{VE}_{Ii} = 1 - \frac{Y_i(v)}{Y_i(p)},$$

for individuals within the doomed principal stratum only. Following the development of VE_S earlier, we define the population postinfection causal vaccine effect VE_I within the doomed principal stratum as

$$\text{VE}_I = 1 - \frac{E\{Y(v)|S^{P_0} = (1, 1)\}}{E\{Y(p)|S^{P_0} = (1, 1)\}}. \quad (4)$$

Like VE_S , (4) could equivalently be given in terms of probabilities, because the postinfection random variables $Y(v)$ and $Y(p)$ are assumed to be binary, such that VE_I can be interpreted as the causal estimand measuring the relative reduction in the probability of the worse postinfection outcome given vaccine compared with placebo in those individuals who would be infected under either treatment assignment. The challenge in estimating principal effects such as (4) lies in identifying those individuals in the doomed basic principal stratum, because at most one potential infection outcome is ever observed for a particular individual. To help overcome this obstacle, we make the following monotonicity assumption.

Assumption 3. Monotonicity. The individual causal effect of vaccination on infection is nonnegative; that is, $\text{VE}_{Si} \geq 0$ for all i .

Under this assumption, an individual who would become infected under vaccine also would become infected under placebo, so the harmed principal stratum $S^{P_0} = (1, 0)$ is empty. In the context of possible noncompliance with treatment assignment, Angrist et al. (1996) and Imbens and Rubin (1997) made a monotonicity assumption akin to Assumption 3, which immediately eliminates the basic principal stratum of defiers (persons who would always take the opposite of their treatment assignment). By Assumption 3, we are not assuming that the individual postinfection causal effect of vaccination is nonnegative; rather, we allow that vaccination could have a negative (i.e., harmful) causal effect on the postinfection outcome. Indeed, one of our goals is to identify the conditions under which it could be determined the postinfection causal effect is negative.

Under Assumption 3, for each possible observed combination of (Z_i, S_i^{obs}) , we can deduce to which basic principal stratum an individual must belong:

$$\begin{aligned} (Z_i = p, S_i^{\text{obs}} = 0) &\rightarrow S^{P_0} = (0, 0) \quad (\text{because } \text{VE}_{Si} \geq 0), \\ (Z_i = p, S_i^{\text{obs}} = 1) &\rightarrow S^{P_0} = (0, 1) \quad \text{or} \quad S^{P_0} = (1, 1), \\ (Z_i = v, S_i^{\text{obs}} = 0) &\rightarrow S^{P_0} = (0, 1) \quad \text{or} \quad S^{P_0} = (0, 0), \\ (Z_i = v, S_i^{\text{obs}} = 1) &\rightarrow S^{P_0} = (1, 1) \quad (\text{because } \text{VE}_{Si} \geq 0). \end{aligned}$$

Thus the infected vaccine recipients will immediately provide information about the distribution of the postinfection outcome $Y(v)$ in the doomed principal stratum. However, getting a handle on $Y(p)$ in $S^{P_0} = (1, 1)$ is still problematic, because infected placebo recipients may be members of either the protected or the doomed principal stratum. In Section 4 we derive models that allow us to identify the distribution of $Y(p)$ in the doomed principal stratum and, in turn, the postinfection causal vaccine effect VE_I given in (4).

3. IDENTIFIABILITY AND ESTIMATION OF VACCINE EFFECTS

In this section we investigate identifiability and estimation of the estimands defined in the previous section under Assumptions 1–3 only.

3.1 Parameterization

Let the parameters θ govern the probabilities associated with the basic principal stratum. In our case the stratum $S^{P_0} = (1, 0)$ is empty, so let $\theta = (\theta_{00}, \theta_{01}, \theta_{11})$, where

$$\Pr\{S^{P_0} = (i, j); \theta\} = \theta_{ij} \quad \text{for } i, j = 0, 1; i \leq j. \quad (5)$$

Next, let the parameters $\phi = (\phi_{00}, \phi_{01}, \phi_{10}, \phi_{11})$ govern the probabilities associated with the joint potential postinfection outcomes in the doomed basic principal stratum $S^{P_0} = (1, 1)$, where

$$\Pr\{(Y(v), Y(p)) = (k, m)|S^{P_0} = (1, 1); \phi\} = \phi_{km} \quad \text{for } k, m = 0, 1. \quad (6)$$

Let the parameters $\boldsymbol{\gamma} = (\gamma_0, \gamma_1)$ govern the probabilities associated with the two possible potential postinfection outcomes under placebo in the protected basic principal stratum, $S^{P_0} = (0, 1)$, where

$$\Pr\{Y(p) = i | S^{P_0} = (0, 1); \boldsymbol{\gamma}\} = \gamma_l \quad \text{for } l = 0, 1. \quad (7)$$

Finally, let the law of Z be given by $\Pr\{Z = z; \boldsymbol{\psi}\} = \psi_z$ for $z = v, p$.

Under this parameterization, the causal estimand of vaccine efficacy for susceptibility is

$$\text{VE}_S = 1 - \frac{\theta_{11}}{\theta_{01} + \theta_{11}}.$$

Based on the definition of the causal estimand VE_I given in (4), we are interested not in the joint probabilities ϕ_{km} ($k, m = 0, 1$), but rather in just two of the marginal probabilities. In particular, under the foregoing parameterization,

$$\text{VE}_I = 1 - \frac{\phi_{1\cdot}}{\phi_{\cdot 1}}, \quad (8)$$

where

$$\Pr\{Y(v) = 1 | S^{P_0} = (1, 1)\} = \phi_{10} + \phi_{11} = \phi_{1\cdot}$$

and

$$\Pr\{Y(p) = 1 | S^{P_0} = (1, 1)\} = \phi_{01} + \phi_{11} = \phi_{\cdot 1}.$$

We also have

$$\text{VE}_I^{\text{net}} = 1 - \frac{\phi_{1\cdot}}{\gamma_1 \text{VE}_S + \phi_{\cdot 1}(1 - \text{VE}_S)}$$

and

$$\text{VE}_I^{\text{ITT}} = 1 - \frac{\phi_{1\cdot}(1 - \text{VE}_S)}{\gamma_1 \text{VE}_S + \phi_{\cdot 1}(1 - \text{VE}_S)}.$$

3.2 Identifiability

Without further assumptions, VE_I is not identifiable from the observable random variables $(Y^{\text{obs}}, S^{\text{obs}}, Z)$. To see this, first write the probability function for the observable random variables $\Pr\{Z = z, S^{\text{obs}} = s, Y^{\text{obs}} = y; \boldsymbol{\theta}, \boldsymbol{\gamma}, \boldsymbol{\phi}, \boldsymbol{\psi}\}$ as

$$\begin{aligned} & \theta_{00} \psi_p I(z = p, s = 0, y = *) \\ & + \{\theta_{01} + \theta_{11} - (\theta_{01} \gamma_1 + \theta_{11} \phi_{\cdot 1})\} \psi_p I(z = p, s = 1, y = 0) \\ & + (\theta_{01} \gamma_1 + \theta_{11} \phi_{\cdot 1}) \psi_p I(z = p, s = 1, y = 1) \\ & + (\theta_{00} + \theta_{01}) \psi_v I(z = v, s = 0, y = *) \\ & + \theta_{11}(1 - \phi_{1\cdot}) \psi_v I(z = v, s = 1, y = 0) \\ & + \theta_{11} \phi_{1\cdot} \psi_v I(z = v, s = 1, y = 1), \end{aligned} \quad (9)$$

where $I(\cdot)$ is the usual indicator function. Next, suppose that we have $\theta_{01} = \theta_{11}$, $\phi_{\cdot 1} = \epsilon$, and $\gamma_1 = 1 - \epsilon$ for some $\epsilon \in (0, 1)$. Then the probability function given in (9) is constant for different values of ϵ . Moreover, $\text{VE}_I = 1 - \phi_{1\cdot}/\epsilon$ depends on the choice of ϵ , indicating that VE_I is not identifiable. In contrast, the estimands VE_S , VE_I^{ITT} , and VE_I^{net} are identifiable without further assumptions.

3.3 Estimation

Suppose that we observe n independent and identically distributed realizations of $(Z, S^{\text{obs}}, Y^{\text{obs}})$, where Y^{obs} is undefined or does not exist if $S^{\text{obs}} = 0$. There are six possible observed combinations of $(Z, S^{\text{obs}}, Y^{\text{obs}})$. Let $n_{sy}(z)$ be the number of each combination observed in the study population, where $s = 0, 1$ is the observed infection outcome S^{obs} ; $y = 0, 1, *$ is the observed postinfection outcome Y^{obs} ; and $z = v, p$. That is,

$$n_{0*}(p) = \sum_i I(Z_i = p, S_i^{\text{obs}} = 0, Y_i^{\text{obs}} \text{ does not exist}),$$

$$n_{10}(p) = \sum_i I(Z_i = p, S_i^{\text{obs}} = 1, Y_i^{\text{obs}} = 0),$$

$$n_{11}(p) = \sum_i I(Z_i = p, S_i^{\text{obs}} = 1, Y_i^{\text{obs}} = 1)$$

and

$$n_{0*}(v) = \sum_i I(Z_i = v, S_i^{\text{obs}} = 0, Y_i^{\text{obs}} \text{ does not exist}),$$

$$n_{10}(v) = \sum_i I(Z_i = v, S_i^{\text{obs}} = 1, Y_i^{\text{obs}} = 0),$$

$$n_{11}(v) = \sum_i I(Z_i = v, S_i^{\text{obs}} = 1, Y_i^{\text{obs}} = 1),$$

where the summations are over $i = 1, \dots, n$. (Note that the double subscripts do not have the same meaning for the n 's as for the ϕ 's and θ 's.) We assume each of the six combinations is observed at least once, that is, $n_{sy}(z) > 0$ for $(s, y) = (0, *), (1, 0), (1, 1)$ and $z = v, p$. Let $n(p) = n_{0*}(p) + n_{10}(p) + n_{11}(p)$ and $n(v) = n_{0*}(v) + n_{10}(v) + n_{11}(v)$ denote the number of individuals assigned to placebo and to vaccine. Let $n_{1\cdot}(p) = n_{10}(p) + n_{11}(p)$ and $n_{1\cdot}(v) = n_{10}(v) + n_{11}(v)$ denote the number of infected individuals assigned to placebo and vaccine. Let

$$\text{AR}_z = \frac{n_{1\cdot}(z)}{n(z)} \quad \text{for } z = v, p;$$

that is, AR_z is the observed attack rate in the group assigned treatment z . Finally, let

$$\text{SAR}_z = \frac{n_{11}(z)}{n_{1\cdot}(z)} \quad \text{for } z = v, p;$$

that is, SAR_z is the observed secondary attack rate in the group infected given treatment z .

MLEs of the identifiable vaccine efficacy estimands can be found by maximizing the likelihood

$$L(\boldsymbol{\theta}, \boldsymbol{\gamma}, \boldsymbol{\phi}) \propto \prod_{i=1}^n \Pr[Y_i^{\text{obs}} = y_i, S_i^{\text{obs}} = s_i | Z_i = z_i; \boldsymbol{\theta}, \boldsymbol{\gamma}, \boldsymbol{\phi}],$$

subject to constraints on $\boldsymbol{\theta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\phi}$ that ensure that (5)–(7) are probability functions. Using the results from the Appendix, it follows that the MLE of VE_S is given by

$$\widehat{\text{VE}}_S = \begin{cases} 1 - \frac{\text{AR}_v}{\text{AR}_p} & \text{if } \text{AR}_v \leq \text{AR}_p \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Further, the MLEs of VE_I^{ITT} and VE_I^{net} are

$$\widehat{VE}_I^{net} = 1 - \frac{SAR_v}{SAR_p} \tag{11}$$

and

$$\widehat{VE}_I^{ITT} = 1 - (1 - \widehat{VE}_S) \frac{SAR_v}{SAR_p} \tag{12}$$

or, equivalently,

$$\widehat{VE}_I^{ITT} = \begin{cases} \widehat{VE}_I^{net} & \text{if } \widehat{VE}_S = 0 \\ 1 - \frac{n_{11}(v)/n(v)}{n_{11}(p)/n(p)} & \text{if } \widehat{VE}_S > 0. \end{cases} \tag{13}$$

The causal estimand VE_I is not identifiable because $\phi_{.1}$, the denominator of the right side of (8), is not identifiable. But $\phi_{1.}$, the numerator of the right side of (8), can be identified by the observable random variables. The corresponding MLE is given by

$$\widehat{\phi}_{1.} = SAR_v, \tag{14}$$

that is, the observed secondary attack rate in the vaccine arm.

Finally, although $\phi_{.1}$ is not identifiable, we can identify

$$\Pr[Y(p) = 1 | S(p) = 1; \theta, \boldsymbol{\gamma}, \boldsymbol{\phi}] = \gamma_1 VE_S + \phi_{.1}(1 - VE_S). \tag{15}$$

The MLE of (15) is SAR_p such that any feasible pair $(\widehat{\gamma}_1, \widehat{\phi}_{.1})$ satisfying

$$SAR_p = \widehat{\gamma}_1 \widehat{VE}_S + \widehat{\phi}_{.1}(1 - \widehat{VE}_S) \tag{16}$$

is an MLE of $(\gamma_1, \phi_{.1})$. In the following section we introduce selection models that identify $(\gamma_1, \phi_{.1})$. These models add an additional constraint to the parameters such that there is only one feasible pair $(\widehat{\gamma}_1, \widehat{\phi}_{.1})$ that satisfies (16). Because the value of the likelihood remains the same, we have the unique MLE of $\phi_{.1}$ and, therefore, VE_I . Before proceeding, we revisit our examples.

3.4 Applications

3.4.1 Rotavirus Candidate Vaccine. In the rotavirus candidate vaccine study (Vesikari et al. 1990), the observed data were

$$\begin{aligned} n_{0*}(p) &= 84, & n_{0*}(v) &= 90, \\ n_{10}(p) &= 3, & n_{10}(v) &= 5, \\ n_{11}(p) &= 13, & n_{11}(v) &= 5. \end{aligned}$$

From (10), $\widehat{VE}_S = 1 - (10/100)/(16/100) = .375$. It then follows from (13) that $\widehat{VE}_I^{ITT} = 1 - (5/100)/(13/100) = .62$. The secondary attack rates are $SAR_v = \widehat{\phi}_{1.} = 5/10 = .50$ and $SAR_p = 13/16 = .81$, such that $\widehat{VE}_I^{net} = 1 - (5/10)/(13/16) = .385$.

To consider estimation of VE_I , Table 2 shows the relation of the observed data to the basic principal strata and the strata of joint potential postinfection outcomes within each basic principal stratum. By Assumptions 1–3, we know the following:

- All $n_{10}(v) + n_{11}(v) = 10$ belong to the doomed stratum $S^{P_0} = (1, 1)$.
- All $n_{0*}(p) = 84$ belong to the immune stratum $S^{P_0} = (0, 0)$.
- The $n_{0*}(v) = 90$ could belong to $S^{P_0} = (0, 0)$ or $S^{P_0} = (0, 1)$.
- The $n_{10}(p) + n_{11}(p) = 16$ could belong to $S^{P_0} = (0, 1)$ or $S^{P_0} = (1, 1)$.

Ignoring statistical variability, by Assumption 2 (independence), because there are 10 vaccine recipients in the doomed stratum, there are 10 placebo recipients in the doomed stratum as well. Because there are 84 placebo recipients in the immune stratum, there are 84 vaccine recipients in the immune stratum as well. Thus there must be 6 recipients from each of the vaccinated and unvaccinated groups in the protected stratum. Thus we can estimate the size of the unobserved principal stratum $S^{P_0} = (0, 1)$. However, we do not know which 6 of the 16 infected placebo recipients are in the protected stratum $S^{P_0} = (0, 1)$ or which 10 of the 16 are in the doomed stratum $S^{P_0} = (1, 1)$. This illustrates the need for further assumptions to identify VE_I .

Table 2. Rotavirus Vaccine Example (Vesikari et al. 1990): Relation of Observed Data to the Basic Principal Strata of the Joint Potential Infection Outcomes and the Strata of Joint Potential Postinfection Outcomes

Basic principal stratum, S^{P_0}	Potential infection outcomes $(S(v), S(p))$	Vaccinated $n(v) = 100$		Potential postinfection outcomes $(Y(v), Y(p))$	Placebo $n(p) = 100$	
		No. in S^{P_0}	Observed data		Observed data	No. in S^{P_0}
Immune	(0, 0)	[84]		(*, *)	$n_{0*}(p) = 84$	84
Protected	(0, 1)	[6]		(*, 0) (*, 1)	$n_{10}(p) = 3$ $n_{11}(p) = 13$	[6]
Doomed	(1, 1)		10	(0, 0) (1, 0) (0, 1) (1, 1)		

NOTE: The shaded regions correspond to membership within the basic principal strata with observations outside the shaded regions having unknown principal stratum. Numbers in [] are not observed, but follow from Assumptions 1–3 and the observed data.

3.4.2 Pertussis Vaccine. The pertussis vaccine analysis of Préziosi and Halloran (2003b) included exactly 1 year of follow-up (the calendar year 1993), so that the person-years at risk are a close approximation to the number of persons at risk. Thus we use the person-years at risk for $n(v)$ and $n(p)$. Although vaccine status was not randomized, there was no evidence of systematic differences between the vaccinated and unvaccinated groups, indicating that Assumption 2 (independence) might be reasonable. Active surveillance to find cases makes it unlikely that ascertainment bias plays a role. (See Préziosi and Halloran 2003b for further discussion.) The observed data are

$$\begin{aligned} n_{0*}(p) &= 814, & n_{0*}(v) &= 3,297, \\ n_{10}(p) &= 77, & n_{10}(v) &= 372, \\ n_{11}(p) &= 129, & n_{11}(v) &= 176. \end{aligned}$$

From (10), $\widehat{VE}_S = 1 - (548/3,845)/(206/1,020) = .29$. The secondary attack rates are $SAR_v = \widehat{\phi}_1 = 176/548 = .32$ and $SAR_p = 129/206 = .63$, such that $\widehat{VE}_I^{\text{net}} = 1 - (176/548)/(129/206) = .49$, which is the same as \widehat{VE}_P given in (1). Finally $\widehat{VE}_I^{\text{ITT}} = 1 - (176/3,845)/(129/1,020) = .64$, which Préziosi and Halloran (2003b) called \widehat{VE}_S for severity.

4. SELECTION BIAS MODELS

In the previous section we showed that VE_I is not identifiable under Assumptions 1–3 (SUTVA, independence, and monotonicity). In particular, the parameterization (5)–(7) defines a family of joint distribution functions of $(Z, S(v), S(p), Y(v), Y(p))$ for which the resulting conditional distributions of $Y(p)|S^{P_0} = (1, 1)$ are not identified from the observable random variables $(Z, S^{\text{obs}}, Y^{\text{obs}})$. In this section we consider different additional assumptions about the selective effect of the vaccine on susceptibility to infection that restrict inference to a subset of this family such that the distribution of $Y(p)|S^{P_0} = (1, 1)$ and, in turn, VE_I are identifiable. The lack of identifiability of VE_I derives from $\phi_{\cdot 1}$ and γ_1 occurring in (9) only through the term

$$\theta_{01}\gamma_1 + \theta_{11}\phi_{\cdot 1} = \Pr\{Y^{\text{obs}} = 1, S^{\text{obs}} = 1|Z^{\text{obs}} = p\}. \quad (17)$$

Thus for any fixed values of θ_{01} , θ_{11} , and $\Pr\{Y^{\text{obs}} = 1, S^{\text{obs}} = 1|Z^{\text{obs}} = p\}$, all pairs of parameters

$$\{(\gamma_1, \phi_{\cdot 1}) : 0 \leq \gamma_1 \leq 1, 0 \leq \phi_{\cdot 1} \leq 1, \text{ and (17) holds}\}, \quad (18)$$

will yield the same distribution of $(Z, S^{\text{obs}}, Y^{\text{obs}})$. The selection models presented in this section place additional constraints on the parameter space such that only one pair of parameters satisfies (18). These selection models are nonparametric in the sense that no additional restrictions are placed on the distribution of $(Z, S^{\text{obs}}, Y^{\text{obs}})$ (Scharfstein, Rotnitzky, and Robins 1999). Although the observed data do not allow us to distinguish among selection models, through sensitivity analyses we can estimate VE_I under each selection model in a collection of models that represent a range of plausible selective effects of vaccine. Moreover, if we can show that VE_I is significantly greater (less) than 0 over all values of possible selective effects, the observed data provide sufficient evidence to conclude the vaccine has a positive (negative) causal effect on the postinfection endpoint under Assumptions 1–3 only.

We begin in Section 4.1 by assuming there is no selective vaccine effect. In Section 4.2 we present two extreme models that assume maximum possible levels of positive and negative selection bias and, in turn, yield bounds on the MLE of the causal estimand in a fashion similar to the approaches of Manski (1990), Balke and Pearl (1997), and Hudgens et al. (2003). In Section 4.3 we present three approaches to sensitivity analysis that allow a range of selective effects. We discuss making inferences on VE_I accounting for statistical variability in Section 4.4, and then apply the different methods to the two vaccine examples in Section 4.5.

4.1 No Selection Bias

The simplest assumption is that there is no selection; that is, the probability of the postinfection outcome conditional on infection under placebo is independent of infection status under vaccine,

$$\begin{aligned} \Pr\{Y(p) = y|S^{P_0} = (1, 1); \phi\} \\ = \Pr\{Y(p) = y|S^{P_0} = (0, 1); \gamma\} \quad \text{for } y = 0, 1, \end{aligned} \quad (19)$$

which implies $\phi_{\cdot 1} = \gamma_1$. It follows immediately that assumption (19) identifies VE_I and, from (16), that the resulting MLE is

$$\widehat{\phi}_{\cdot 1} = SAR_p. \quad (20)$$

From (14), (20), and the definition of VE_I given by (8), it follows that the MLE of VE_I equals $\widehat{VE}_I^{\text{net}}$ as given in (11). In other words, under the additional assumption of no selection bias as specified by (19), the MLE of the causal vaccine effect is the usual secondary attack rate ratio estimator one obtained when conditioning on infection as in (1).

4.2 Upper and Lower Bounds

In this section we present selection models that provide the largest and smallest estimators of VE_I . In particular, the upper bound selection model yields the parameter pair satisfying (18) that has the largest $\phi_{\cdot 1}$ (and thus the largest VE_I). Because (18) is simply the intersection of the unit square and a line with negative slope, it follows that the pair with maximal $\phi_{\cdot 1}$ must be on the edge of the square, that is, either when

$$\Pr\{Y(p) = 1|S^{P_0} = (1, 1)\} = \phi_{\cdot 1} = 1 \quad (21)$$

or

$$\Pr\{Y(p) = 1|S^{P_0} = (0, 1)\} = \gamma_1 = 0. \quad (22)$$

In words, the upper bound selection bias model assumes either that (a) all placebo recipients in the doomed principal stratum have the worse postinfection outcome or (b) all placebo recipients in the protected principal stratum have the better postinfection outcome. Under the model assuming either (21) or (22), $\phi_{\cdot 1}$ is identifiable; the proof is not shown but is available from the authors. From (16), it follows that the unique MLE of VE_I assuming either (21) or (22) is given by

$$\widehat{VE}_I^{\text{upper}} = \begin{cases} 1 - SAR_v & \text{if } \widehat{VE}_S > 1 - SAR_p \\ \widehat{VE}_I^{\text{ITT}} & \text{if } 0 < \widehat{VE}_S \leq 1 - SAR_p \\ \widehat{VE}_I^{\text{net}} & \text{if } \widehat{VE}_S = 0. \end{cases} \quad (23)$$

All MLEs obtained under Assumptions 1–3 must be less than or equal to $\widehat{VE}_I^{\text{upper}}$.

Similarly, the lower bound selection bias model assumes that under assignment to placebo, the worse postinfection outcome occurs either with probability 0 in the doomed principal stratum,

$$\Pr[Y(p) = 1 | S^{P_0} = (1, 1)] = \phi_{\cdot 1} = 0, \tag{24}$$

or with probability 1 in the protected principal stratum,

$$\Pr[Y(p) = 1 | S^{P_0} = (0, 1)] = \gamma_1 = 1. \tag{25}$$

This model renders $\phi_{\cdot 1}$ identifiable, and from (16) the resulting unique MLE of VE_I is

$$\widehat{VE}_I^{\text{lower}} = \begin{cases} -\infty & \text{if } \widehat{VE}_S > SAR_p \\ 1 - SAR_v / \left\{ \frac{SAR_p - \widehat{VE}_S}{1 - \widehat{VE}_S} \right\} & \text{if } 0 < \widehat{VE}_S \leq SAR_p \\ \widehat{VE}_I^{\text{net}} & \text{if } \widehat{VE}_S = 0. \end{cases} \tag{26}$$

In summary, even though VE_I is not identifiable from the data at hand, these extreme selection bias models provide bounds on the MLE. Moreover, we have derived the circumstances when the upper bound will be negative (suggesting harm) and when the lower bound will be positive (suggesting benefit). For example, $\widehat{VE}_I^{\text{upper}}$ will be negative if and only if $\widehat{VE}_S \leq 1 - SAR_p$ and $\widehat{VE}_I^{\text{ITT}} < 0$. Similarly, $\widehat{VE}_I^{\text{lower}}$ will be positive if and only if $\widehat{VE}_S \leq SAR_p$ and $SAR_v < (SAR_p - \widehat{VE}_S)/(1 - \widehat{VE}_S)$. In contrast, for $\widehat{VE}_S > \max\{SAR_p, 1 - SAR_p\}$, we know that $\widehat{VE}_I^{\text{upper}}$ will be always positive and $\widehat{VE}_I^{\text{lower}}$ will be always negative. In other words, for large enough \widehat{VE}_S we cannot determine the sign of VE_I unless further assumptions are made beyond SUTVA, independence, and monotonicity.

4.3 Sensitivity Analyses

Estimating the causal vaccine effect under an extreme degree of selection bias is useful in bounding the estimate of the postinfection effect above and beyond any possible selective effects. However, the true degree of selection bias is likely less than the foregoing extreme models, such that using $\widehat{VE}_I^{\text{upper}}$ or $\widehat{VE}_I^{\text{lower}}$ may be too conservative. Therefore, we consider classes of selection models that include the foregoing no selection and extreme models. Through sensitivity analysis over a particular class, we can explore the relationship between the assumed degree of selection bias and our inference about VE_I . This approach allows for potentially more realistic models and less conservative inference.

4.3.1 Log Odds Ratio of Infection. Our first approach is similar to that of Scharfstein et al. (1999) and Robins, Rotnitzky, and Scharfstein (2000) and adapted by Gilbert et al. (2003b). We define a sensitivity model in terms of the log odds ratio of having the severe postinfection endpoint under placebo in the doomed versus protected principal strata,

$$\exp(\beta) = \frac{\Pr[Y(p) = 1 | S^{P_0} = (1, 1)] / \Pr[Y(p) = 0 | S^{P_0} = (1, 1)]}{\Pr[Y(p) = 1 | S^{P_0} = (0, 1)] / \Pr[Y(p) = 0 | S^{P_0} = (0, 1)]}. \tag{27}$$

For example, if $\exp(\beta) = 2$, then doomed individuals have twice the odds of having the worse postinfection outcome under placebo compared with protected individuals. In terms of our parameterization, this implies that

$$\phi_{\cdot 1} = \frac{\gamma_1 \exp(\beta)}{\gamma_0 + \gamma_1 \exp(\beta)}. \tag{28}$$

For fixed β , we can solve (16) and (28) for $\phi_{\cdot 1}$ and, in turn, VE_I . A sensitivity analysis can then be performed by repeating this process over a range of different β 's.

The selection models given in Sections 4.1 and 4.2 are special cases of the class of logistic selection models given by (27) when $\beta = 0$ (no selection), $\beta \rightarrow \infty$ (upper bound), or $\beta \rightarrow -\infty$ (lower bound). For example, letting $\beta \rightarrow \infty$ implies either

$$\Pr[Y(p) = 1 | S^{P_0} = (1, 1)] \rightarrow 1 \tag{29}$$

or

$$\Pr[Y(p) = 1 | S^{P_0} = (0, 1)] \rightarrow 0. \tag{30}$$

Because (29) implies (21) and (30) implies (22), we recover the extreme upper selection bias model. The relation among the log odds ratio, upper bound, and no selection models are depicted graphically in Figure 1.

4.3.2 Conditioning on γ_1 as the Sensitivity Analysis Parameter. Section 4.3.1 posits a relation between the distributions of the postinfection endpoint under placebo in the doomed and protected strata. Alternatively, we might simply condition on the nuisance parameter γ_1 , which governs the postinfection endpoint distribution in the protected stratum. Assuming that γ_1 is known renders $\phi_{\cdot 1}$ and, hence, VE_I identifiable. From (16),




Figure 1. Graphical Depiction of Selection Bias Models Described in Section 4. The four dotted curves depict the log odds ratio selection model for $\beta = 0, 1, 2, 4$. The $\beta = 0$ model is equivalent to the no selection model. The thick black lines on the edge of the unit square correspond to the upper-bound selection model. The solid thin line with negative slope represents an example of a family of joint distribution functions of $(Z, S(v), S(p), Y(v), Y(p))$ for which all combinations of parameters $(\phi_{\cdot 1}, \gamma_1)$ on this line give rise to the same distribution of the observable random variables $(Z, S^{\text{obs}}, Y^{\text{obs}})$. Each selection model identifies exactly one pair of parameters $(\phi_{\cdot 1}, \gamma_1)$ from this family, rendering $VE_I = 1 - \phi_{\cdot 1} / \phi_{\cdot 1}$ identifiable.

the resulting MLE of VE_I is

$$\widehat{VE}_I = 1 - SAR_v / \left\{ \frac{SAR_p - \gamma_1 \widehat{VE}_S}{1 - \widehat{VE}_S} \right\}, \quad (31)$$

where γ_1 varies between

$$\max \left\{ 0, \frac{SAR_p - (1 - \widehat{VE}_S)}{\widehat{VE}_S} \right\} \leq \gamma_1 \leq \min \left\{ 1, \frac{SAR_p}{\widehat{VE}_S} \right\}, \quad (32)$$

with the left side of (32) giving rise to $\widehat{VE}_I^{\text{upper}}$ and the right side of (32) giving rise to $\widehat{VE}_I^{\text{lower}}$.

4.3.3 Complete-Data Model. A third approach to sensitivity analysis regards the unknown basic principal stratum membership of the infected placebo recipients as missing data and formulates the sensitivity analysis in terms of the complete-data likelihood. The observed data are $n_{10}(p)$ and $n_{11}(p)$. If we could know the basic principal stratum membership, then the complete data would be $n_{10}^d(p)$ and $n_{11}^d(p)$, the number of infected placebo recipients in the doomed stratum with $Y(p) = 0$ and $Y(p) = 1$, and $n_{10}^p(p)$ and $n_{11}^p(p)$, the corresponding number in the protected stratum. Given the complete data, $\phi_{.1}$ becomes identifiable. The complete-data log-likelihood for (θ, ϕ, γ) is given by

$$\begin{aligned} & n_{0*}(p) \log(\theta_{00}) \\ & + n_{10}^p(p) \log(\theta_{01}\gamma_0) \\ & + n_{11}^p(p) \log(\theta_{01}\gamma_1) \\ & + n_{10}^d(p) \log(\theta_{11}(1 - \phi_{.1})) \\ & + n_{11}^d(p) \log(\theta_{11}\phi_{.1}) \\ & + n_{0*}(v) \log(\theta_{00} + \theta_{01}) \\ & + n_{10}(v) \log(\theta_{11}(1 - \phi_{.1})) \\ & + n_{11}(v) \log(\theta_{11}\phi_{.1}), \end{aligned} \quad (33)$$

which factors in terms of θ , ϕ , and γ , say $l(\theta)$, $l(\phi)$, and $l(\gamma)$. Maximizing $l(\phi)$ yields the same MLE $\widehat{\phi}_{.1} = SAR_v$ as before. It also yields the MLE

$$\widehat{\phi}_{.1} = \frac{n_{11}^d(p)}{n_{10}^d(p) + n_{11}^d(p)}, \quad (34)$$

the secondary attack rate under placebo in the doomed stratum. The sensitivity analysis involves estimating VE_I using (34) for all possible complete-data configurations consistent with the assumptions and the constraints implied by the observed data. Molenberghs, Kenward, and Goetghebeur (2001) called this set of point estimates the *region of ignorance*, and called the collection of CIs or other measures of precision together with the region of ignorance the *region of uncertainty*.

4.4 Statistical Variability

Once we assume a particular selection model, we have a finite-dimensional parametric inference problem with a unique MLE. Therefore, conditional on a selection model, we use standard methods to obtain CI estimates for VE_I . For example, CIs can be computed assuming the usual chi-squared limiting distribution of the profile likelihood ratio (Barndorff-Nielsen and Cox 1994). Alternatively, using the observed information

and the delta method, Wald-type CIs for VE_I can be determined. The resulting CIs can then be used to determine a region of uncertainty for any of the sensitivity analyses described earlier. A region of uncertainty that excludes 0 will then imply a statistically significant postinfection causal effect of the vaccine.

4.5 Applications, Continued

4.5.1 Rotavirus Candidate Vaccine. For these data, $\widehat{VE}_S > 1 - SAR_p$, so from (23), $\widehat{VE}_I^{\text{upper}} = 1 - SAR_v = .50$. In contrast, $0 < \widehat{VE}_S \leq SAR_p$, so, from (26),

$$\widehat{VE}_I^{\text{lower}} = 1 - \frac{5/10}{\frac{\{13/16 - (1 - (10/100)(16/100))\}}{(10/100)/(16/100)}} = .29.$$

Figure 2(a) shows the sensitivity analysis of \widehat{VE}_I as a function of the odds ratio e^β as described in Section 4.3.1. Figure 3(a) shows the sensitivity analysis of \widehat{VE}_I as a function of γ_1 as described in Section 4.3.2. From (32), the range for γ_1 is $[\cdot5, 1.0]$ in this example. For these figures, profile likelihood-based CIs are presented; Wald-type CIs gave qualitatively similar results. The vertical dotted line in Figures 2 and 3 corresponds to the assumption of no selection bias. In particular, in Figure 3, $\widehat{\gamma}_1$ corresponds to the MLE of γ_1 under the assumption of no selection bias.

Table 3 gives all possible configurations of the complete data for the placebo arm consistent with the observed data and Assumptions 1–3 as described in Section 4.3.3. On the one hand, the complete data are constrained by the observed data to satisfy $n_{11}(p) = n_{11}^d(p) + n_{11}^p(p) = 13$ and $n_{10}(p) = n_{10}^d(p) + n_{10}^p(p) = 3$. On the other hand, Assumptions 1–3 imply that of the total of 16 infected placebo recipients, 6 must be in the protected stratum and 10 must be in the doomed stratum. There are only four complete-data configurations that satisfy these constraints. These four configurations can be envisaged by assigning the three infected placebo recipients with mild infection to either the doomed stratum or the protected stratum. At one extreme, the three are all in the doomed stratum; at the other extreme, they are all in the protected stratum. The resulting four MLEs of VE_I given in the bottom row of Table 3 compare the region of ignorance. For this dataset, the assumption of no selection bias does not correspond to any possible complete-data configuration. But the extreme lower and upper bound estimates, $\widehat{VE}_I^{\text{lower}}$ and $\widehat{VE}_I^{\text{upper}}$, are included in the region of ignorance.

This example demonstrates the utility of the sensitivity analyses. For these data, the 95% CI for VE_I covers 0 for the lower bound selection model, indicating that one cannot conclude that the vaccine has a protective causal effect against severe disease at the $\alpha = .05$ level. In contrast, the upper bound selection model CI excludes 0, indicating a significant protective effect. In this situation, sensitivity analyses can be used to assess at what point significance is attained. For example, Figures 2(a) and 3(a) indicate that even under the assumption of no selection bias, the effect remains nonsignificant.

Finally, the lack of statistical significance in this example may be attributed simply to small sample size. For instance, had the study been such that there were 1,000/arm with the same observed marginal distributions (i.e., 3% mild and 13% severe in the placebo arm, 5% mild and 5% severe in the vaccine arm),

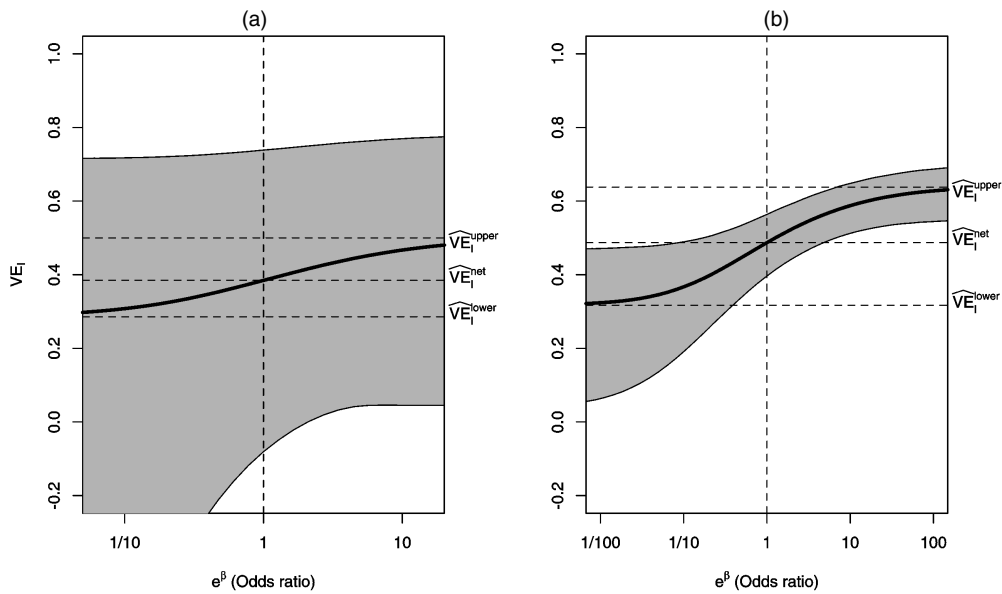


Figure 2. Sensitivity Analysis Using the Odds Ratio of Having the Severe Postinfection Endpoint Under Placebo in the Doomed versus Protected Principal Strata as Described in Section 4.3.1. (a) Rotavirus; (b) pertussis. The vertical dotted line corresponds to the assumption of no selection bias.

then the 95% CI for VE_I under the lower bound model would have been [.09, .46], indicating a significant causal vaccine effect on rotavirus disease severity in individuals who would be infected regardless of vaccination status.

4.5.2 Pertussis Vaccine. For these data, $0 \leq \widehat{VE}_S \leq 1 - SAR_p$, so from (23), $\widehat{VE}_I^{\text{upper}} = \widehat{VE}_I^{\text{ITT}} = .64$. In contrast, $0 \leq \widehat{VE}_S \leq SAR_p$, so from (26), $\widehat{VE}_I^{\text{lower}} = .32$. Figures 2(b) and 3(b) show the sensitivity analyses of \widehat{VE}_I as described in Sections 4.3.1 and 4.3.2. In both figures, the lower limits of the 95% CIs are well above 0 over the range of both selection

models, suggesting that pertussis vaccination causes significant protection against severe disease in children who would develop pertussis regardless of vaccination. The sensitivity analysis using the complete-data likelihood would proceed similarly as for the rotavirus vaccine candidate, with many more possible data configurations and taking the unequal sizes of the vaccinated and unvaccinated arms into account.

5. DISCUSSION

For vaccine studies in which the effect of vaccination on binary postinfection outcomes is of interest, our results will en-

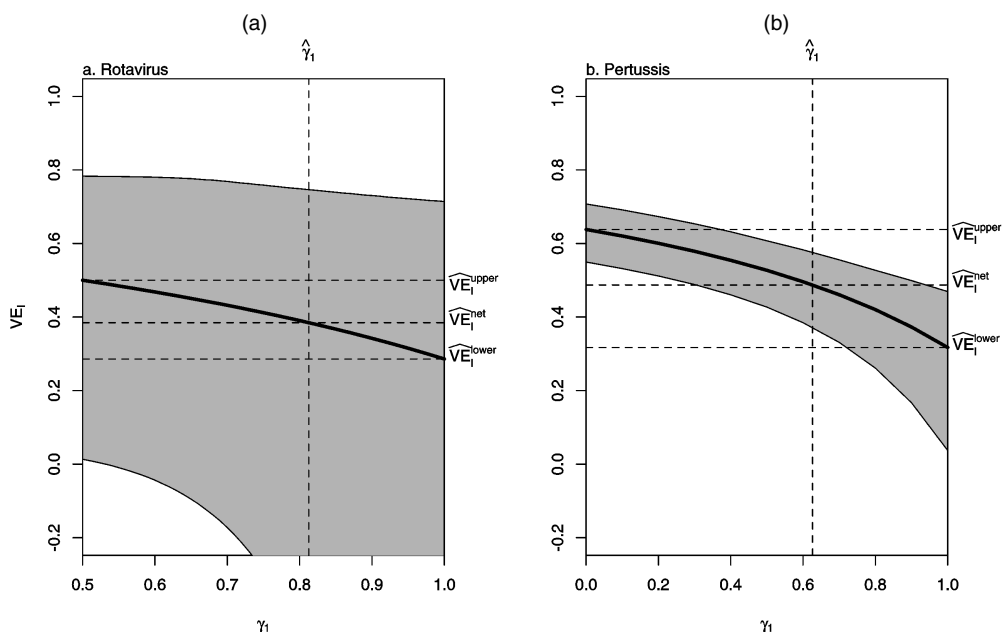


Figure 3. Sensitivity Analysis Assuming $\gamma_1 = \Pr\{Y(p) = 1 | S^p_0 = (0, 1); \gamma\}$ Is Known as Described in Section 4.3.2. (a) Rotavirus; (b) pertussis. The vertical dotted line gives the MLE of γ_1 under the assumption of no selection bias.

Table 3. Rotavirus Vaccine Example (Vesikari et al. 1990): All Possible Configurations of the Underlying Complete Data for the Placebo Arm Consistent With the Observed Data and Assumptions 1–3

Basic principal stratum, S^{P_0}	Potential infection outcomes $(S(v), S(p))$	Potential postinfection outcomes $(Y(v), Y(p))$	Infected placebo recipients observed	No. in S^{P_0} $z = p$	Assumed underlying complete data for the placebo arm							
					No selection bias							
Protected	(0, 1)	(*, 0)	}	$n_{10}(p) = 3$	}	[6] =	0	1	1.125	2	3	
		(*, 1)					6	5	4.875	4	3	
Doomed	(1, 1)	(0, 0)	}	$n_{11}(p) = 13$	}	[10] =	3	2	1.875	1	0	
		(1, 0)					7	8	8.125	9	10	
		(0, 1)										
		(1, 1)										
					\widehat{VE}_I	.29	.375	.38	.44	.50		

NOTE: The shaded regions correspond to membership within the basic principal strata with observations outside the shaded regions having unknown principal stratum. Numbers in [] are not observed, but follow from the observed data and Assumptions 1–3. Maximum likelihood estimates of causal VE_I are given for each possible complete data configuration. The no selection bias model is also included, although it does not correspond to any of the possible complete data configurations.

able researchers to quantify the soundness and limitations of their findings. Combining principal stratification with a maximum likelihood approach, we have developed methods for estimating the causal effect of vaccination on postinfection outcomes in those individuals who would be infected under vaccine and placebo. Closed forms for the MLE of the causal estimand have been derived assuming extreme upper, extreme lower, and no selection bias models. These results indicate specific assumptions under which currently used efficacy estimators have a causal interpretation; for example, the usual net postinfection efficacy estimator is the MLE of our causal estimand under the assumption of no selection bias. In contrast, the usual intent-to-treat estimator provides an upper bound on the postinfection causal vaccine effect when the protective effect of the vaccine against infection is not too large (i.e., $\widehat{VE}_S \leq 1 - SAR_p$). We have also presented three approaches to sensitivity analysis that allow selection models to range from no selection to the extreme maximum possible levels. Coupled with confidence intervals, these sensitivity analyses provide a comprehensive description of the uncertainty regarding VE_I . For example, in the analysis of the pertussis vaccine study, the estimated postinfection vaccine effect was significantly greater than 0 for all selection bias models, including the extreme lower bound model, indicating a significant causal effect in reducing disease severity. In other words, under only the standard assumptions of SUTVA, monotonicity, and independence of treatment assignment, we have shown that there is a definitive causal vaccine effect on severity of pertussis in individuals who would be infected whether vaccinated or not.

Several extensions of this work and avenues for further research merit mention. For instance, further research is needed on the consequences of relaxing several key assumptions. Regarding Assumption 1, SUTVA is clearly violated in many vaccine field studies (Rubin 1990a; Halloran and Struchiner 1995). When SUTVA is violated, an individual can have different potential infection outcomes under the same treatment

assignment. Therefore, defining the basic principal stratum for an individual needs to account for the possible treatment assignments of other individuals. Assumption 2 (independence) clearly holds in the rotavirus example because the trial was randomized, and likely also holds in the pertussis field study, as discussed in Section 3.4.2. In general, the effect of nonrandom allocation of the vaccine can be explored through additional sensitivity analyses and other approaches for observational studies (Rosenbaum 1995). In addition, exclusion restrictions not considered here may be necessary if the independence assumption is not made. Assumption 3 (monotonicity) also likely holds for the rotavirus and pertussis examples, because $\widehat{VE}_S > 0$ in both studies. Relaxing Assumption 3 to allow for a harmful vaccine effect on infection would make identifiability and estimation of the postinfection causal estimand VE_I more difficult. However, a vaccine that increases susceptibility to infection would be unacceptable from a public health standpoint, so that concerns about postinfection selection bias become less relevant. Finally, we have assumed that assignment to a treatment is equivalent to receipt of that treatment. The increased complexity of defining, identifying, and estimating causal effects in the presence of noncompliance (Angrist et al. 1996) would likely necessitate further exclusion restriction assumptions (e.g., given the actual treatment received, assignment to treatment carries no additional information).

Additional research is also needed to incorporate baseline covariates and allow for missing or mismeasured infection or postinfection outcomes. The operating characteristics of conditional and unconditional (i.e., ITT) tests for binary postinfection vaccine effects remain unknown. Our preliminary comparisons (results not shown) suggest a conditional approach [similar to that of Hudgens et al. (2003)] may be more powerful. These results are consistent with the findings of Mehrotra, Li, and Gilbert (2005). Finally, as pointed out by a reviewer, the multivariate literature on characterizing the set of possible joint distributions given fixed marginal or conditional distributions (see,

e.g., Joe 1997) may have direct application in the principal stratification setting.

Correct assessment of vaccine effects on postinfection endpoints has important public health consequences. For example, there is precedent and concern for vaccine-induced immune responses causing disease enhancement (Kliks, Nisalak, Brandt, Wahl, and Burke 1989; Burke 1992; Mascola et al. 1993; Nabel 2001). If the estimated net postinfection vaccine effect suggested a possible harmful effect, then a vaccine could be wrongfully discounted if, in fact, the apparent negative effect were due to selection bias. Our methods could be used to determine whether such an observed effect was consistent with selection bias. In contrast, if the estimated net vaccine effect were positive, then one might be less compelled to distinguish (a) a selective effect of the vaccine on infection from (b) the causal effect of the vaccine in the doomed principal stratum, because either effect is beneficial from a public health perspective. Disentangling (a) from (b) could be quite valuable from a scientific standpoint, however. For example, our analysis of the pertussis vaccine study provides evidence that vaccination has an effect on disease severity above and beyond that attributable to selection bias, suggesting the vaccine produces an immune response that directly hinders severe symptoms, such as paroxysms with whoops. This could suggest directions for further biological research. Another example arises in HIV, where researchers are considering both antibody and T-cell-based vaccines (McMichael, Mwau, and Hanke 2002; Nabel 2001) under the hypothesis that the former will protect from infection and the latter will protect from disease or death. If a combination vaccine proves effective in reducing the risk of morbidity or mortality, then researchers will be interested in understanding the mechanism of protection, that is, to which vaccine component is such an effect attributable. Analysis revealing a significant post-infection causal vaccine effect in the doomed stratum would provide evidence of T-cell-based protection. Separating (a) and (b) also might be helpful in predicting the success of particular vaccination programs (Hayes, Alexander, Bennett, and Couzens 2000).

Although in this article we have handled the postinfection outcomes of transmission and disease severity or death similarly, there are important differences. In particular, vaccine efficacy for infectiousness (VE_I) involves transmission to another susceptible individual, whereas vaccine efficacy for severity or progression (VE_P) involves only the infected individual. Thus, unlike VE_P , VE_I can depend on the characteristics of the individuals susceptible to secondary infection, as well as on the type of contacts between individuals. Also, an infected individual could expose more than one individual, so that there could be several potential postinfection outcomes, one for each exposed individual; for example, the study by Préziosi and Halloran (2003a), more than 20 individuals were exposed to one infected individual in some cases. Alternatively, an infected individual may expose no other susceptibles, and then the potential postinfection outcome would be undefined. Our method needs to be extended to such multiple outcomes and possible dependencies.

In most vaccine studies in humans, whether or not an individual becomes infected is not under the control of the investigator. The potential postinfection outcomes for individuals who

would not become infected under either vaccine or placebo are undefined, or are a priori counterfactual (Frangakis and Rubin 2002), resulting in our defining the postinfection causal estimand in the doomed stratum only. However, there are a few infectious agents that can be used in infection-challenge studies in either human or animal. For these types of studies, a more appropriate approach might be that of Robins and Greenland (1992, 1994), which assumes that the posttreatment variable is controllable; that is, the investigator can force the individuals to take on any possible value of the posttreatment variable. Comparison of the different approaches merits further research.

APPENDIX: MAXIMUM LIKELIHOOD ESTIMATION

Let $\xi = \theta_{01}\gamma_1 + \theta_{11}\phi_{\cdot 1}$ and consider maximizing the log-likelihood,

$$\begin{aligned} n_{0*}(p) \log(\theta_{00}) &+ n_{10}(p) \log(\theta_{01} + \theta_{11} - \xi) \\ &+ n_{11}(p) \log(\xi) \\ &+ n_{0*}(v) \log(\theta_{00} + \theta_{01}) \\ &+ n_{10}(v) \log(\theta_{11}(1 - \phi_{\cdot 1})) \\ &+ n_{11}(v) \log(\theta_{11}\phi_{\cdot 1}), \end{aligned}$$

subject to the constraints

$$\theta_{00} + \theta_{01} + \theta_{11} = 1, \quad (\text{A.1})$$

$$0 \leq \theta_{ij}, \quad i, j = 0, 1; i \leq j, \quad (\text{A.2})$$

$$0 \leq \xi \leq \theta_{01} + \theta_{11}, \quad (\text{A.3})$$

$$0 \leq \phi_{\cdot 1} \leq 1. \quad (\text{A.4})$$

The maximum of $\phi_{\cdot 1}$ is given by $\hat{\phi}_{\cdot 1} = n_{11}(v)/n_{1\cdot}(v)$, such that we are left with the simpler problem of maximizing

$$\begin{aligned} l(\theta, \xi) &= n_{0*}(p) \log(\theta_{00}) \\ &+ n_{10}(p) \log(\theta_{01} + \theta_{11} - \xi) \\ &+ n_{11}(p) \log(\xi) \\ &+ n_{0*}(v) \log(\theta_{00} + \theta_{01}) \\ &+ n_{1\cdot}(v) \log(\theta_{11}), \end{aligned}$$

subject to (A.1)–(A.3).

The function $l(\theta, \xi)$ is strictly concave over the closed, convex parameter space defined by (A.1)–(A.3) such that a unique MLE exists and the (first-order) Kuhn–Tucker conditions are necessary and sufficient for optimality (Fletcher 1987). In particular, define Lagrange multipliers $\lambda_0, \lambda_1, \dots, \lambda_5$ and Lagrangian function

$$\begin{aligned} L(\theta, \xi) &= l(\theta, \xi) + \lambda_0 \left(1 - \sum \theta_{ij}\right) + \lambda_1 \theta_{00} + \lambda_2 \theta_{01} \\ &+ \lambda_3 \theta_{11} + \lambda_4 \xi + \lambda_5 (\theta_{01} + \theta_{11} - \xi), \end{aligned}$$

where the summation is over $i, j = 0, 1$ such that $i \leq j$. Then $(\hat{\theta}, \hat{\xi})$ is an MLE iff it is a feasible point according to (A.1)–(A.3) and satisfies

$$\frac{\partial L(\theta, \xi)}{\partial \theta_{ij}} = 0, \quad i, j = 0, 1; i \leq j, \quad (\text{A.5})$$

$$\frac{\partial L(\theta, \xi)}{\partial \xi} = 0, \quad (\text{A.6})$$

$$\lambda_i \geq 0, \quad i = 1, \dots, 5, \quad (\text{A.7})$$

and

$$\begin{aligned} \lambda_0 \left(1 - \sum \theta_{ij}\right) &= \lambda_1 \theta_{00} = \lambda_2 \theta_{01} = \lambda_3 \theta_{11} = \lambda_4 \hat{\xi} \\ &= \lambda_5 (\theta_{01} + \theta_{11} - \hat{\xi}) \\ &= 0. \end{aligned} \tag{A.8}$$

Under the assumption $n_{sy}(z) > 0$ for all possible combinations of s , y , and z , we must have $\hat{\theta}_{00}$, $\hat{\theta}_{11}$, and $\hat{\xi}$ strictly greater than 0 to maximize $l(\theta, \xi)$, such that $\lambda_1 = \lambda_3 = \lambda_4 = 0$. Similarly, we must also have $\hat{\xi} < \hat{\theta}_{01} + \hat{\theta}_{11}$, such that $\lambda_5 = 0$. Multiplying (A.5) and (A.6) by θ_{ij} and ξ and summing,

$$\sum \theta_{ij} \frac{\partial L(\theta, \xi)}{\partial \theta_{ij}} + \xi \frac{\partial L(\theta, \xi)}{\partial \xi} = 0$$

can be shown to yield $\lambda_0 = n$. It follows that $(\hat{\theta}, \hat{\xi})$ is an MLE iff it satisfies (A.1)–(A.3) and

$$\frac{\partial l(\theta, \xi)}{\partial \theta_{00}} = \frac{n_{0*}(p)}{\theta_{00}} + \frac{n_{0*}(v)}{\theta_{00} + \theta_{01}} = n, \tag{A.9}$$

$$\frac{\partial l(\theta, \xi)}{\partial \theta_{01}} = \frac{n_{10}(p)}{\theta_{01} + \theta_{11} - \xi} + \frac{n_{0*}(v)}{\theta_{00} + \theta_{01}} \leq n, \tag{A.10}$$

$$\frac{\partial l(\theta, \xi)}{\partial \theta_{11}} = \frac{n_{10}(p)}{\theta_{01} + \theta_{11} - \xi} + \frac{n_{10}(v) + n_{11}(v)}{\theta_{11}} = n, \tag{A.11}$$

and

$$\frac{\partial l(\theta, \xi)}{\partial \xi} = \frac{n_{11}(p)}{\xi} - \frac{n_{10}(p)}{\theta_{01} + \theta_{11} - \xi} = 0. \tag{A.12}$$

We can show the MLE of θ has the following closed form that depends on the data. In particular, if

$$\frac{n_{0*}(v)}{n(v)} \geq \frac{n_{0*}(p)}{n(p)},$$

then the MLEs are given by

$$\begin{aligned} \hat{\theta}_{00} &= \frac{n_{0*}(p)}{n(p)}, \\ \hat{\theta}_{01} &= \frac{n_{0*}(v)}{n(v)} - \frac{n_{0*}(p)}{n(p)}, \end{aligned}$$

and

$$\hat{\theta}_{11} = 1 - \frac{n_{0*}(v)}{n(v)}.$$

Otherwise,

$$\begin{aligned} \hat{\theta}_{00} &= \frac{n_{0*}(p) + n_{0*}(v)}{n}, \\ \hat{\theta}_{01} &= 0, \end{aligned}$$

and

$$\hat{\theta}_{11} = 1 - \frac{n_{0*}(p) + n_{0*}(v)}{n}.$$

In either case, the MLE of ξ is given by

$$\hat{\xi} = \frac{n_{11}(p)}{n_{1.}(p)} (\hat{\theta}_{01} + \hat{\theta}_{11}).$$

REFERENCES

Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996), "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 444–455.

Balke, A., and Pearl, J. (1997), "Bounds on Treatment Effects From Studies With Imperfect Compliance," *Journal of the American Statistical Association*, 92, 1171–1176.

Barndorff-Nielsen, O. E., and Cox, D. R. (1994), *Inference and Asymptotics*, New York: Chapman & Hall.

Burke, D. S. (1992), "Human HIV Vaccine Trials: Does Antibody-Dependent Enhancement Pose a Genuine Risk?" *Perspectives in Biology and Medicine*, 35, 511–530.

Clements-Mann, M. L. (1998), "Lessons for AIDS Vaccine Development From Non-AIDS Vaccines," *AIDS Research and Human Retroviruses*, 14 (Suppl. 3), S197–203.

Cox, D. R. (1958), *Planning of Experiments*, New York: Wiley.

Fletcher, R. (1987), *Practical Methods of Optimization*, New York: Wiley.

Frangakis, C. E., and Rubin, D. B. (2002), "Principal Stratification in Causal Inference," *Biometrics*, 58, 21–29.

Gilbert, P. B., Bosch, R., and Hudgens, M. G. (2003a), "Sensitivity Analysis for the Assessment of Causal Vaccine Effects on Viral Load in HIV Vaccine Trials," *Biometrics*, 59, 531–541.

Gilbert, P. B., DeGruttola, V., Hudgens, M. G., Self, S. G., Hammer, S. M., and Corey, L. (2003b), "What Constitutes Efficacy for a Human Immunodeficiency Virus Vaccine That Ameliorates Viremia: Issues Involving Surrogate End Points in Phase 3 Trials," *Journal of Infectious Diseases*, 188, 179–193.

Halloran, M. E., Préziosi, M.-P., and Chu, H. (2003), "Estimating Vaccine Efficacy From Secondary Attack Rates," *Journal of the American Statistical Association*, 98, 38–46.

Halloran, M. E., and Struchiner, C. J. (1995), "Causal Inference in Infectious Diseases," *Epidemiology*, 6, 142–151.

Halloran, M. E., Struchiner, C. J., and Longini, I. M. (1997), "Study Designs for Evaluating Different Efficacy and Effectiveness Aspects of Vaccines," *American Journal of Epidemiology*, 146, 789–803.

Hayes, R. J., Alexander, N. D. E., Bennett, S., and Cousens, S. N. (2000), "Design and Analysis Issues in Cluster-Randomized Trials of Interventions Against Infectious Diseases," *Statistical Methods in Medical Research*, 9, 95–116.

Holland, P. W. (1986), "Statistics and Causal Inference," *Journal of the American Statistical Association*, 81, 945–960.

Hudgens, M. G., Hoering, A., and Self, S. G. (2003), "On the Analysis of Viral Load Endpoints in HIV Vaccine Trials," *Statistics in Medicine*, 22, 2281–2298.

Imbens, G. W., and Rubin, D. B. (1997), "Bayesian Inference for Causal Effects in Randomized Experiments With Noncompliance," *The Annals of Statistics*, 25, 305–327.

Joe, H. (1997), *Multivariate Models and Dependence Concepts*, London: Chapman & Hall.

Kendrick, P., and Eldering, G. (1939), "A Study in Active Immunization Against Pertussis," *American Journal of Hygiene*, Sect. B, 38, 133.

Kliks, S. C., Nisalak, A., Brandt, W. E., Wahl, L., and Burke, D. S. (1989), "Antibody-Dependent Enhancement of Dengue Virus Growth in Human Monocytes as a Risk Factor for Dengue Hemorrhagic Fever," *American Journal of Tropical Medicine and Hygiene*, 40, 444–451.

Longini, I. M., Datta, S., and Halloran, M. E. (1996), "Measuring Vaccine Efficacy for Both Susceptibility to Infection and Reduction in Infectiousness for Prophylactic HIV-1 Vaccines," *Journal of Acquired Immune Deficiency Syndromes and Human Retrovirology*, 13, 440–447.

Manski, C. F. (1990), "Nonparametric Bounds on Treatment Effects," *American Economic Review, Papers and Proceedings*, 80, 319–323.

Mascola, J. R., Mathieson, B. J., Zack, P. M., Walker, M. C., Halstead, S. B., and Burke, D. S. (1993), "Summary Report: Workshop on the Potential Risks of Antibody-Dependent Enhancement in Human HIV Vaccine Trials," *AIDS Research and Human Retroviruses*, 9, 1175–1184.

McMichael, A., Mwau, M., and Hanke, T. (2002), "Design and Tests of an HIV Vaccine," *British Medical Bulletin*, 62, 87–98.

Mehrotra, D. V., Li, X., and Gilbert, P. B. (2005), "Dual-Endpoint Evaluation of Vaccine Efficacy: Application to a Proof-of-Concept Clinical Trial of a Cell-Mediated Immunity-Based HIV Vaccine," BARDS Technical Report 111, Merck Research Laboratories.

Molenberghs, G., Kenward, M. G., and Goetghebeur, E. (2001), "Sensitivity Analysis for Incomplete Contingency Tables: The Slovenian Plebiscite Case," *Applied Statistics*, 50, 15–29.

Nabel, G. J. (2001), "Challenges and Opportunities for Development of an AIDS Vaccine," *Nature*, 410, 1002–1007.

Préziosi, M.-P., and Halloran, M. E. (2003a), "Effects of Pertussis Vaccination on Transmission: Vaccine Efficacy for Infectiousness," *Vaccine*, 21, 1853–1861.

- (2003b), "Effects of Pertussis Vaccination on Severity: Vaccine Efficacy for Clinical Severity," *Clinical Infectious Diseases*, 37, 772–779.
- Robins, J. M., and Greenland, S. (1992), "Identifiability and Exchangeability for Direct and Indirect Effects," *Epidemiology*, 3, 143–155.
- (1994), "Adjusting for Differential Rates of Prophylaxis Therapy for PCP in High- versus Low-Dose AZT Treatment Arms in an AIDS Randomized Trial," *Journal of the American Statistical Association*, 89, 737–749.
- Robins, J. M., Rotnitzky, A., and Scharfstein, D. O. (2000), "Sensitivity Analysis for Selection Bias and Unmeasured Confounding in Missing Data and Causal Inference Models," in *Statistics in Epidemiology, Environment and Clinical Trials*, eds. M. E. Halloran and D. A. Berry, New York: Springer-Verlag, pp. 1–94.
- Rosenbaum, P. R. (1984), "The Consequences of Adjustment for a Concomitant Variable That Has Been Affected by the Treatment," *Journal of the Royal Statistical Society, Ser. A*, 147, 656–666.
- (1995), *Observational Studies*, Berlin: Springer-Verlag.
- Rubin, D. B. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization," *The Annals of Statistics*, 6, 34–58.
- (1990a), "Comment: Neyman (1923) and Causal Inference in Experiments and Observations Studies," *Statistical Science*, 5, 472–480.
- (1990b), "Formal Modes of Statistical Inference for Causal Effects," *Journal of Statistical Planning and Inference*, 25, 279–292.
- (2000), Comment on "Causal Inference Without Counterfactuals," by A. P. Dawid, *Journal of the American Statistical Association*, 95, 435–437.
- Scharfstein, D. O., Rotnitzky, A., and Robins, J. M. (1999), "Adjusting for Non-ignorable Drop-Out Using Semiparametric Nonresponse Models," *Journal of the American Statistical Association*, 94, 1096–1146.
- Struchiner, C. J., and Halloran, M. E. (1996), "Randomization and Baseline Transmission in Vaccine Field Trials," technical report, Emory University, Dept. of Biostatistics.
- Vazquez, M., LaRussa, P. S., Gershon, A. A., Steinberg, S. P., Freudigman, K., and Shapiro, E. D. (2001), "The Effectiveness of the Varicella Vaccine in Clinical Practice," *New England Journal of Medicine*, 344, 955–960.
- Vesikari, T., Rautanen, T., Varis, T., Beards, G. M., and Kapikian, A. Z. (1990), "Rhesus *Rotavirus* Candidate Vaccine: Clinical Trial in Children Vaccinated Between 2 and 5 Months of Age," *American Journal of Diseases of Children*, 144, 285–289.