9

Vaccine Effects on Post-infection Outcomes

9.1 Measures of Vaccine Effects on Post-infection Outcomes

9.1.1 Vaccine efficacy for post-infection outcomes

A post-infection outcome is an outcome that occurs after a person has been infected. As an example, a post-infection outcome of interest could be symptoms or serious disease in those who become infected. In this chapter, we present methods for estimating the effect of vaccination on post-infection outcomes conditional on infection having occurred, denoted VE_P.

VE_P can be measured as one minus the ratio of the mean of the post-infection outcome in the infected vaccinated people and the post-infection outcome in the infected unvaccinated people:

$$\text{VE}_P = 1 - \frac{\frac{\text{vaccinated post-infection outcome}}{\text{infected vaccinated people}}}{\frac{\text{unvaccinated post-infection outcome}}{\text{infected unvaccinated people}}} \tag{9.1}$$

Similarly, if a post-clinical outcome in the clinical cases is of interest, then VE_P can be measured as one minue the ratio of the mean measure of the post-clinical outcome in the vaccinated cases and the post-clinical outcome in the unvaccinated cases:

$$\text{VE}_P = 1 - \frac{\frac{\text{vaccinated post-clinical outcome}}{\text{vaccinated clinical cases}}}{\frac{\text{unvaccinated post-clinical outcome}}{\text{unvaccinated clinical cases}}} \tag{9.2}$$

Throughout the discussion in this chapter, the methods are applicable to post-infection outcomes given infection as well as to post-clinical outcomes given a clinical case. We do not repeat everywhere the result for both situations. In this chapter, if the interest is on post-infection outcomes, then VE_S, VE_{SP},

and VE_P denote vaccine efficacy for susceptibility to infection, vaccine efficacy for susceptibility to the post-infection outcome not conditional on infection, and vaccine efficacy for the post-infection outcome conditional on being infected. If the interest is in some outcome in clinical cases, the vaccine efficacies are defined analogously.

In randomized studies, approaches that use the originally randomized populations in the denominators, in this case VE_S and VE_{SP}, enjoy the statistical validity associated with an intent-to-treat analysis. This validity holds whether the ascertainment is on infection, clinical disease, or severe disease. The efficacy estimate may be higher for more stringent case definitions (Chapter 6), but the statisitical validity of the comparison is not compromised. However, VE_P conditions on infection (clinical disease) to estimate a net effect of the vaccine on the post-infection or post-disease endpoint in just those people who become infected. The infected vaccinated group and the infected control group may not be comparable, so the comparison may not be statistically valid. In the first sections in this chapter, we assume that the comparison is valid. In Sections 9.3 and 9.4, we relax this assumption and show the implications.

9.1.2 Scientific questions of interest

A common question of interest is whether clinical cases in vaccinated people are less severe than clinical cases in unvaccinated people. As early as 1939, Kendrick and Eldering described less severe disease in children inoculated with pertussis vaccine compared to children who had not been inoculated. Children vaccinated against chickenpox develop less severe disease if they do develop clinical symptoms (Vazquez et al. 2001).

In studying malaria vaccine candidates, the density of malaria parasites in the blood or level of anemia are post-infection outcomes of interest. Sometimes the post-infection outcome is a surrogate endpoint for a clinical endpoint of interest. In assessing HIV vaccine effects on duration of progression to clinical AIDS disease, the clinical endpoint of AIDS take years after infection. As a consequence, the post-infection outcomes viral load and CD4 count are used as surrogates for the clinical endpoint of interest.

Although asymptomatic infections have not traditionally been ascertained in most vaccine studies, for understanding the overall public health effects of vaccination programs and for dynamic models, ascertaining people with asymptomatic infections is important. Asymptomatic people may still be infectious for others. It is important to estimate what proportion of infected people remain asymptomatic. Pathogenicity is a measure of the ability of an infectious agent to cause disease in an infected person. Pathogenicity can be measured by the probability of developing disease if infected. If the vaccine reduces the probability of developing symptomatic infection, VE_P measures the vaccine efficacy for pathogenicity.

Transmissibility for others is also a post-infection outcome. Thus, vaccine efficacy for infectiousness, VE_I, is a special case of a vaccine effect on a

post-infection outcome. If measured by level of viral shedding or some other laboratory measure, then VE_I is similar to a VE_P measure, or at least a surrogate measure. If VE_I is measured epidemiologically based on the transmission probability or secondary attack rate in others, it is more complex than simple VE_P measures (Chapters 10 through 12).

Just as with VE_S, different post-infection outcomes can be used to measure VE_P. Depending on the scientific question of interest, the outcome could be dichotomous (0,1), continuous, as with parasite density, or time-to-event, such as the time between ascertaining infection and developing a clinical outcome of interest. Thus, we could differentiate VE_P measures based on different outcomes by the notation, such as $\text{VE}_{P,\lambda}$ analogous to $\text{VE}_{S,\lambda}$ if based on the time-to-event, though we do not present that here. The outcome could be defined conditional on having been infected or conditional on having developed a clinical case, depending on the method of ascertainment. Conditional on developing infection, the outcome may be whether the infected person develops disease or not. It may be whether the infected person develops disease within some time period after infection. Both of these are dichotomous outcomes. If the outcome is the time to an event after infection, an incidence rate or survival analysis that begins with the observation at the time of infection might be appropriate. For continuous or time-to-event post-infection outcomes, the mean, median or some other summary measure in the two groups could be used. The exact form of the VE_P estimator depends on the choice of outcome. Some options for outcomes in VE_P are summarized in Table 9.1.

As an example of VE_P based on a dichotomous outcome, Préziosi and Halloran (2003b) proposed a method of estimating the efficacy of vaccine in reducing the probability of developing severe disease in clinical cases:

$$\text{VE}_P = 1 - \frac{\dfrac{\text{severe vaccinated cases}}{\text{all vaccinated cases}}}{\dfrac{\text{severe unvaccinated cases}}{\text{all unvaccinated cases}}}. \tag{9.3}$$

If the post-infection outcome is dichotomous, we can define the post-infection attack rate (PAR) as the number with the post-infection outcome of interest divided by the number of infections:

$$\text{PAR} = \frac{\text{number with post-infection outcome}}{\text{number infected}}. \tag{9.4}$$

Letting p denote the control group and v denote the vaccinated group, then VE_P using a dichotomous outcome can be defined as

$$VE_P = 1 - \frac{PAR(v)}{PAR(p)}. \tag{9.5}$$

9.1.3 Relation of VE_P, VE_S, and VE_{SP}

For dichotomous infection outcomes and dichotomous postinfection outcomes, there is a simple relation between VE_P, VE_S, and VE_{SP}. Let ψ denote the

Table 9.1. Different types of postinfection and postclinical outcomes, VE_P. Ascertainment can be on infection or on clinical disease, which determines the VE_S

VE_S outcome	Postinfection VE_P outcome	Examples
Infection 0,1	dichotomous	clinical case (0,1) clinical case within time interval (0,1) transmission to other (0,1)
	continuous	malaria parasite density HIV viral load
	time-to-event	time to developing symptoms
Clinical case 0,1	dichotomous	severe disease (0,1) death transmission to other (0,1)
	continuous	malaria parasite density chickenpox: number of lesions
	time-to-event	time to clearing infection

relative risk of the post-infection outcome in the infected vaccinated peoples compared with the infected unvaccinated people, and θ be the relative risk of infection in the vaccinated compared with the unvaccinated people. Then VE_P is

$$VE_P = 1 - \frac{\frac{\text{vaccinated cases}}{\text{vaccinated infections}}}{\frac{\text{unvaccinated cases}}{\text{unvaccinated infections}}} = 1 - \psi. \tag{9.6}$$

Letting

$$VE_{S,CI} = 1 - \frac{\frac{\text{infected vaccinated people}}{\text{vaccinated people}}}{\frac{\text{infected unvaccinated people}}{\text{unvaccinated people}}} = 1 - \theta, \tag{9.7}$$

then

$$VE_{SP,CI} = 1 - \frac{\frac{\text{vaccinated cases}}{\text{vaccinated people}}}{\frac{\text{unvaccinated cases}}{\text{unvaccinated people}}}$$

$$= 1 - \frac{\frac{\text{infected vaccinated people}}{\text{vaccinated people}}}{\frac{\text{infected unvaccinated people}}{\text{unvaccinated people}}} \times \frac{\frac{\text{vaccinated cases}}{\text{vaccinated infections}}}{\frac{\text{unvaccinated cases}}{\text{unvaccinated infections}}}$$

$$= 1 - (1 - VE_S)(1 - VE_P) = 1 - \theta\psi. \tag{9.8}$$

These relations hold under the assumption that the infected people in the control arm are comparable to the infected people in the vaccine arm. To estimate (9.8), asymptomatic infections need to be ascertained. A similar relation as that in equation (9.8) holds for $\text{VE}_{SP,CI}$, when the postclinical outcome is severe cases of the disease of interest, and VE_S is based on clinical cases.

9.2 Effect of vaccination on disease severity

9.2.1 Pertussis vaccine study in Niakhar

Préziosi and Halloran (2003b) analyzed a study of pertussis vaccination in the Niakhar study area of Senegal to estimate the effect of the vaccination on reducing the severity of clinical pertussis cases. The study population and surveillance for pertussis in the Niakhar study area is described in Section 10.2.3. Briefly, the Niakhar study area is 150 km southeast of Dakar, Senegal, and includes 30 villages. Extended families reside in compounds. In January 1993, there were 26,306 residents living in 1800 compounds. Surveillance began in March 1983 with annual, after 1987 weekly visits to compounds. Pertussis was endemic, with epidemics every 3–4 years, and 1993 was a pertussis epidemic year. Active surveillance was conducted in children <15 years of age by weekly visits to the compounds by trained field workers. They reported cases in children <15 years old who had potential pertussis (cough of >7 days duration) A physician then visited to confirm clinically and collect laboratory samples.

A case of pertussis was defined by confirmation of pertussis infection by presence of at least one of three laboratory criteria: (1) isolation of *B. pertusiss* from a nasopharyngeal aspirate (culture positive) (2) significant increase or decrease in pertussis toxin or filamentous hemagglutinin antibodies (serology positive), or (3) signs and symptoms of disease in an individual who lived in the same compound as a child who had onset of culture-positive disease within 28 days (epilink).

9.2.2 Global score of disease severity

Estimating VE_P requires defining the disease outcomes of interest carefully. To compare the severe to non-severe cases, definitions of a severe case and a non-severe case, or other levels of severity, such as moderate severity, are needed.

Table 9.2. Scale used to assess the severity of illness among children with symptoms of pertussis.

Variable	No. of points
Severity of cough	
Typical paroxysms with whoops	4
Typical paroxysms without whoops	3
Atypical paroxysms only	1
Apnea	6
Pulmonary sign[a]	3
Mechanical complication[b]	3
Facial swelling	3
Conjunctival injection	3
Post-tussive vomiting	2
Total score (severity)[c]	
Mild disease	≤ 6
Severe disease	>6

[a] Bronchitis or bronchopneumonia.
[b] Subjunctival hemorrhage or umbilical or inguinal hernia.
[c] The overall median total score was 6 in this study.

Préziosi and Halloran (2003b) proposed a scale to assess the global clinical severity of pertussis cases, rather than analyzing each individual symptom. Severity of illness was assessed according to the scale in Table 9.2. Death is not included, because there was only one death due to pertussis in the study period. Each relevant symptom was given a score based on the judged severity of the symptom. The global symptom score for each child was obtained by the simple sum of the individual's clinical signs and symptoms scores. Severe disease was defined by a score greater than a particular threshold value. The main outcome measure was defined using the overall severity median score in the confirmed clinical cases.

Sex, age, and type of case (primary or postinfection) were included in a multivariate analysis using logistic regression and then backtransformed to the relative risk scale (Halloran et al 2003) (Chapter 11). Confidence intervals were obtained using the bootstrap (Efron and Tibshirani 1993).

In 1993, 2,123 individuals with potential cases of pertussis were identified in 518 of 1800 residential compounds, 98% under 15 years of age. Nearly all children under 6 months or 9 years and older were unvaccinated, so these age groups could not be included in comparison. Cultures were done on 99% of all suspected cases, and serologic testing in 69% of unvaccinated and 83% of vaccinated suspected cases. In all, 834 children with 837 cases of laboratory-confirmed pertussis were identified. Details of confirmation criteria and clinical signs and symptoms are in Préziosi and Halloran (2003b).

Table 9.3. Number of cases of severe pertussis, among 834 children who had or had not received pertussis vaccine, and efficacy of the vaccine in reducing severity, according to severity score (Préziosi and Halloran 2003b).

	No. (%) of cases			
Score[a]	All ($n = 837$)	In unvaccinated children ($n = 243$)	In vaccinated children ($n = 594$)	Vaccine efficacy, % VE_P (95% CI)
>0	738 (88)	233 (96)	505 (85)	11 (8–15)
>1	728 (87)	231 (95)	497 (84)	12 (8–16)
>2	677 (81)	227 (93)	450 (76)	19 (14–23)
>3	559 (67)	205 (84)	354 (60)	29 (23–35)
>4	529 (63)	194 (80)	335 (56)	29 (22–36)
>5	443 (53)	178 (73)	265 (45)	39 (32–46)
>6	339 (41)	149 (61)	190 (32)	48 (39–55)
>7	315 (38)	139 (57)	176 (30)	48 (39–56)
>8	268 (32)	119 (49)	149 (25)	49 (38–58)
>9	151 (18)	76 (31)	75 (13)	60 (47–70)
>10	147 (18)	75 (31)	72 (12)	61 (48–71)
>11	130 (16)	67 (28)	63 (11)	62 (48–72)
>12	31 (4)	20 (8)	11 (2)	78 (54–89)
>13	30 (4)	19 (8)	11 (2)	76 (51–89)
>14	24 (3)	17 (7)	7 (1)	83 (60–93)

[a] The scale used to assign the severity score is shown in Table 9.2. The overall median score was 6. A score ≤6 indicates mild disease; a score >6 indicates severe disease.

9.2.3 VE_P for severity of pertussis disease

Based on the median threshold for mild versus severe disease of 6, 61% of unvaccinated children and 32% of vaccinated children had severe disease. Based on this threshold,

$$\begin{aligned} VE_P &= 1 - \frac{\dfrac{\text{severe vaccinated cases}}{\text{all vaccinated cases}}}{\dfrac{\text{severe unvaccinated cases}}{\text{all unvaccinated cases}}}. \\ &= 1 - \frac{190/594}{149/243} = 0.48 \quad (95\%\ \text{CI} \quad 0.39, \quad 0.55) \end{aligned} \tag{9.9}$$

Thus, unvaccinated children were twice as likely as vaccinated children to have severe disease. Table 9.3 presents a sensitivity analysis of the estimate of VE_P to the choice of threshold for defining a severe case. The threshold varies from 1 to >14. The estimated VE_P varies from 11% to 83%, becoming higher as the threshold for defining a severe case gets higher. The lower limit of the 95% CI was greater than 0 for all thresholds. The results indicate that pertussis

vaccination substantially decreases the severity of breakthrough disease in children who receive 3 doses of vaccine, compared with that in unvaccinated children. The majority of vaccinated children who developed pertussis had mild disease.

Because this is an observational study, there is a potential for selection bias, particularly in (1) ascertainment and (2) laboratory confirmation. Both are minimal in this case because (1) active surveillance, and (2) most children with suspected cases had laboratory tests done. To assess potential bias in the selection of the confirmed cases, they examined clinical illnesses among children with potential case of pertussis whose biological tests were negative and among children for whom no laboratory samples were available. A comparison of the vaccinated and unvaccinated children who either had no biological test done and in whom the tests were all negative was also done. The vaccine showed no appreciable effect in these groups.

In a secondary analysis, VE_{SP} was also computed, first using all cases, and then using just severe cases. Child-years at risk were computed for 1993 among susceptible children 6 months up to 8 years old. Standard CIs were computed assuming log-normality of relative risks. In the secondary analysis, VE_{SP} for all cases was 29% (95% CI, 19% – 39%), and VE_{SP} for severe cases was 64% (95% CI, 55% – 71%). It is typical that VE_{SP} is higher for more severe or stringent case definitions.

9.2.4 Rotavirus vaccine in Finland

Vesikari et al. (1990) provide a a second example in their analysis of a randomized, double-blinded, placebo controlled trial of a *Rhesus* rotavirus candidate vaccine. The trial was conducted in children two to five months of age from 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. Combining the severe and moderately severe cases, five of 10 cases in the vaccinated group were severe or moderately severe, and 13 of 16 cases in the placebo group were severe or moderately severe. Using equation (9.3) yields $\widehat{\text{VE}}_P = 0.38$, (95% CI -0.11,0.74).

9.3 Causal Effects in Post-Infection Outcomes

9.3.1 Postinfection selection bias

In Sections 9.1 and 9.2 the assumption was made that the infected vaccinated group and the infected unvaccinated group were comparable. In most studies up to recently the assumption was not seriously questioned. However, conditioning on an event, such as infection, that occurs subsequent to receipt of

vaccine or control could result in selection bias, even if the study were randomized (Halloran and Struchiner 1995). With the development of HIV vaccine candidates, the assumption about the comparability of the infected vaccinated and infected unvaccinated groups gained considerable attention (Hudgens et al., 2003; Gilbert et al., 2003b). The initial HIV vaccine candidates were hoped to protect against infection and also slow progression to AIDS postinfection. The HIV vaccine trials were designed to draw blood from all the participants to ascertain infection at about three month intervals. Since the incubation period to the development of AIDS after HIV infection is usually several years, post-infection measures in the blood such as viral load and CD4 cell count are used as surrogates of potential future development of AIDS. Concern grew that the infected people in the vaccinated group and infected people in the unvaccinated group might not be comparable.

For example, assume that the potential immune response to HIV has a distribution in the population before individuals are randomized to vaccine or control. Randomization would assure that in large samples, the potential distribution of the immune response to HIV would be the same in the vaccine and the control groups. It could be that the vaccine enhances protection only in people who have the stronger immune system, conferring some level of protection against infection if exposed. Then the people in the vaccinated group who become infected would be the ones with the weaker immune system, whereas the infected people in the unvaccinated group would be those with a weaker immune system as well as those with the stronger immune system. In this situation, if we compare a postinfection outcome in the vaccinated group with that in the unvaccinated group, it could appear that the vaccine makes things worse, even if vaccination has absolutely no effect on anything after infection.

For example, if people with a weaker immune system tend to have a higher viral load after being infected than those with a stronger immune system, then the mean viral load in the infected vaccinated group would be higher than the mean viral load in the infected unvaccinated group (Figure 9.1). The resulting VE_P estimate would be negative. This observation could lead to the false conclusion that the vaccine made the postinfection outcome worse, possibly resulting in rejection of a potentially useful vaccine candidate (Hudgens et al., 2003; Gilbert et al., 2003b). However, the vaccine in this case actually does not make anything worse. The problem is that the infected vaccinated group and infected control group are no longer comparable because of selection bias.

A similar problem exists in principal for other diseases and vaccines other than HIV, but it received considerably less attention. When the benefits of vaccination are clearly positive, selection bias might not lead to discarding the vaccine, but to either an over- or an underestimate of the public health benefits. Thus, it is important both scientifically and for public health purposes to be able to differentiate the effects of vaccines on infection from their effects on post-infection outcomes, and to account for potential selection bias.

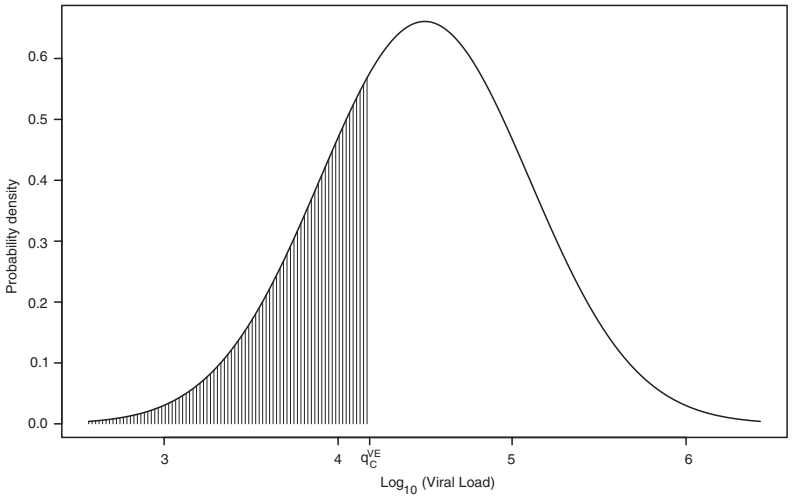

Fig. 9.1. Viral load distribution for infected participants under a selection model. The normal distribution represents the viral loads of the infected controls. The shaded area represents the potential viral loads of the vaccine efficacy $\times$ 100 per cent that are protected by the vaccine. The unshaded area (after appropriate scaling) represents the viral load distribution of the infected vaccinees (from Hudgens, Hoering, Self 2003).

.

9.3.2 Defining causal estimands for post-infection outcomes

How do we account for possible selection in estimating VE_P if we do not know if it is present? Different methods have been used to adjust analyses for post-treatment variables such as infection (Robins and Greenland 1992, 1994; Rosenbaum 1984). The method presented here is based on the potential outcomes approach to causal inference introduced in Chapter 1.4. In Table 1.1, four types of people are defined based on their joint potential outcome under vaccine and control, namely immune, harmed, protected, and doomed. If infection is the potential outcome of interest, then the four types of people are defined by their joint potential infection outcome under vaccine and control. 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 (Rosenbaum 1984; Frangakis and Rubin 2002).

Frangakis and Rubin (2002) propose a method to adjust for post-treatment variables, called *principal stratification*, that stratifies on the joint potential post-treatment variables under each of the treatments being considered. The causal effects of one treatment compared to the other on a main outcome of interest are defined within each of these principal strata and are called principal

effects. If infection is considered as a post-treatment variable, then the post-infection outcome is defined under vaccine and placebo only in the doomed stratum, in which people would be infected under both vaccine and placebo. The post-infection outcome is not defined for anyone in the immune stratum. It is defined only under placebo in the protected stratum, and only under vaccine for the harmed stratum. The importance of estimating quantities defined only in a subpopulation in which the outcomes are defined was presented in the context of outcomes censored by death (Kalbfleisch and Prentice 1980). Robins (1986, remark 12.2,) considered inference about causal effects in the stratum that would survive under either treatment.

Several papers have been published using this approach to assess vaccine effects on post-infection outcomes. 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 post-infection outcome viral load. Hudgens et al. (2003) developed bounds. Gilbert et al.(2003b) adapted methods for sensitivity similar to that of Scharfstein, Robins, and Rotnitzky (1999) and Robins, Rotnitzky, and Scharfstein (2000). Shepherd et al (2006) considered sensitivity analyses comparing outcomes only existing in a subset selected post-randomization, conditional on covariates, with application to HIV. Jemiai et al (2007) develop extensions of Gilbert et al (2003) that allow the estimation of treatment effects conditional on covariates. Shepherd et al (2007) developed the methods for a time-to-event postinfection outcome, also with application to HIV vaccine. The time-to-event postinfection outcome was the time from infection diagnosis to initiation of antiretroviral therapy. Hudgens and Halloran (2006) developed methods for the causal vaccine effects on binary postinfection outcomes with applications to pertussis and rotavirus vaccines. Table 9.4 summarizes literature on bounds and sensitivity analyses of causal vaccine effects for different types of postinfection outcomes.

Because the development for continuous and time-to-event postinfection outcomes involves complex integral equations, we focus on the development of the dichotomous post-infection outcome from Hudgens and Halloran (2006). The approach for the continuous and time-to-event outcomes is similar. The common steps of the argument regardless of the type of post-infection outcome are as follows:

1. Assume SUTVA and an assignment mechanism independent of the potential outcomes, for example, randomization.
2. Define the causal VE_P in the doomed (always-infected) principal stratum, which is not identifiable from the observed data without further assumptions.
3. Assume that the harmed principal stratum is empty, called the monotonicity assumption.
4. The monotonicity assumption implies that all infected vaccine recipients are in the doomed stratum, so the numerator of the causal VE_P is identifiable. However, the infected placebo recipients could be in either the

Table 9.4. Bounds and sensitivity analyses of causal vaccine effects, VE_P, for different types of postinfection outcomes assuming SUTVA, randomization, and monotonocity.

Infection outcome VE_S	Postinfection outcome VE_P	Analysis	Reference
0, 1	continuous	bounds	Hudgens, Hoering, Self 2003
		sensitivity analysis	Gilbert, Bosch, Hudgens 2003
		covariates	Shepherd, Gilbert, Jemiai, Rotnitzky 2006
			Jemiai, Rotnitzky, Shepherd, Gilbert 2007
	dichotomous	bounds and sensitivity analysis	Hudgens and Halloran 2006
	time-to-event	bounds and sensitivity analysis	Shepherd, Gilbert, Lumley 2007

protected or the doomed stratum, so the denominator of the causal VE_P is not identifiable.

5. Bounds can be set on the estimates of the causal VE_P by extreme assumptions about the distribution of post-infection outcome in the infected placebo recipients in the protected compared with the distribution in the the doomed stratum.
6. Sensitivity analyses can be done by varying a selection bias parameter over reasonable ranges of selection bias, with the assumption of no selection bias being a special case.

To formalize these concepts, we use an extension of the causal model introduced in Chapter 1.4 and Table 1.1. Let $Z_i = v$ if the ith individual is assigned vaccine, and $Z_i = p$ if assigned control. Denote the potential infection outcome of the ith individual if assigned Z_i as $S_i(Z_i)$, where $S_i(Z_i) = 0$ if uninfected and $S_i(Z_i) = 1$ if infected. The focus is on evaluating the causal effect of vaccine on the outcome Y that occurs after an individual becomes infected. Y could be a continuous random variable, a time-to-event variable, or a dichotomous outcome. Here we develop the notation for a dichotomous outcome. If $S_i(Z_i) = 1$, then $Y_i(Z_i) = 1$ if the ith individual has the worse, or more severe postinfection outcome, and $Y_i(Z_i) = 0$ otherwise. If an individual's potential infection outcome for an assignment is uninfected, that is, $S_i(Z_i) = 0$, then $Y_i(Z_i)$ is undefined and denoted by $*$. 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 post-infection outcome.

Table 9.5. Basic principal stratification P_0 based on the potential infection outcomes $(S(v), S(p))$ with potential post-infection strata based on $(Y(v), Y(p))$ (Hudgens and Halloran 2006).

Potential infection strata		Potential post-infection strata	
Basic principal stratum, S^{P_0}	Potential infection outcomes $(S(v), S(p))$	Potential post-infection outcomes $(Y(v),Y(p))$	Post-infection interpretation
immune	(0,0)	$(*,*)$	always undefined
harmed	(1,0)	$(0,*)$	not severe vaccine, undefined placebo
		$(1,*)$	severe vaccine, undefined placebo
protected	(0,1)	$(*,0)$	undefined vaccine, not severe placebo
		$(*,1)$	undefined vaccine, severe placebo
doomed	(1,1)	(0,0)	never severe
		(1,0)	harmed by vaccine
		(0,1)	helped by vaccine
		(1,1)	always severe

In the following, we assume that the potential outcomes for each individual are independent of the treatment assignment of other individuals, that is, there is no interference between individuals (SUTVA). We further assume that the assignment to vaccine or control is independent of the potential infection outcomes and the potential postinfection outcomes. Randomization is one assignment mechanism where the treatment assignment is independent of the potential outcomes.

A *basic principal stratification* P_0 is defined according to the joint potential infection outcomes $S^{P_0} = (S(v), S(p))$ (Frangakis and Rubin 2002). Table 9.5 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 post-infection 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 (infected under both vaccine and placebo). Since membership in a basic principal stratum is not affected by whether an individual is actually assigned vaccine or placebo, the strata can be used in the same way as pre-treatment covariates, with causal post-infection vaccine effects defined within a basic principal stratum S^{P_0}.

In general, causal effects are defined in terms of potential outcomes. From Table 9.5, we see the doomed basic principal stratum, $S^{P_0} = (1, 1)$, is the only stratum in which both potential post-infection endpoints, and thus their joint

distribution, are defined. For this reason, defining individual post-infection 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 post-infection outcome only for an individual who would become infected whether vaccinated or not. Thus, 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 post-infection outcome (Y) for those participants who would be infected under both treatment assignments.

Regardless of the type of of post-infection outcome, the population casual vaccine efficacy to prevent infection $S = 1$ can be defined as

$$\mathrm{VE}_S = 1 - \frac{\Pr(S(1) = 1)}{\Pr(S(0) = 1)}, \tag{9.10}$$

the relative average causal effect (RACE) of vaccination on infection (Hudgens and Halloran 2006). Under randomization, it follows that

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

The causal vaccine efficacy for a binary post-infection for those participants who would be infected under both treatment assignments, VE_P, is defined in equation (9.11). The form is different for continuous and time-to-event outcomes. The problem with this approach is that it is not possible to tell to which stratum any individual belongs, at least without further assumptions. For example, a person who is vaccinated and becomes infected could belong to either the doomed or the harmed stratum. A person who receives control and is infected could belong to either the doomed or the protected stratum. Thus it is not possible to estimate the causal VE_P from the observed data without further assumptions.

One assumption that is plausible for most vaccines is helpful in this situation. If we assume that the vaccine does not harm people with respect to infection, then we can claim that the harmed stratum is empty. This assumption is called the monotonicity assumption. Under the monotonicity assumption, a vaccinated person who becomes infected must be in the doomed stratum. The monotonicity assumption does not help with the people who receive control and become infected. Infected people in the control arm can still be in either the protected or the doomed stratum.

Although it is not possible to identify who of the infected control group is in the protected or doomed stratum, it is possible to set upper and lower bounds on the vaccine effect on the postinfection outcome, $\widehat{\mathrm{VE}}_P^{upper}$ and $\widehat{\mathrm{VE}}_P^{lower}$. Estimating the causal vaccine effect under an extreme degree of selection bias is useful in bounding the estimate of the post-infection effect above and beyond any possible selective effects. However, the true degree of selection bias is likely less than the extreme models, such that using $\widehat{\mathrm{VE}}_P^{upper}$ or $\widehat{\mathrm{VE}}_P^{lower}$

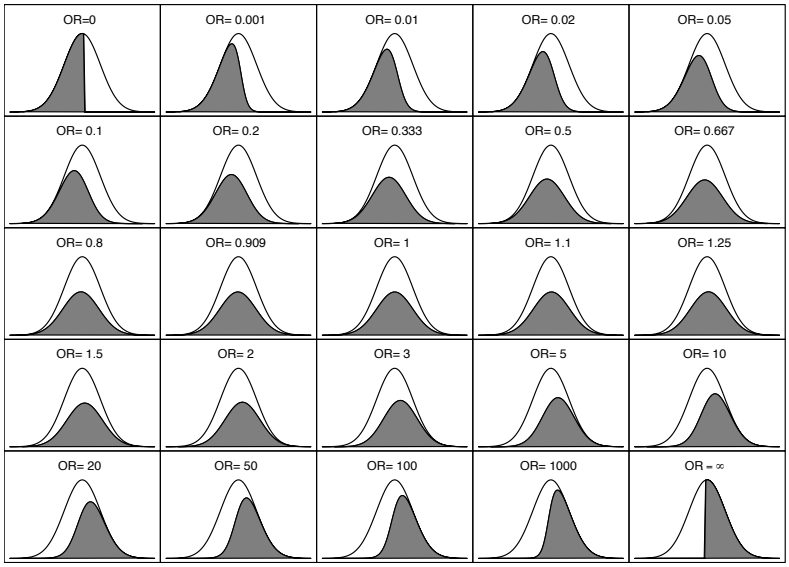

Fig. 9.2. Distribution of the potential post-infection outcome Y in the infected control group in the protected stratum and the infected control group in the doomed stratum for different values of the selection bias odds ratio $\exp(\beta)$. The shaded area represents the distribution of the potential Y outcome in the infected control group in the doomed stratum. The area under the clear distribution is that in the protected stratum. (courtesy of B. Shepherd)

.

may be too conservative. Therefore, it is useful to do sensitivity analyses by varying the amount of selection bias, in which the case of no selection and extreme bounds are included as special cases.

Gilbert et al.(2003b), Shepherd et al (2006), and Shepherd et al (2007) adapted methods for sensitivity similar to that of Scharfstein, et al (1999) and Robins, et al (2000) for continuous outcomes. In this approach, the sensitivity analysis is performed by varying a selection bias parameter β over a range. In particular the odds ratio, $\text{OR} = \exp(\beta)$, is varied from 0 to $+\infty$, with no selection bias being at $\text{OR} = 1$. The odds ratio is interpreted as given infection in the placebo arm, for a one unit increase in the Y outcome, the odds of being infected if randomized to the vaccine arm multiplicatively increases by $\text{OR} = \exp(\beta)$.

Figure 9.2 illustrates different degrees of selection bias associated with varying the odds ratio, showing the distributions of the potential Y outcome in the infected control group in the protected stratum and the infected control group in the doomed stratum. The shaded area represents the distribution of

the potential Y outcome in the infected control group in the doomed stratum. The area under the clear distribution is that in the protected stratum. When the odds ratio equals 1, there is no selection bias, and the distributions in the two strata are the same. As the odds ratio tends to 0, the distribution of the Y outcome in the doomed stratum tends to be lower than the Y outcome in the protected stratum. As the odds ratio tends to ∞, the distribution of the Y outcome in the doomed stratum tends to be higher than the Y outcome distribution in the protected stratum. The data do not provide information about the degree of selection bias. Then outside knowledge or expert opinion can be used to choose a plausible range for the selection bias (Shepherd et al 2006).

9.4 Causal Effects for Binary Post-infection Outcomes

9.4.1 Defining vaccine effects

Three estimands regarding the effect of vaccination on the binary post-infection outcome Y can be formally defined (Hudgens and Halloran 2006)). The approach to assessing vaccine effects on post-infection endpoints based on the observed data in Section 9.1 is the *net* vaccine effect estimand which conditions on infection, i.e.,

$$\mathrm{VE}_P^{net} = 1 - \frac{E\left\{Y^{obs}|S^{obs}=1, Z=v\right\}}{E\left\{Y^{obs}|S^{obs}=1, Z=p\right\}} = 1 - \frac{E\left\{Y(v)|S(v)=1\right\}}{E\left\{Y(p)|S(p)=1\right\}},$$

with the second equality following from the independence, e.g. randomization assumption. As discussed in Section 9.3.1, in general, VE_P^{net} does not have a causal interpretation since the set of individuals with $S(v)=1$ is not necessarily identical to the set of individuals with $S(p)=1$.

An estimand that defines the effect of vaccination on disease rather than infection, or severe disease rather than disease, as in Chapter 6 might be considered intent-to-treat (ITT) because it does not condition on the post-treatment variable S^{obs}. It incorporates all individuals according to their treatment assignment. $\mathrm{VE}_{SP,CI}$ in equation (9.8) is an example of an ITT estimand. Formally,

$$\mathrm{VE}_{SP,CI} = \mathrm{VE}_P^{ITT} = 1 - \frac{E\left\{Y(v) \times S(v)\right\}}{E\left\{Y(p) \times S(p)\right\}},$$

where the convention sets $Y(z) \times S(z) = 0$ if $S(z) = 0$, $z = v, p$. This is a general form for what Préziosi and Halloran (2003b) called "VE_{SP} for severity." The VE_P^{ITT} estimand has a causal interpretation, but it combines vaccine effects on susceptibility and the post-infection outcome. Formally, equation (9.8) can be written as

$$\mathrm{VE}_P^{ITT} = 1 - (1 - \mathrm{VE}_S)(1 - \mathrm{VE}_P^{net}).$$

Hudgens and Halloran (2006) propose a third estimand for the effect of vaccination on a binary post-infection outcome that has a causal interpretation and is separate from vaccine effects on susceptibility. Using the basic principal stratification shown in Table 9.5, they define the causal VE_P. In particular, the individual causal vaccine effect on the post-infection outcomes is defined as

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

for individuals within the doomed principal stratum only. Following the development of VE_S above, define the population post-infection causal vaccine effect VE_P within the doomed principal stratum as

$$\mathrm{VE}_P = 1 - \frac{E\{Y(v)|S^{P_0} = (1,1)\}}{E\{Y(p)|S^{P_0} = (1,1)\}}. \tag{9.11}$$

All of the papers listed in Table 9.4 define an analogous causal estimand for post-infection outcomes based on the individuals within the doomed principal stratum only. Like VE_S, (9.11) could equivalently be given in terms of probabilities since the post-infection random variables $Y(v)$ and $Y(p)$ are assumed to be binary such that VE_P can be interpreted as the causal estimand measuring the relative reduction in the probability of the worse post-infection outcome given vaccine compared to placebo in those individuals who would be infected under either treatment assignment.

9.4.2 Parameterization

Let the parameters $\boldsymbol{\theta}$ govern the probabilities associated with the basic principal strata. By the monotonicity assumption, the harmed stratum $S^{P_0} = (1,0)$ is empty, so let $\boldsymbol{\theta} = (\theta_{00}, \theta_{01}, \theta_{11})$ where

$$\Pr\{S^{P_0} = (i,j); \boldsymbol{\theta}\} = \theta_{ij} \text{ for } i,j = 0,1; i \leq j. \tag{9.12}$$

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

$$\Pr\{(Y(v), Y(p)) = (k,m)|S^{P_0} = (1,1); \boldsymbol{\phi}\} = \phi_{km} \text{ for } k,m = 0,1. \tag{9.13}$$

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

$$\Pr\{Y(p) = i|S^{P_0} = (0,1); \boldsymbol{\gamma}\} = \gamma_l \text{ for } l = 0,1. \tag{9.14}$$

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

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

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

Based on the definition of the causal estimand VE_P given in (9.11), we are not interested in the joint probabilities ϕ_{km} $(k, m = 0, 1)$, but rather just two of the marginal probabilities. In particular,

$$\mathrm{VE}_P = 1 - \frac{\phi_{1\cdot}}{\phi_{\cdot 1}}, \tag{9.15}$$

where

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

and

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

Under this parameterization,

$$\mathrm{VE}_P^{net} = 1 - \frac{\phi_{1\cdot}}{\gamma_1 \mathrm{VE}_S + \phi_{\cdot 1}(1 - \mathrm{VE}_S)},$$

and

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

9.4.3 Estimation

Suppose we observe n independent and identically distributed realizations of (Z, S^{obs}, Y^{obs}), where Y^{obs} is undefined or does not exist if $S^{obs} = 0$. There are six possible observed combinations of (Z, S^{obs}, Y^{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 post-infection outcome Y^{obs}; and $z = v, p$. That is,

$$\begin{aligned}
n_{0*}(p) &= \textstyle\sum_i I(Z_i = p, S_i^{obs} = 0, Y_i^{obs} \text{ does not exist}) \\
n_{10}(p) &= \textstyle\sum_i I(Z_i = p, S_i^{obs} = 1, Y_i^{obs} = 0) \\
n_{11}(p) &= \textstyle\sum_i I(Z_i = p, S_i^{obs} = 1, Y_i^{obs} = 1)
\end{aligned}$$

$$\begin{aligned}
n_{0*}(v) &= \textstyle\sum_i I(Z_i = v, S_i^{obs} = 0, Y_i^{obs} \text{ does not exist}) \\
n_{10}(v) &= \textstyle\sum_i I(Z_i = v, S_i^{obs} = 1, Y_i^{obs} = 0) \\
n_{11}(v) &= \textstyle\sum_i I(Z_i = v, S_i^{obs} = 1, Y_i^{obs} = 1)
\end{aligned}$$

where the summations are over $i = 1, \ldots, n$. The double subscripts for the n's do not have the same meaning as for the ϕ's and θ's. Assume that each of the six combinations is observed at least once. 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 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 placebo and vaccine. Let

$$AR_z = \frac{n_{1\cdot}(z)}{n(z)} \text{ for } z = v, p,$$

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

$$PAR_z = \frac{n_{11}(z)}{n_{1\cdot}(z)} \text{ for } z = v, p,$$

i.e., PAR_z is the observed postinfection attack rate in the group infected given treatment z.

Maximum likelihood estimators (MLEs) of the identifiable vaccine efficacy estimands can be found by maximizing the likelihood

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

subject to constraints on $\boldsymbol{\theta}, \boldsymbol{\gamma}, \boldsymbol{\phi}$ that ensure (9.12-9.14) are probability functions. Hudgens and Halloran (2006) show that the MLE of VE_S is given by

$$\widehat{\mathrm{VE}}_S = \begin{cases} 1 - \frac{AR_v}{AR_p} \text{ if } AR_v \leq AR_p, \\ \\ 0 \qquad\quad \text{otherwise.} \end{cases} \tag{9.16}$$

This is the usual estimator of VE_S based on the attack rates, or cumulative incidence. Further, the MLE of VE_P^{net} is

$$\widehat{\mathrm{VE}}_P^{net} = 1 - \frac{PAR_v}{PAR_p}, \tag{9.17}$$

the same as in equations (9.2) and (9.6). The MLE of VE_I^{ITT} is

$$\widehat{\mathrm{VE}}_P^{ITT} = 1 - (1 - \widehat{\mathrm{VE}}_S)\frac{PAR_v}{PAR_p}, \tag{9.18}$$

or equivalently

$$\widehat{\mathrm{VE}}_P^{ITT} = \begin{cases} \widehat{\mathrm{VE}}_P^{net} \qquad\quad \text{if } \widehat{\mathrm{VE}}_S = 0, \\ \\ 1 - \frac{n_{11}(v)/n(v)}{n_{11}(p)/n(p)} \text{ if } \widehat{\mathrm{VE}}_S > 0, \end{cases} \tag{9.19}$$

the same as in equation (9.8). In summary, the three MLEs $\widehat{\mathrm{VE}}_S$, $\widehat{\mathrm{VE}}_P^{net}$, and $\widehat{\mathrm{VE}}_P^{ITT}$ derived formally by the methods of causal inference correspond to the usual estimators associated with these measures.

The causal estimand VE_P is not identifiable because $\phi_{\cdot 1}$, the denominator of the right side of (9.15), is not identifiable. On the other hand, $\phi_{1\cdot}$, the numerator of the right side of (9.15), can be identified by the observable random variables. The corresponding MLE is given by

$$\widehat{\phi}_{1\cdot} = PAR_v, \tag{9.20}$$

i.e., the observed postinfection attack rate in the vaccine arm.

Finally, while $\phi_{\cdot 1}$ is not identifiable, we can identify

$$\Pr[Y(p) = 1 | S(p) = 1; \boldsymbol{\theta}, \boldsymbol{\gamma}, \boldsymbol{\phi}] = \gamma_1 \mathrm{VE}_S + \phi_{\cdot 1}(1 - \mathrm{VE}_S). \tag{9.21}$$

The MLE of (9.21) is PAR_p such that any feasible pair $(\widehat{\gamma}_1, \widehat{\phi}_{\cdot 1})$ satisfying

$$PAR_p = \widehat{\gamma}_1 \widehat{\mathrm{VE}}_S + \widehat{\phi}_{\cdot 1}(1 - \widehat{\mathrm{VE}}_S), \tag{9.22}$$

is an MLE of $(\gamma_1, \phi_{\cdot 1})$.

9.4.4 Applications

Rotavirus candidate vaccine

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

$n_{0*}(p) = 84$	$n_{0*}(v) = 90$
$n_{10}(p) = 3$	$n_{10}(v) = 5$
$n_{11}(p) = 13$	$n_{11}(v) = 5$

From (9.16), $\widehat{\mathrm{VE}}_S = 1 - (10/100)/(16/100) = 0.375$. It then follows from (9.19) that $\widehat{\mathrm{VE}}_P^{ITT} = 1 - (5/100)/(13/100) = 0.62$. The postinfection attack rates are $\mathrm{PAR}_v = \widehat{\phi}^{1\cdot} = 5/10 = 0.50$ and $\mathrm{PAR}_p = 13/16 = 0.81$ such that $\widehat{\mathrm{VE}}_P^{net} = 1 - (5/10)/(13/16) = 0.385$.

To consider estimation of the causal VE_P, we examine the relation of the observed data to the basic principal strata and the strata of joint potential post-infection outcomes within each basic principal stratum. By the assumptions of SUTVA, independence, and monotonicity, 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 the immune stratum $S^{P_0} = (0, 0)$ or the protected stratum $S^{P_0} = (0, 1)$.
- The $n_{10}(p) + n_{11}(p) = 16$ could belong to the protected $S^{P_0} = (0, 1)$ or the doomed $S^{P_0} = (1, 1)$.

Ignoring statistical variability, by the independence assumption, since there are 10 vaccine recipients in the doomed stratum, there are 10 placebo recipients in the doomed stratum. Since there are 84 placebo recipients in the immune stratum, there are 84 vaccine recipients in the immune stratum. So there must be 6 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 protected stratum $S^{P_0} = (0, 1)$ or which 10 of the 16 are in the doomed stratum $S^{P_0} = (1, 1)$. Why do we care? Because to estimate the causal VE_P, we need to know the post-infection outcomes of those in the doomed stratum. This illustrates the need for further assumptions to identify VE_P.

Pertussis vaccine

The pertussis vaccine analysis presented in Section 9.2.1 included exactly one year of follow-up, the calendar year 1993, so 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 $n(p)$. During that one calendar year, there were 3845 and 1020 person-years at risk in the vaccinated and unvaccinated children (Préziosi, pers. comm.). Using slightly different inclusion criteria for cases than in Section 9.2.1, of 548 cases in the vaccinated group, 176 were severe, and of 206 cases in the unvaccinated group, 129 were severe. Based on equation (9.3), $\widehat{\text{VE}}_P = 0.49$, (95% CI 0.40,0.56). Although vaccine status was not randomized, there was no evidence of systematic differences between the vaccinated and unvaccinated groups, so that the independence assumption might be reasonable. The observed data are

$$\begin{array}{ll} n_{0*}(p) = 814 & n_{0*}(v) = 3297 \\ n_{10}(p) = 77 & n_{10}(v) = 372 \\ n_{11}(p) = 129 & n_{11}(v) = 176 \end{array}$$

From (9.16), $\widehat{\text{VE}}_S = 1 - (548/3845)/(206/1020) = 0.29$. The postinfection attack rates are $\text{PAR}_v = \widehat{\phi^{1\cdot}} = 176/548 = 0.32$ and $\text{PAR}_p = 129/206 = 0.63$ such that $\widehat{\text{VE}}_P^{net} = 1-(176/548)/(129/206) = 0.49$, which is the same as $\widehat{\text{VE}}_P$ yielded by using (9.3). Finally $\widehat{\text{VE}}_P^{ITT} = 1 - (176/3845)/(129/1020) = 0.64$, which in Préziosi and Halloran (2003b) was $\widehat{\text{VE}}_{SP}$ for severity.

9.4.5 Selection bias models

The inability to identify the causal VE_P is due to $\phi_{\cdot 1}$ and γ_1 not being separated in the term

$$\theta_{01}\gamma_1 + \theta_{11}\phi_{\cdot 1} = \Pr[Y^{obs} = 1, S^{obs} = 1 | Z^{obs} = p]. \tag{9.23}$$

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

$$\{(\gamma_1, \phi_{\cdot 1}) : 0 \leq \gamma_1 \leq 1, 0 \leq \phi_{\cdot 1} \leq 1, \text{ and (9.23) holds}\}, \tag{9.24}$$

will yield the same distribution of (Z, S^{obs}, Y^{obs}). The selection models presented in this section place additional constraints on the parameter space such that only one pair of parameters satisfy (9.24).

No selection bias

The assumption of no selection implies that the probability of the post-infection outcome conditional on infection under placebo is independent of infection status under vaccine:

$$\Pr\left\{Y(p) = y | S^{P_0} = (1,1); \boldsymbol{\phi}\right\} = \Pr\left\{Y(p) = y | S^{P_0} = (0,1); \boldsymbol{\gamma}\right\} \text{ for } y = 0, 1, \tag{9.25}$$

which implies $\phi_{\cdot 1} = \gamma_1$. Under assumption (9.25), from (9.22), the resulting MLE is

$$\widehat{\phi}_{\cdot 1} = \text{PAR}_p.$$

From (9.20), (9.26), and the definition of VE_P given by (9.15), it follows that the MLE of the causal VE_P equals $\widehat{\text{VE}}_P^{net}$ as given in (9.17). In other words, under the additional assumption of no selection bias as specified by (9.25), the MLE of the causal vaccine effect is the usual postinfection attack rate ratio estimator one obtains when conditioning on infection as in equation (9.2).

Upper and lower bounds

The upper bound selection model yields the parameter pair $(\gamma_1, \phi_{\cdot 1})$ consistent with the observed data that has the largest $\phi_{\cdot 1}$, thus largest VE_P. Since (9.24) 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, i.e., either when

$$\Pr[Y(p) = 1 | S^{P_0} = (1,1)] = \phi_{\cdot 1} = 1, \tag{9.26}$$

or

$$\Pr[Y(p) = 1 | S^{P_0} = (0,1)] = \gamma_1 = 0. \tag{9.27}$$

In words, the upper bound selection bias model assumes either (i) all placebo recipients in the doomed principal stratum have the worse post-infection outcome or (ii) all placebo recipients in the protected principal stratum have the better post-infection outcome. From (9.22) it follows that the unique MLE of VE_P assuming either (9.26) or (9.27) is given by:

$$\widehat{\text{VE}}_P^{upper} = \begin{cases} 1 - PAR_v \text{ if } \widehat{\text{VE}}_S > 1 - PAR_p, \\ \widehat{\text{VE}}_P^{ITT} \quad \text{if } 0 < \widehat{\text{VE}}_S \leq 1 - PAR_p, \\ \widehat{\text{VE}}_P^{net} \quad \text{if } \widehat{\text{VE}}_S = 0. \end{cases} \tag{9.28}$$

All MLEs obtained under the three key assumptions must be less than or equal to $\widehat{\text{VE}}_I^{upper}$.

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

$$\Pr[Y(p) = 1 | S^{P_0} = (1,1)] = \phi_{\cdot 1} = 0, \tag{9.29}$$

or with probability one in the protected principal stratum,

$$\Pr[Y(p) = 1 | S^{P_0} = (0,1)] = \gamma_1 = 1. \tag{9.30}$$

The resulting unique MLE of VE_I is

$$\widehat{\text{VE}}_P^{lower} = \begin{cases} -\infty & \text{if } \widehat{\text{VE}}_S > PAR_p, \\ 1 - PAR_v / \left\{ \frac{PAR_p - \widehat{\text{VE}}_S}{1 - \widehat{\text{VE}}_S} \right\} & \text{if } 0 < \widehat{\text{VE}}_S \leq PAR_p, \\ \widehat{\text{VE}}_P^{net} & \text{if } \widehat{\text{VE}}_S = 0. \end{cases} \tag{9.31}$$

Hudgens and Halloran (2006) derived the circumstances when the upper bound will be negative (suggesting harm) and when the lower bound will be positive (suggesting benefit). For example, $\widehat{\text{VE}}_P^{upper}$ will be negative if and only if $\widehat{\text{VE}}_S \leq 1 - PAR_p$ and $\widehat{\text{VE}}_P^{ITT} < 0$. Similarly, $\widehat{\text{VE}}_P^{lower}$ will be positive if and only if $\widehat{\text{VE}}_S \leq PAR_p$ and $PAR_v < (PAR_p - \widehat{\text{VE}}_S)/(1 - \widehat{\text{VE}}_S)$. On the other hand, for $\widehat{\text{VE}}_S > \max\{PAR_p, 1 - PAR_p\}$, $\widehat{\text{VE}}_P^{upper}$ will be always positive and $\widehat{\text{VE}}_P^{lower}$ will be always negative. In other words, for large enough $\widehat{\text{VE}}_S$ the sign of VE_P cannot be determined unless further assumptions are made beyond SUTVA, independence, and monotonicity.

Sensitivity analysis for selection bias

Hudgens and Halloran (2006) present three approaches to sensitivity analysis that allow selection models to range from no selection to the extreme maximum possible levels.

Log odds ratio of infection

The first approach is similar to that of Scharfstein, et al (1999) and Robins, et al (2000). The sensitivity model is defined in terms of the log odds ratio of having the severe post-infection endpoint under placebo in the doomed versus protected principal strata:

$$\exp(\beta) = \frac{\Pr[Y(p)=1|S^{P_0}=(1,1)]/\Pr[Y(p)=0|S^{P_0}=(1,1)]}{\Pr[Y(p)=1|S^{P_0}=(0,1)]/\Pr[Y(p)=0|S^{P_0}=(0,1)]}. \quad (9.32)$$

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

$$\phi_{\cdot 1} = \frac{\gamma_1 \exp(\beta)}{\gamma_0 + \gamma_1 \exp(\beta)}. \quad (9.33)$$

For fixed β, one can solve equations (9.22) and (9.33) for $\phi_{\cdot 1}$ and, in turn, VE_P. The sensitivity analysis is done by repeating this process over a range of different βs. The bounds are given above.

Conditioning on γ_1 as the sensitivity analysis parameter

The second approach to a sensitivity analysis conditions on the nuisance parameter γ_1 which governs the post-infection endpoint distribution in the protected stratum. If γ_1 is assumed known, from (9.22), the resulting MLE of VE_I is

$$\widehat{\mathrm{VE}}_P = 1 - PAR_v / \left\{ \frac{PAR_p - \gamma_1 \widehat{\mathrm{VE}}_S}{1 - \widehat{\mathrm{VE}}_S} \right\}, \quad (9.34)$$

where γ_1 varies between

$$\max\left\{0, \frac{PAR_p - (1-\widehat{\mathrm{VE}}_S)}{\widehat{\mathrm{VE}}_S}\right\} \le \gamma_1 \le \min\left\{1, \frac{PAR_p}{\widehat{\mathrm{VE}}_S}\right\}, \quad (9.35)$$

with the left side of (9.35) giving rise to $\widehat{\mathrm{VE}}_P^{upper}$ and the right side of (9.35) giving rise to $\widehat{\mathrm{VE}}_P^{lower}$.

Complete data model

The 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, 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, $\phi_{\cdot 1}$ becomes identifiable. Maximizing the complete data log likelihood for $(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\gamma})$ yields the MLE

$$\widehat{\phi_{\cdot 1}} = \frac{n_{11}^d(p)}{n_{10}^d(p) + n_{11}^d(p)}, \quad (9.36)$$

the postinfection attack rate under placebo in the doomed stratum. The sensitivity analysis involves estimating VE_P using (9.36) for all possible complete data configurations consistent with the assumptions and the constraints implied by the observed data. Molenberghs et al. (2001) call this set of point estimates the *region of ignorance*. They call the collection of confidence intervals (CIs) or other measures of precision together with the region of ignorance the *region of uncertainty*.

Statistical variability

Once a particular selection model has been assumed, it becomes a finite-dimensional parametric inference problem with a unique MLE. Conditional on a selection model, standard methods can be used to obtain CI estimates for VE_P. For example, CIs can be computed assuming the usual χ^2 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_P can be determined. The resulting CIs can then be used to determine a region of uncertainty for any of the sensitivity analyses described above. A region of uncertainty which excludes zero implies a statistically significant post-infection causal effect of the vaccine.

Applications, continued

Rotavirus candidate vaccine

For these data, $\widehat{\mathrm{VE}}_S > 1-PAR_p$, so from (9.28), $\widehat{\mathrm{VE}}_P^{upper} = 1-PAR_v = 0.50$. On the other hand, $0 < \widehat{\mathrm{VE}}_S \leq PAR_p$, so from (9.31),

$$\widehat{\mathrm{VE}}_P^{lower} = 1 - \frac{5/10}{\frac{\{13/16-(1-(10/100)(16/100))\}}{(10/100)/(16/100)}} = 0.29.$$

Figure 9.3a shows the sensitivity analysis of $\widehat{\mathrm{VE}}_P$ as a function of the odds ratio e^{β}. For this figures, profile likelihood based CIs are presented; Wald-type CIs give qualitatively similar results. The vertical dotted line in Figure 9.3 corresponds to the assumption of no selection bias. The lack of statistical significance in this example may be due simply to small sample size. If the study had had 1000 participants in each arm with the same observed marginal distributions, then the 95% CI for VE_P under the lower bound model would have been [0.09, 0.46], indicating a significant causal vaccine effect on rotavirus disease severity in individuals who would have been infected under assignment to either vaccine or control.

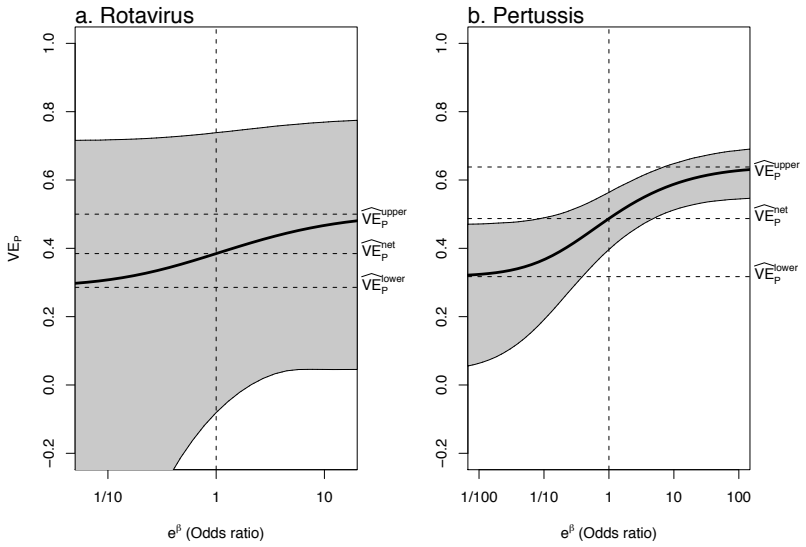

Fig. 9.3. Sensitivity analysis using the odds ratio of having the severe post-infection endpoint under placebo in the doomed versus protected principal strata. (a) rotavirus; (b) pertussis. The vertical dotted line corresponds to the assumption of no selection bias (from Hudgens and Halloran 2006).

Pertussis vaccine

For the pertussis vaccine data, $0 \leq \widehat{\text{VE}}_S \leq 1 - PAR_p$, so from (9.28), $\widehat{\text{VE}}_P^{upper} = \widehat{\text{VE}}_P^{ITT} = 0.64$. On the other hand, $0 \leq \widehat{\text{VE}}_S \leq PAR_p$, so from (9.31), $\widehat{\text{VE}}_P^{lower} = 0.32$. Figure 9.3b shows the sensitivity analysis of $\widehat{\text{VE}}_P$ as a function of the odds ratio e^{β}. The lower limit of the 95% CIs are well above zero over the range of the selection model, suggesting pertussis vaccination causes significant protection against severe disease in children who would develop pertussis regardless of vaccination status.

Problems

9.1. Complete-data likelihood for sensitivity analysis
(a) Write out the complete data log likehood discussed in Section 9.4.5.
(b) How many complete-data configurations are there for the rotavirus example?
(c) What are the different values of $\widehat{\text{VE}}_P$ corresponding to those configurations?

(d) 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. Sketch out the sensitivity analysis for the pertussis vaccine example.

9.2. Varying γ_1 for sensitivity analysis
(a) What is the range for γ_1 in the rotavirus vaccine example? in the pertussis vaccine example?
(b) Produce a graph of VE_P over the range of γ_1 for the rotavirus vaccine example. For the pertussis vaccine example. Mark VE^{upper}, VE^{lower}, and VE^{net}, and the γ_1 corresponding to the assumption of no selection bias.

9.3. Antiviral efficacy against progression, AVE_P
In two studies of the influenza antiviral agent oseltamivir (Hayden et al 2004; Welliver et al 2001), because asymptomatic infections, as well as symptomatic disease (Sec. 10.3.5), had been ascertained in all household contacts of index cases, it was possible to estimate the influenza pathogenicity and the antiviral efficacy in reducing pathogenicity, AVE_P. In the contacts receiving prophylaxis, 10 symptomatic cases occurred in 46 infected people. In the contacts not receiving prophylaxis, 33 symptomatic cases occurred in 75 infected people (Halloran et al 2007). Compute the pathogenicity of the influenza virus in the two groups and the efficacy of postexposure prophylaxis against progression, AVE_P.