7

Modes of Action and Time-varying VE_S

7.1 Mode of action and choice measures

7.1.1 Type I and Type II modes of action

In Chapter 6, vaccine efficacy results were reported simply based on relative cumulative incidence or rates without any further interpretation of the meaning of the efficacy estimate. Suppose you have just been vaccinated against an infectious agent. Your physician or health practitioner tells you that the protective efficacy of the vaccine is 90 percent. You might then wonder if that means that the vaccine reduces your probability of contracting the infection (or disease) by 0.90 at each exposure to infection. In other words, you still might have a finite probability of contracting the infection or disease each time you are exposed, but it is much less than it would have been if you had not been vaccinated. Alternatively, you might think that it means that you have a 0.90 probability of being completely protected against the disease, but still a 0.10 probability that you received absolutely no protection against infection or disease compared to what if would have been had you not been vaccinated. That is, the vaccine fails to elicit a protective immune response in 10% of the vaccinated people. Would you behave differently if you knew which of these possibilities were actually true? Would it make a difference if the vaccine were against a disease with a high fatality ratio? Would it make a difference if the efficacy were 60 percent rather than 90 percent?

In 1984, Smith, Rodriguez, and Fine published a paper that grew out of a student exercise that altered the discussion about interpreting and evaluating protective efficacy of vaccines. They consider two models of vaccine mechanism they called Type I and Type II. In the Type I mechanism, vaccination is assumed to reduce the instantaneous disease rate in all the vaccinated people by a constant proportion. That is, Type I assumes that protection is multiplicative on the baseline hazard of infection. The effect is homogeneous in the vaccinated population. In the Type II mechanism, vaccination is assumed to provide a constant proportion of individuals with complete immunity from

the disease. That is, it completely protects a portion of the vaccinated people, but completely fails to protect in the other portion. Under the Type II mechanism, the distribution of protection is heterogeneous in the vaccinated population.

Smith et al (1984) consider how these assumed models affect the choice of analysis in cohort studies and sampling of controls in case-control studies. Their discussion considered two of the measures of vaccine efficacy that do not condition on exposure to infection, that is vaccine efficacy measures based on Level II and III information and vaccine efficacy based on level IV information. The first, which they considered as one, was based on the hazard rate or cases per person-time at risk, $VE_{S,\lambda}$ or $VE_{S,IR}$. The second was based on the number of cases per person at risk, the cumulative incidence or attack rate, $VE_{S,CI}$ or $VE_{S,AR}$. Their motivation was originally in the design of alternatives to randomized studies, but the results for the cohort studies apply as well to randomized controlled trials.

7.1.2 Leaky and all-or-none modes of action

Before continuing the discussion of the implications of the two models of vaccine action for choice of measures in vaccine studies, we divert to explain why we prefer the use of the term *leaky* for Type I and *all-or-none* for Type II models. In the early 1980's the possibility of developing effective malaria vaccines created a great deal of excitement. The malaria parasite has a complex life cycle with separate antigenic stages. Malaria sporozoites are injected by the mosquito into the human and are the stage infective for humans. Asexual blood stages, or merozoites, subsequently develop and are responsible for malaria disease. Sexual blood stage parasites, or gametocytes, develop from the asexual blood stage and are the stage that are infective for mosquitos. Malaria vaccines were being developed against each of the three main stages, so vaccine candidates were directed at blocking infection, modifying disease once infected, and blocking transmission to the mosquito, corresponding to VE_S, VE_P, and VE_I. A sporozoite vaccine was expected to prevent infection either by inhibiting invasion of liver cells or by impairing effective reproduction once the parasite was in the liver. If the inhibition were not complete, then essentially the liver would let parasites through and be leaky. Struchiner, Halloran, and Spielman (1989) and Halloran, Struchiner and Spielman (1989) developed models of malaria vaccination that separately considered its effect on infection, disease, and transmission to mosquitos. The mechanism of the vaccine model's effect on susceptibility to infection reduces the probability of infection given a bite by an infected mosquito, corresponding to the expected direct effects of a leaky sporozoite vaccine that does not provide sterile immunity. Thus the term "leaky" for a multiplicative effect on the transmission probability comes from the image of the malaria parasites getting through a leaky immunity in the liver. Halloran et al (1989) and Struchiner et al (1989)

also considered waning of immunity and the role of natural boosting of infection in their dynamics models.

In a study of the use of case-control studies under complex disease transmission patterns, Struchiner et al (1990) adopted the term "leaky" rather than Model 1 as suggested by Smith et al (1984). The motivation was partly because it is more descriptive than the term Model 1, partly because it also was meant to take into account the effect on the transmission probability, and partly because the approach grew out the malaria vaccine research. Smith et al (1984) did not discuss exposure to infection or any biological mechanism for the different models of vaccine action. There was considerable resistance in the early 1990's in parts of the vaccine community against the term leaky because of its potentially negative connotations. However, as recognition increased that vaccines often protect more against disease than infection, the term leaky has gained wider acceptance.

We also prefer the term "all-or-none" (Halloran et al 1991) to Model 2 for a vaccine that protects a portion of the vaccinated people completely and the rest of the vaccinated not at all because it is more descriptive. As early as 1915, Greenwood and Yule discussed possible heterogeneities in susceptibility in the vaccinated and unvaccinated groups. Correlation of the antibody response with the distribution of infection rates also suggests that there is heterogeneity in protective response. For example, in the first 17 months of follow-up in a hepatitis B vaccine trial, 10 of the 11 infections in vaccinees accrued in the hypo- or nonresponders (Francis et al 1982). In a live virus varicella vaccine trial, incidence in the 17 percent of vaccinees with low antibody titer was between 5 and 13 percent per year, whereas in vaccinees with high antibody titer, incidence averaged less than 2 percent per year (White et al 1992). We use the terms leaky and all-or-none rather than Model 1 and Model 2 while wanting to give credit to Smith et al (1984) for the important and interesting discussion begun with their arguments presented in the next section.

7.1.3 Implications for choice of efficacy measures

Consider a randomized controlled trial with equal numbers of individuals in the placebo and vaccinated groups with both groups followed for an equal period of time. In this simple example, assume there is no loss to follow-up or deaths, that all cases of disease are ascertained and the time of onset of each case is known (Table 6.1). The measures of interest are based on the hazard rate or cases per person-time at risk, $\mathrm{VE}_{S,\lambda}$ or $\mathrm{VE}_{S,IR}$, and the cumulative incidence or attack rate, $\mathrm{VE}_{S,CI}$ or $\mathrm{VE}_{S,AR}$. If the incidence rates and the attack rates are low, then the two measures will be approximately equal, and it makes little difference which measure is used to compute vaccine efficacy, whereby the approach using the hazard rate or person-time at risk does allow for different follow-up times. However, in many cases the appropriate choice of vaccine efficacy measure may depend on whether the mode of action is leaky or all-or-none.

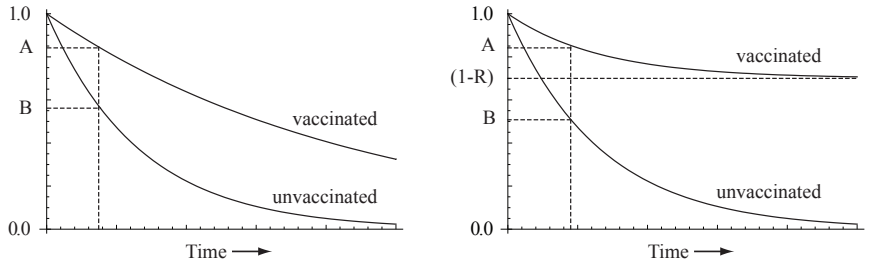

Fig. 7.1. Proportion of individuals without disease by time since start of trial under two models of vaccine action, (a) leaky, (b) all-or-none (adopted from Smith, Rodriguez, and Fine 1984).

For the leaky model, Smith et al (1984) considered an essentially continuous time model. Suppose that in a small interval of time $(t, t+\delta t)$ the probability of an unvaccinated person contracting disease is $\lambda_t \delta t$. Also, suppose that vaccination reduces the probability to $R\lambda_t \delta t$, with R assumed constant over time. Consider also that $\lambda_t = \lambda$ for all t, although this assumption is not necessary. Figure 7.1a shows the survival curves for the unvaccinated and vaccinated groups. The proportion of individuals in each group who would be expected to develop disease by time T would be $1 - e^{-\lambda T}$ and $1 - e^{-\lambda RT}$. Thus the calculated $VE_{S,CI}(T)$ given here as a function of T (Greenland and Frerichs 1988) is

$$VE_{S,CI}(T) = 1 - \frac{1 - e^{-\lambda RT}}{1 - e^{-\lambda T}} \ . \tag{7.1}$$

$VE_{S,CI}(T)$ in equation (7.1) decreases to zero as the follow-up time T increases. That is, this model allows everyone to get disease if the follow-up time is long enough. However, based on $VE_{S,IR}(T) = 1 - (c_1/Y_1)/(c_0/Y_0)$, or $VE_{S,\lambda}$, it is easy to show (see Problem 7.1) that $VE_{S,IR}(T)$ does not change with time, thus

$$VE_{S,IR}(T) = 1 - R. \tag{7.2}$$

Under the all-or-none model, the assumption is that vaccination provides a proportion $(1 - R)$ of the vaccinated group with complete immunity from the disease. The probability of an unvaccinated person contracting disease in the small time interval $(t, t+\delta t)$ is still $\lambda_t \delta t$, and once again it is simple to assume that $\lambda_t = \lambda$ for all t. In the people who were vaccinated but in whom the vaccine provides no protection, the probability of contracting disease in a short interval is the same as in the unvaccinated people, $\lambda \delta t$. Figure 7.1b shows the survival curves for the vaccinated and unvaccinated groups under the all-or-none model. From the initiation of the trial up to time T, the proportions in each group expected to have developed disease would be $1 - e^{-\lambda T}$ and $1 - (1 - R) - Re^{-\lambda T} = R(1 - e^{-\lambda T})$.

$$\mathrm{VE}_{S,CI}(T) = 1 - R. \tag{7.3}$$

$$\mathrm{VE}_{S,IR}(T) = 1/[1 + R(1 - e^{\lambda T})/T\lambda(1 - R)]. \tag{7.4}$$

So under the all-or-none model, the time invariant measure of vaccine efficacy is $\mathrm{VE}_{S,CI}(T)$. The value of $\mathrm{VE}_{S,IR}(T)$ or $\mathrm{VE}_{S,\lambda}(T)$ will tend to one as the people in the vaccinated group who are still susceptible to disease are depleted, leaving only those who are completely immune.

Smith et a (1984) showed that if in a randomized study $\mathrm{VE}_{S,CI}(T)$ decreases with time, but $\mathrm{VE}_{S,IR}(T)$ ($\mathrm{VE}_{S,\lambda}(T)$) remains constant, the result would be suggestive of a leaky, multiplicative mechanism. On the other hand, if $\mathrm{VE}_{S,CI}(T)$ is constant, but $\mathrm{VE}_{S,IR}(T)$ increases with time, the result would suggest an all-or-none mechanism. This result is the same for randomized prospective studies or observational cohort studies. Other mechanisms could explain time-varying efficacy estimates. In Section 7.3, we consider the situation that the efficacy within individuals actually does wane with time, which provides a biological mechanism for a time-varying vaccine efficacy. Also, heterogeneities in exposure to infection could play a role. In Chapter 8, we consider case-control studies, including the findings of Smith et al (1984). More general distributions of protection were developed (Halloran et al 1992) and Brunet et al (1993) developed a method of estimation based on state space models.

7.1.4 Attack rates versus transmission probabilities

Suppose that the infection process occurs as discrete exposures to infection (Halloran et al 1991) rather than in continuous time models as in Smith et al (1984). The question then is to define a direct protective effect of vaccination given a specific amount of exposure to infection, not just comparable exposure to infection. To be biologically interpretable and to be robust to different transmission conditions, the parameters of interest might need to take account of the type and amount of exposure. The following argument shows how vaccine efficacy measured using the attack rate can depend on the number of exposures to infection, thus could vary from population to population (Halloran et al 1991).

Let p_0 be the probability of transmission to an unvaccinated person after one exposure. Let $p_1 = \theta p_0$ be the probability of transmission to a vaccinated person after one exposure, where θ is the multiplicative, leaky effect on the transmission probability in the vaccinated person. Let $\mathrm{AR}_1(n)$, $\mathrm{AR}_0(n)$ and $\mathrm{VE}_{AR}(n)$ denote the attack rates and vaccine efficacy based on the attack rates that would be observed after everyone had n exposures to infection. Assume that everyone in the population receives actually one exposure to infection, and that there are N_0 and N_1 individuals in the unvaccinated and vaccinated groups. Then the attack rates in the vaccinated and unvaccinated groups are

$$AR_1(1) = \frac{p_1 N_1}{N_1} = p_1 = \theta p_0, \quad AR_0(1) = \frac{p_0 N_0}{N_0} = p_0 \ ,$$

so that the $\mathrm{VE}_{AR}(1)$ and VE_p are the same,

$$VE_{AR}(1) = 1 - \frac{AR_1(1)}{AR_0(1)} = 1 - \frac{p_1}{p_0} = VE_p = 1 - \theta. \tag{7.5}$$

Now assume that everyone in the population is exposed to infection a second time. We assume a discrete model of infection and that each exposure is independent of the previous exposures. The attack rates in the unvaccinated would now be given by the probability of having been infected by the first infective plus the probability of being infected by the second infective given that a person was not infected by the first infective. Then

$$AR_1(2) = p_1 + (1 - p_1)p_1 = p_1(2 - p_1) = \theta p_0(2 - \theta p_0) \tag{7.6}$$
$$AR_0(2) = p_0 + (1 - p_0)p_0 = p_0(2 - p_0) \tag{7.7}$$

Thus, after two exposures,

$$VE_{AR}(2) = 1 - \frac{p_1(2 - p_1)}{p_0(2 - p_0)} = 1 - \frac{\theta(2 - \theta p_0)}{(2 - p_0)}, \tag{7.8}$$

so that $\mathrm{VE}_{AR}(2) < \mathrm{VE}_p$. In general, for n exposures to infection,

$$VE_{AR}(n) = 1 - \frac{1 - (1 - p_1)^n}{1 - (1 - p_0)^n} = 1 - \frac{1 - (1 - \theta p_0)^n}{1 - (1 - p_0)^n}, \tag{7.9}$$

It can be shown by induction that for $n > 1$, $\mathrm{VE}_{AR}(n) < \mathrm{VE}_p$.

Example

Suppose a vaccine has a multiplicative, leaky effect that is the same in everyone and reduces the probability of transmission per potentially infective exposure by 80 percent, $\mathrm{VE}_p = 0.80$. Then the transmission probability in vaccinated people would be 20 percent of that in unvaccinated people, so that $p_1 = 0.20 p_0$. Suppose we want to evaluate the efficacy of the vaccine in a study population of 2000, where 1000 individuals are vaccinated and 1000 are not. Assume for this disease that $p_0 = 0.25$, so that $p_1 = 0.20 \cdot 0.25 = 0.05$. At the end of one month, assume that every person in the study has had exactly five exposures to infection. What is the expected attack rate in each group and the $\mathrm{VE}_{AR}(5)$ after one month?

In the unvaccinated group, the probability of becoming infected is $1 - (1 - p)^5 = 1 - 0.75^5 = 0.76$, so the expected number of infections in the unvaccinated group is 1000 people $\times 0.76 = 760$. In the vaccinated group, the probability of becoming infected after five exposures is $1 - (1 - 0.05)^5 = 1 - 0.95^5 = 0.23$, so the expected number of infections in that group is 1000

people $\times 0.23 = 230$. Then $VE_{AR}(5) = 1-(230/1000)/(760/1000) = 1-0.30 = 0.70$, which is lower than the vaccine effect on the transmission probability, $VE_p = 0.80$.

Suppose that after two months, each individual has had exactly ten exposures. Now the expected number of infections in the unvaccinated group is $(1-0.75^{10})\times 1000 = 943$, and in the vaccinated group, it is $(1-0.95^{10})\times 1000 = 401$. After ten exposures, the $VE_{AR}(10) = (401/1000)/(943/1000) = 0.57$. The vaccine seems less efficacious after two months even though the effect of the vaccination on the transmission probability has not waned.

As the number of exposures in the two groups increases, the observed vaccine efficacy based on the attack rate will decrease to zero. Eventually everyone in both groups will become infected under the multiplicative assumption if they are exposed often enough, illustrating the meaning of a multiplicative or leaky model at the transmission probability level. In principle, people can still become infected if exposed often enough.

Suppose we use the model in continuous time (Chapter 4.3.3) and similar to Smith et al (1984), but take into account the number of exposures to infection. Assume that c is the contact rate with infectives. Then $\lambda_0 = cp$ and $\lambda_1 = 0.20cp$. In continuous time, $VE_{S,\lambda} = 1 - \frac{\lambda_1}{\lambda_0} = 0.20$, giving the same answer as the multiplicative effect on the transmission probability. In the unvaccinated group, the probability of being infected after five exposures in the first month is $1 - \exp(-5 \cdot 0.25) = 0.713$ and in the vaccinated group is $1 - \exp(-5 \cdot 0.05) = 0.221$, so the expected number of infections is 713 in the unvaccinated group and 221 in the vaccinated group. The number of expected infections is different than calculated above from the discrete model. The observed $VE_{S,AR}(5) = 1 - 0.221/0.713 = 0.69$, similar but not identical to that calculated from the discrete model. After 10 exposures, the $VE_{AR}(10) = 1 - (0.393/0.918) = 0.57$, the same as using the discrete model, though the expected number infected in the vaccinated and unvaccinated groups are different when calculated using the discrete model above.

7.2 Frailty mixture models for $VE_{S,\lambda}$

7.2.1 Mixing models

In this section we consider estimation and interpretation of vaccine efficacy when the distribution of protection can include some people who are completely protected, some who have no protection, and the rest having a continuous distribution of protection (Longini and Halloran 1996; Halloran et al 1996). A frailty model is a survival analysis model that allows for unmeasured heterogeneity in the population. The frailty mixing model developed here falls into the general category of frailty models (Vaupel 1979) used in survival analysis, but like cure models (Farewell 1982), it allows for a point mass at 0.

In the following development it is important to distinguish between heterogeneity in the distributions of the hazard rates and heterogeneity in vaccine effects. Assume that the heterogeneity in susceptibility in the unvaccinated and vaccinated groups is described by the nonnegative mixing random variables Z_0 and Z_1. Assume that the distribution of susceptibility in each group is such that α_0 and α_1 are the proportion of people in each group that are highly protected, that is, not susceptible to infection such that Z_ν, $\nu = 0, 1$ has point mass α_ν at 0. The susceptibility in the susceptible proportion follows a continuous distribution $f_\nu()$ with probability $1 - \alpha_\nu$. Thus,

$$\begin{aligned} &P(Z_\nu = 0) = \alpha_\nu, \\ &Z_\nu | Z_\nu > 0 \equiv X_\nu \sim f_\nu(), \quad \text{with probability} \quad 1 - \alpha_\nu. \end{aligned} \tag{7.10}$$

The distribution f_ν allows flexibility to model the shape and spread of the continuous part of the distribution of Z_ν. However, in the estimation problem here the mean is not identifiable. Thus, let $f_\nu()$ be from a two-parameter family, but with $E(X_\nu) = 1$. Furthermore, let $\text{var}(X_\nu) = \delta_\nu$. Then $E(Z_\nu = 1 - \alpha_\nu$, and $\text{var}(Z_\nu = (1 - \alpha_\nu)(\delta_\nu + \alpha_\nu)$.

An example of the distribution of susceptibility in the vaccinated and unvaccinated groups if X_ν follows a gamma distribution is shown in Figure 7.2. In this example, $\alpha_0 = 0.1$ and $\alpha_1 = 0.5$. The expectation of the random variable in the susceptible proportion of each group equals one. In the vaccinated group, the susceptibility is reduced by the factor $\theta = 0.5$ in the people still susceptible, The area under each curve of susceptibles is α_0 and α_1 in the unvaccinated and vaccinated groups. For a vaccine that highly protects some people while conferring partial protection on the rest, there are several measures of vaccine efficacy. The difference between the proportion highly protected in each group, $VE_\alpha = \alpha_1 - \alpha_0$, measures the proportion of the population highly protected due to vaccination. The measure $VE_\theta = 1 - \theta$ is the efficacy of the vaccine in conferring partial protection conditional both on a specified exposure to infection and on remaining to some degree susceptible.

The summary measure of protective vaccine efficacy is the expected relative reduction in susceptibility conferred by the vaccine at the beginning of observation,

$$\text{VE}(0)_{S,SUM} = 1 - \frac{(1 - \alpha_1)\theta}{1 - \alpha_0}. \tag{7.11}$$

If the $\alpha_0 = 0$, that is, no one in the unvaccinated group is completely protected, then the summary measure of vaccine efficacy under heterogeneity is

$$\text{VE}(0)_{S,SUM} = 1 - (1 - \alpha)\theta. \tag{7.12}$$

7.2.2 Frailty model

Following the ideas of dependent happenings (Chapter 2.4), let $P(t)$ be the infection point prevalence at time t. Then the individual level hazard rate to

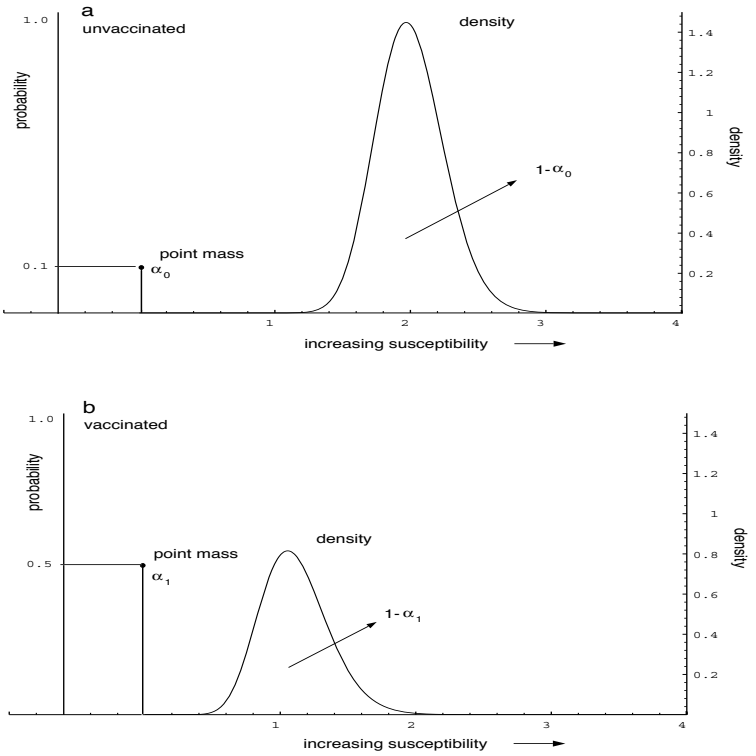

Fig. 7.2. Schematic distribution of susceptibility in the (a) unvaccinated and (b) vaccinated groups. The proportion highly protected is $\alpha_0 = 0.1$ in the unvaccinated and $\alpha_1 = 0.5$ in the vaccinated. The expectation of the random variable in the susceptible proportion of each group is equal to one. The area under each curve of susceptibles is $1 - \alpha_0$ and $1 - \alpha_1$ in the unvaccinated and vaccinated groups. In the vaccinated group the susceptibility is reduced by the factor $\theta = 0.5$.

an unvaccinated and vaccinated person at time t, respectively, is

$$\lambda_0(t) = Z_0 cpP(t), \quad \lambda_1(t) = Z_1 cpP(t). \tag{7.13}$$

To derive the survival function, let $S_\nu(t)$ be the fraction of the stratum ν that is considered to be at risk of infection at time t, $t \geq 0$. The assumption is made here that the population is closed to immigration, but open to emigration (that is, right censoring), so that $S_\nu(t)$ is a survival function. In addition, the assumption is made that vaccination takes place at or before time 0, and that the effects of vaccination do not wane over time. Then the population survival functions are

$$S_\nu(t) = E[\exp\{-Z_\nu \Lambda_\nu(t)\}] = L_{Z_\nu}\{\Lambda_\nu(t)\}, \tag{7.14}$$

where

$$\Lambda_0(t) = cp \int_0^t P(\tau) d\tau,$$

$\Lambda_1(t) = \theta \Lambda_0(t)$, and $L_Z(\ )$ is the Laplace transform (Aalen 1988, 1992). The Laplace transform of Z_ν is

$$L_{Z_\nu}(s) = \alpha_\nu + (1 - \alpha\nu) L_{X_\nu}(s). \tag{7.15}$$

If X_ν follows a gamma distribution with both scale and shape parameters equal to $1/\delta_\nu$, then from equations (7.14) and (7.15),

$$S_\nu(t) = \alpha_\nu + (1 - \alpha_\nu) \left\{ \frac{1}{1 + \Lambda_\nu(t)\delta_\nu} \right\}^{1/\delta_\nu} \tag{7.16}$$

When $\delta_\nu = 0$, then X_ν is degenerate at 1, and

$$S_\nu(t) = \alpha_\nu + (1 - \alpha_\nu) \exp\{-\Lambda_\nu(t)\}. \tag{7.17}$$

Statistical inference

This approach is for data in grouped survival form with observations made at times $t_0(=0), t_1, \ldots, t_k$. Define the time intervals as $[t_{i-1}, t_i)$, $i = 1, \ldots, k$. Then let $P(t)$ be piecewise constant on these intervals, where $P(t) = P_i$ in interval i. Then from equation (7.14),

$$\Lambda_0(t) = cp \int_0^t P(\tau) d\tau = cp\kappa \left\{ \sum_{j=1}^{i} (t_j - t_{j-1}) P_j + (t - t_i) P_i \right\}, \quad t \in [t_i, t_{i+1}),$$

where κ is a proportionality constant related to the proportion of a time interval that infected individuals are infectious (Halloran et al 1996). Here the P_i are treated as observed quantitities and not as parameters to be estimated.

The parameters to be estimated are c, p, κ, α_0, α_1, δ_0, δ_1, and θ. Set $a = cp\kappa$, since c and p cannot be separately estimated from data with no contact information (Rhodes et al 1996), and κ is simply a proportionality constant. To formulate the likelihood function for observations from the population under study, let $r_{i\nu}$ be the number of people at risk in group ν at the beginning of interval i, minus half those who are lost to follow-up during the interval i, and let $m_{i\nu}$ be the number infected during that interval. Then the likelihood function is

$$L(\text{data}|a, \alpha_0, \alpha_1, \delta_0, \delta_1, \theta) = \prod_{i=1}^{k} \prod_{\nu=0}^{1} \left\{ \frac{S_\nu(t_i)}{S_\nu(t_{i-1})} \right\}^{r_{i\nu} - m_{i\nu}} \left\{ \frac{S_\nu(t_i)}{1 - S_\nu(t_{i-1})} \right\}^{r_{i\nu}} \tag{7.18}$$

(see Aalen 1988). The likelihood function (7.18) can be maximized using standard methods.

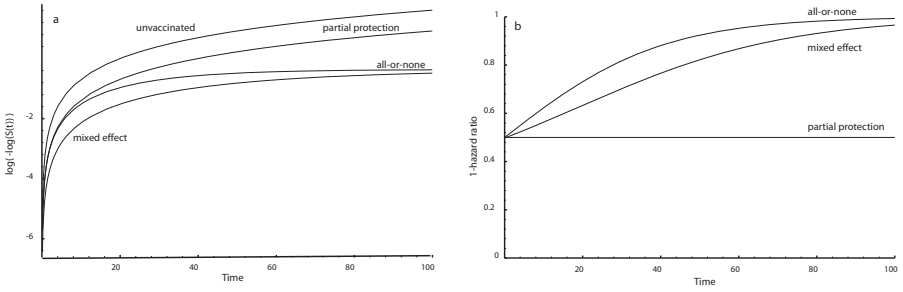

Fig. 7.3. a, diagnostic natural log minus log survival plots checking the proportional hazards assumption for vaccine conferring homogeneous partial protection, an all-or-none vaccine, and a mixed degenerate vaccine model compared with the unvaccinated group. b, plots of 1− hazard ratios for homogeneous partial protection ($\theta = 0.5$), the all-or-none vaccine ($\alpha_1 = 0.5$), and the mixed degenerate model ($\theta = 0.75, \alpha_1 = 0.33$). $VE_{SUM} = 0.5$ at time $t_0 = 0$ in these three cases (Halloran et al 1996).

Halloran, et al (1996) explored the potential use of the above described frailty mixture model for the estimation of $VE_{S,\lambda}$ over the parameter space that covers the possibilities of most vaccine studies. They showed that the parameters are identifiable under reasonable field conditions as long as there is not too much right censoring. Most importantly they showed that the conventional VE_S estimators based on proportional hazards and cumulative incidence, can be considerably biased when unmeasured heterogeneity is present. This bias is removed when the correct frailty mixture model is used. Violation of the proportional hazards assumption under frailty distributions is illustrated in Figure 7.3. The model is also applicable if there is heterogeneity in exposure to infection, though the interpretation of the estimates is different.

7.2.3 Measles outbreak in Burundi

A measles outbreak started in Muyinga, Burundi, in April 1988. The outbreak peaked in October 1988 and was over by December of that year (Chen et al 1994; Longini et al 1993). Measles illness histories were compiled after the outbreak for children aged 9 to 60 months. Only the month of onset of the measles illness was accurately recorded for most of the children (Table 7.1). Monthly measles incidence is given in the last column of Table 7.1. Initially 1,436 children had no previous history of measles illness and known measles vaccination status, that is, they had childhood immunization cards. Of these 1,436 children, 857 (60%) were vaccinated against measles before the outbreak. An additional 140 children were vaccinated during the outbreak. During the outbreak, 129 of the unvaccinated and 93 of the vaccinated children developed measles illness.

In the analysis of the data in Table 7.1, the vaccination times of the 140 children who were vaccinated during the outbreak were treated as right-

Table 7.1. Numbers at risk, ill and monthly exposure for the measles epidemic Muyinga, Burundi, April–November 1998 (from Longini and Halloran 1996).

		Unvaccinated			Vaccinated			Exposure
i	Month	At risk	Ill	%	At risk	Ill	%	$p \times 100\%$
1	April	579	10	1.7	857	9	1.1	1.3
2	May	551	13	2.4	848	13	1.5	1.9
3	June	517	13	1.9	835	2	0.2	0.9
4	July	483	12	2.5	833	20	2.4	2.4
5	August	451	22	4.9	813	18	2.2	3.2
6	September	408	50	12.3	795	24	3.0	6.4
7	October	337	12	3.6	771	7	0.9	1.7
8	November	317	0	0.0	764	4	0.0	0.0
Total			129		93			

censored times. The measles incidence in September was aberrantly high. There were only seven months, ($k = 7$), of measles incidence data, so δ_0 and δ_1 could not be estimated and were set to 0, so that the survival functions (7.17) were used in the likelihood function (7.18). In this case the summary vaccine efficacy is $VE_{SUM}(0) = 1 - (1 - \alpha_1)\theta$. The values calculated for $\{P_i\}$ were taken from the study population and are shown in Table 7.1. The maximum likelihood estimates and their standard errors are $\hat{a} = 1.66 \pm 0.14$, $\hat{\alpha}_1 = 0.805 \pm 0.060$, and $\hat{\theta} = 2.76 \pm 1.24$. The measles vaccine completely protected an estimated $\hat{\alpha}_1 = 0.805$,(95% CI, 0.687– 0.924) of the vaccinated children. The estimate of θ is greater than 1, suggesting that, assuming equal exposure, the relative per contact risk of contracting measles is higher in the vaccinated children who do not receive complete protection than it is in the unvaccinated children. The estimated summary measure of vaccine efficacy at time 0 is $\widehat{VE}_{SUM}(0) = 0.462$ (95% CI, 0.318–0671).

Table 7.2 gives the observed and expected (based on the fitted model) number of measles illness. The χ^2 goodness of fit statistic is 12.8 with 11 degrees of freedom. This yields a p-value of 0.3, so the model fits the data by this criterion. However, the distribution is not strictly χ^2 because of the correlation of the data over time. In future analyses, one might want to fit different models and use model selection tests such as likelihood ratio tests to choose among models (Hudgens and Gilbert 2008).

7.2.4 Model selection in low dose challenge studies

In human field studies, we generally cannot observe the actual number of potentially infectious contacts that each person makes. However, in a recent challenge study in macaques with an HIV vaccine candidate, repeated challenges were made and the infection status monitored (Ellenberger et al 2006). One of the 14 macaques in the control arm and four of the 16 in the vaccine

Table 7.2. Observed and expected frequencies for the model fitted to the data from the measles epidemic Muyinga, Burundi, April–November 1998 (from Longini and Halloran 1996).

		Unvaccinated		Vaccinated	
i	Month	Observed	Expected	Observed	Expected
1	April	10	12.6	9	9.8
2	May	13	16.7	13	12.8
3	June	13	7.6	2	5.8
4	July	12	19.1	20	14.7
5	August	22	23.0	18	16.7
6	September	41.0	12.3	24	27.2
7	October	12	9.4	7	6.0
Total		129		93	

arm were not infected when the study ended. Hudgens and Gilbert (2008) developed a clever method to distinguish using statistical methods whether the protective effect of the vaccine were leaky or all-or-none. The data are in the online supporting material of their paper. They used a discrete time survival model similar to equation (7.9), but also including a term for complete protection, making it a discrete time analogue of the summary measure of vaccine efficacy under heterogeneity in equation (7.12):

$$VE_S(n) = 1 - \frac{(1-\alpha)\{1-(1-\theta p)^n\}}{1-(1-p)^n}. \tag{7.19}$$

They developed maximum likelihood methods to estimate the transmission probability in the unvaccinated group, the proportion with complete protection, and the multiplicative effect on the transmission probability. They used a likelihood ratio test and the Akaike Information Criterion (AIC) to compare the leaky model, the all-or-none model, the summary model under heterogeneity, and the null model. They found that the statistical evidence suggested that the vaccine candidate had a significant leaky effect. These methods could be used for other repeated low dose challenge studies or studies where the exposures to infection are known. The power for detecting an all-or-none effect was observed to be greater than the power to detect a leaky effect.

7.3 Estimating waning efficacy

7.3.1 Waning efficacy in the cholera vaccine trial

Unmeasured heterogeneities in susceptibility, protection, and exposure to infection can produce time-varying estimates of $\mathrm{VE}_{S,IR}(t)$ or $\mathrm{VE}_{S,\lambda}(t)$ that are

Table 7.3. Piecewise constant RR estimates, with approximate 95% confidence intervals for the oral whole cell and oral B-subunit whole cell vaccines, Matlab, Bangladesh, May 1, 1985 to November 30, 1989 (Durham et al 1999)

Year	Dates	Whole cell vaccine		BS whole cell vaccine	
		RR	95% CI	RR	95% CI
1	May 1985–April 1986	0.44	[0.32–0.62]	0.33	[0.23–0.48]
2	May 1986–April 1987	0.45	[0.32–0.65]	0.47	[0.33–0.67]
3	May 1987–April 1988	0.55	[0.34–0.86]	0.86	[0.57–1.29]
4	May 1988–December 1989	1.21	[0.70–2.10]	0.83	[0.45–1.52]

a result of the underlying heterogeneities, while true waning of protection or boosting of protection can lead to real time-varying effects.

A traditional method to examine for waning vaccine efficacy over time has been to partition the time axis into time intervals and to assume that the efficacy is constant within each interval. Then a separate constant $VE_{S,IR}(t)$ is estimated for each time interval. If there is time dependence, then the estimates will vary across the time intervals.

This method of partitioning the time axis was used to estimate $VE_{S,IR}(t)$ of oral killed whole-cell (WC) and oral B subunit killed whole-cell (BS-WC) oral cholera vaccines of a randomized, double-blinded vaccine trial in Matlab, Bangladesh (Clemens 1990) (Chapter 6.4.6). In a longer term follow-up from May 1, 1985 to November 30, 1989, 580 cases of cholera occurred, with 284, 150, and 146 in the placebo, WC vaccine, and BS-WC vaccine groups. The efficacy of both vaccines appeared to wane. The methods used to analyze waning vaccine efficacy from this trial involved partitioning the study duration into discrete time units and comparing piecewise constant incidence rate ratio estimates for successive time periods (Clemens et al 1990, van Loon et al 1996). For example, Table 7.3 gives the piecewise constant incidence rate ratio estimates for the whole cell and B-subunit whole cell vaccines. The incidence rate ratio for each year is calculated by using a ratio of incidence rates, where the incidence among those vaccinated is compared with the incidence among the unvaccinated. The time period called year 4 includes 19 months of follow-up, after which there were no observed cholera cases. The incidence rate ratio (RR) estimates appear to increase, so the efficacy estimates decrease across the time period.

Thus, we see a waning time trend in efficacy, with no significant protection by the fourth year. However, because the data have been grouped into years, it is difficult to be more precise about when and how these changes in efficacy occur. Since the partitioning boundaries are selected at one year intervals, it is not clear if the waning protection is continuous or precisely at what point in time significant protection is lost. With use of a Poisson regression including covariates, the problem still remains of how to partition the time axis into

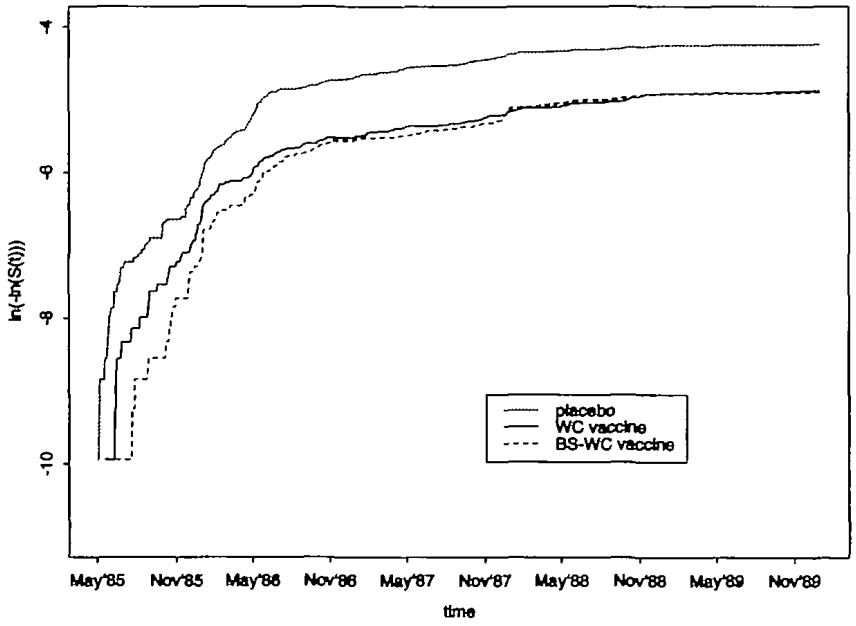

Fig. 7.4. Log-minus-log plots of the Kaplan-Meier estimates of the survival curves for the placebo and two cholera vaccines, Matlab, Bangladesh, May 1, 1985, through November 31, 1989. WC, killed whole cell; BS-WC, B subunit killed whole cell (Durham et al 1998).

piece-wise constant components. The problem can be solved by the use of survival analysis methods.

Figure 7.4 shows plots of the log minus log Kaplan-Meier estimates of the survival curves for the placebo and two oral cholera vaccines. The good separation between the vaccine and placebo curves indicates that the vaccines give protection. The BS-WC vaccine provides better protection during the first year. The curves slowly approach one another indicating the waning of the protective effect, but this is difficult to see with plots based on cumulative incidence.

7.3.2 Nonparametric estimation of time-varying vaccine effects

Durham, et al (1999) adapted and compared two basic approaches for the nonparametric estimation of smoothed curves for $\mathrm{VE}_{S,\lambda}(t) = 1- \mathrm{RR}(t) = 1 - \frac{\lambda_1(t)}{\lambda_0(t)}$. The first is a generalized additive models approach that involves using a time varying coefficient (Hastie and Tibshirani 1993) version of the proportional hazards model assuming a Poisson model (Whitehead 1980).

This approach is useful for diagnostics to ascertain the shape of $\beta(t)$, but it cannot provide an estimator for $VE_{S,\lambda}(t)$.

The other method uses Schoenfeld residuals (Schoenfeld 1982; Grambsch and Therneau 1994). The general idea is to fit an ordinary proportional hazards model to the data, then to compute the scaled differences between the actual and expected covariate values at each event time, called Schoenfeld residuals. The scaled residuals are added to the coefficient from the proportional hazards model. The time-varying regression coefficient $\beta(t)$ is recovered by smoothing the re-scaled Schoenfeld residuals. Conceptually, we are nonparametrically estimating the instantaneous hazard rate ratio $e^{\beta(t)}$, thus $VE_{S,\lambda}(t)$. Both methods provide a hypothesis test for the null $H_0 : \beta(t) = \beta$ for all t, that is, for no time-varying effects. The method using the Schoenfeld residuals is easy to use, provides an estimate of $e^{\beta(t)}$ on the natural scale, and allows easy incorporation of time-dependent covariates, so we recommend this approach in general.

Durham et al (1998) used the method involving Schoenfeld residuals to estimate smooth plots of the $VE_{S,\lambda}(t)$ for the two oral cholera vaccines from the cholera vaccine trial described in Section 6.4.6. Figure 7.5 shows the plot of the $VE_{S,\lambda}(t)$ estimates and the 95% CI's for the two vaccines. Table 7.4 gives the efficacy estimates and the approximate 95% confidence intervals for selected time points throughout the study. Age group (ages 2-5 years, >5 years) was included in the model as a covariate. The bending downward of the curves is indicative of waning. The p-values for the hypothesis test for departures from the proportional hazards assumption are 0.008 and 0.002 for the estimated model of the WC and BS-WC vaccines, respectively. The WC vaccine gives fairly constant and significant protection, with a $VE_{S,\lambda}(t)$ of about 0.50, for the first two and one half years of the trial, but then protection appears to wane rapidly. After three years of the trial (May, 1988), the point estimate of the $VE_{S,\lambda}(t)$ is 0.245 and the 95% CI covers zero. Protection from the WC-BS vaccine starts out higher than that from WC vaccine, *i.e.*, 0.713 versus 0.430, but then gradually wanes at a fairly constant rate, *i.e.*, about two - three percent per month. This analysis provides a more complete description of the VE_S than that based on yearly incidence ratios described above. Two further analyses studied the waning by age group and by biotype of the cases. Analysis and plots were done using modifications of the Splus functions `coxph` and `cox.zph`. Details are in the appendix in Durham et al (1998).

The results of this method must be interpreted carefully. Smoothed values at the beginning and end of the observation period are uncertain, with large CI's. This is a typical effect of smoothing which is exacerbated when the number of events decreases near the end of the observation period. For example, in the cholera vaccine trial, overall cholera incidence began to drop during the last year of the trial. Thus, the $VE_{S,\lambda}(t)$ estimates during the last year become unreliable. Nonetheless, a definite waning effect is apparent in Figure 7.5. This approach for estimating $VE_{S,\lambda}(t)$ provides a graphical inter-

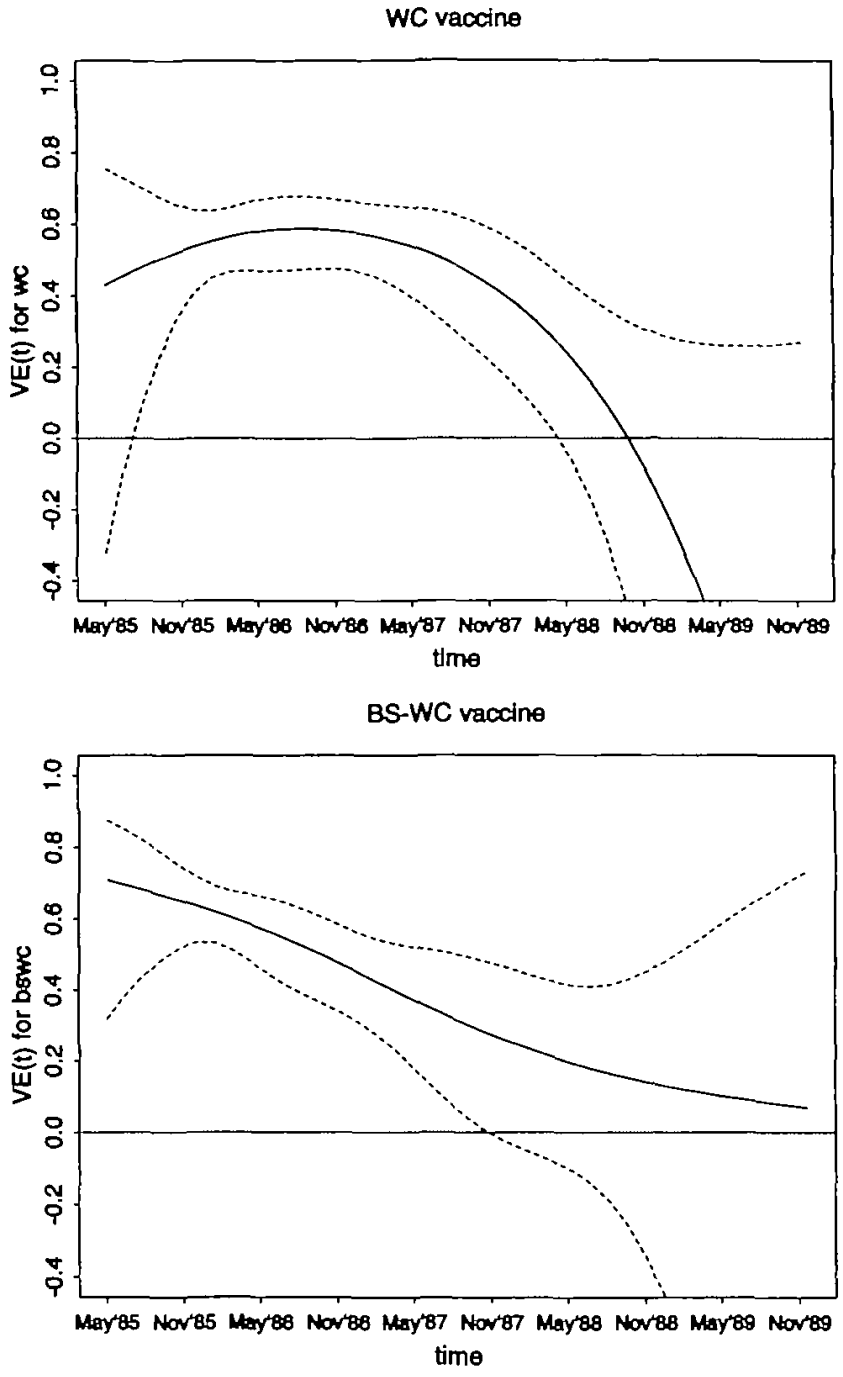

Fig. 7.5. Nonparametric smoothed plots of vaccine efficacy $\widehat{VE}_{S,\lambda}(t)$ versus time t, with 95% confidence intervals, for the killed whole cell (WC) and B subunit killed whole cell (BS-WC) vaccines, Matlab, Banglaesh, May 1, 1985, through November 31, 1989 (Durham et al 1998).

Table 7.4. Estimated vaccine efficacy over time, $VE_{S,\lambda}(t)$, with 95% confidence intervals for the the WC and BS-WC vaccines, Matlab, Bangladesh, May 1, 1985, through November 31, 1989. (durham1998)

Date	Day	Whole cell vaccine		BS whole cell vaccine	
		VE(day)	95% CI	VE(day)	95% CI
May 1985	0	0.430	−0.342–0.758	0.713	0.320–0.879
November 1985	183	0.525	0.356–0.650	0.650	0.523 –0.743
May 1986	365	0.579	0.467–0.667	0.572	0.457–0.662
November 1986	548	0.583	0.478–0.667	0.476	0.344–0.582
May 1987	730	0.538	0.394–0.648	0.374	0.176–0.524
November 1987	913	0.433	0.220–0.588	0.280	0.006–0.478
May 1988	1,095	0.245	−0.028–0.445	0.202	−0.089–0.416
November 1988	1,278	−0.073	−0.664–0.308	0.141	−0.338–0.448
May 1989	1,460	−0.590	−2.400–0.257	0.092	−0.955–0.578

pretation of time-varying vaccine effects as well as a test for departure from the proportional hazards assumption.

Plots of $\ln(-\ln(S(t)))$ are frequently used to assess graphically whether the proportional hazards assumption holds for time-to-event data. Since these are cumulative hazard function plots, they can fail to give a clear picture of time-varying effects that occur later in the study after a substantial number of events have occurred. Figure 7.4 provides a good illustration of this problem. The placebo and vaccine curves should be roughly parallel for all time if there were no time-varying effects. The early waning of protection of the BS-WC vaccine is readily apparent, especially with respect to the WC vaccine. The placebo and BS-WC are further apart around November, 1985 than they are after November, 1986. In addition, the protective effects of the two vaccine appear to be converging after November, 1986. It is harder to see the later waning of both vaccines after the first three years. Close inspection reveals that the vaccine and placebo curves are closer together around May, 1989, than they are around November, 1985. The estimated $VE_{S,\lambda}(t)$, Figure 7.5, gives a clearer picture of the time varying effects. If the proportional hazards assumption were valid, then these curves should be roughly straight lines, with zero slope rather than negative slope. In addition, the null hypothesis of a constant effect over time was rejected for both vaccines. This test, however, can be underpowered for small numbers of events.

7.3.3 Other approaches to estimate waning

Farrington (1992) reviewed some of the problems in estimating the occurrence and extent of waning vaccine protection. Kanaan and Farrington (2002) developed an approach to estimate vaccine efficacy in the presence of waning. The

model is an extension of the all-or-none and leaky model in equation (7.12) to allow for waning. They also focus on observational data, allowing people to be vaccinated over time and also allowing for underreporting. First, assuming that the vaccine does not wane, they formulate a version of the summary VE_S in equation (7.12)

$$\mathrm{VE}_S = 1 - (1 - \alpha)\theta, \tag{7.20}$$

where α is the proportion completely protected and θ is the leaky, or proportional hazards effect. Let a proportion π of the population be vaccinated, all at age τ. Then the model of waning for the all-or-none effect, assuming that $\theta = 1$, called a selection model, assumes that some people who were initially protected lose their protection. If the proportion initially protected when vaccinated is α_0, then one can model the proportion protected at time t after vaccination as

$$\alpha(t) = 1 - \alpha_0 \exp(-\alpha_1 t), \quad t \geq 0, \tag{7.21}$$

so that the age-specific vaccine efficacy at age x is

$$\mathrm{VE}_S(x, \tau) = (1 - \alpha_0) \exp(-\alpha(x - \tau)). \tag{7.22}$$

If $\alpha_1 = 0$, then the vaccine efficacy does not wane. In the deterioration model, the people who are initially completely protected are assumed to remain protected, but the leaky protection θ_0 in the initially partially protected people wanes with time. Under this model, age-specific vaccine efficacy is

$$\mathrm{VE}_S(x, \tau) = 1 - (1 - \alpha_0)\theta_0 \exp(-\theta_1(x - \tau)). \tag{7.23}$$

When $\theta_1 = 0$, the partial protection does not wane. The parameters α_0 and θ_0 represent the efficacy close to the age of vaccination τ. Kanaan and Farrington (2002) analyze two observational data sets of pertussis vaccination. The first is a cohort study of cases of pertussis in children born 1970–1986 done by a general practitioner from 1977–1987 in children 1 to 7 years old (Jenkinson 1988). The second was a case report study from the notifications of pertussis in the United Kingdom in 1989-1990, divided into an epidemic and a nonepidemic period. The vaccine coverage for each age group was known (Ramsay et al 1993). The data for both studies are given in Kanaan and Farrington (2002).

Parametric survival analysis methods are used to estimate the parameters of interest. Calendar time effects were taken into account by allowing for epidemic and nonepidemic periods, E_k and using a parametric approach to the baseline hazard. For the cohort data, the infection hazard is modeled both as an age-independent, time-dependent piecewise constant hazard $\lambda(a, t) = \rho_k$, $t \in E_k$, and as an age- and time-dependent gamma $\lambda(a, t) = \rho_k a \exp(-\beta a)$, $t \in E_k$. For the case report data, the age and time effects are confounded, so the assumption is that $\rho_k = \rho$. Parameters are introduced that allow for complete ascertainment, equal and possibly incomplete ascertainment in the vaccinated

and unvaccinated cases, incomplete ascertainment in the vaccinated cases only, and arbitrary differential ascertainment. The likelihood for the cohort model is an extension of equation (7.18) from Longini and Halloran (1996). The model for the case report data is an extension of the screening model of Farrington (1993).

Not too surprisingly, with this number of parameters, estimating all of the parameters and choosing the model that fits best is somewhat difficult. It was not possible to differentiate between waning of the all-or-none protection or waning of partial protection, but there was strong evidence of waning in the cohort data. In the case report data, there was near-complete lack of identifiability of the vaccine efficacy because of the negative correlation between the proportion completely protected and the ascertainment proportions. The approach to modeling the waning is still valid, and could be used in future observational data sets where estimation of waning of vaccine efficacy was of interest.

7.4 Summary strategy for estimating protective effects

7.4.1 Strategy

We present a general strategy for estimating $VE_{S,\lambda}(t)$ from time-to-event or incidence data (Halloran et al 1999). The first step is to conduct diagnostics. Then, with the help of the diagnostics, we find the best estimator of the VE_S. We begin by constructing log-minus-log plots of the Kaplan-Meier or actuarial estimates of the survival curves for the unvaccinated and vaccinated groups. These plots provide information about whether the vaccine effect is leaky, all-or-none or a mixture. In addition, they provide some information about whether vaccine induced protection is waning. If the curves are parallel, then the effect is mostly leaky (multiplicative), and we should model the vaccine effect with a proportional hazards model. Any divergence from parallelism indicates time-varying effects and the presence of some form of heterogeneity and/or waning protection. In this case, a model other than the proportional hazards model is needed. If the curves tend to diverge, then there is all-or-none effect and if they tend to converge, then the model still may be leaky, but with an unmeasured random effect (heterogeneity). Convergence could also indicate waning protection. Although construction of log-minus-log plots is an important first diagnostic step, they are sometimes difficult to interpret. If there are a sufficient number of events, a more informative plot is a smoothed hazard ratio plot of $VE_{S,\lambda}(t) = 1 - \lambda_1(t)/\lambda_0(t)$ as described in Section 7.3.2. The possible patterns associated with different vaccine effects are shown in Figure 7.3. A line with zero slope indicates a purely leaky or multiplicative effect. The researcher can construct a formal hypothesis test for zero slope (Grambsch and Therneau 1994; Durham et al 1998).

If there is no evidence of time varying effects from the diagnostics, then the $VE_{S,PH} = 1 - e^{\beta}$ can be estimated by fitting a proportional hazards model.

If there is evidence of time varying effects, then the investigator should fit the full family of frailty mixture models. If these models provide an adequate fit to the data, then the estimated parameters may be, but are not necessarily, the appropriate measures of the VE_S. If there is evidence of waning or other time-varying effects not attributable to unmeasured heterogeneity, then the nonparametric estimate of $VE_S(t)$ itself will provide the best estimate. In this case, it may be possible to construct a time-dependent parametric model of the $VE_S(t)$ that would provide tighter confidence intervals than the nonparametric approach.

7.4.2 Interpretation of measures

Which parameter to use to estimate VE_S in a particular study depends on the type and duration of the study, the infectious agent and its transmission mode, the resources available, and the assumptions of the distribution of protection within the vaccinated group. Even if time-dependent effects are detected, knowledge of the underlying biology will need to be used to interpret the effects and to help choose between actual waning, boosting, or heterogeneities. In many contemporary vaccine trials, immune response data are collected that can be used to help estimate and interpret vaccine effects. Also measuring actual or potential exposure to infection in individuals will help identify heterogeneities in exposure to infection. Some trials of vaccines for vector borne diseases have entomologic data. These help in quantifying potential exposure to infection. Trials of vaccines for vector borne diseases that collect entomologic data can use the information to help quantify potential exposure to infection.

Struchiner and Halloran (2007) show that randomization does not control for confounding in randomized vaccine trials, particularly when exposure to infection is an unmeasured confounder (Chapter 14). Differences in transmission intensity, previous exposure to infection, and pre-existing partial immunity and heterogeneities across communities result in different VE_S estimates, even when the actual biologic action of the vaccine is the same conditional on these factors. Reviews of pertussis vaccine trials in different populations using different estimators consider some of these issues (Fine and Clarkson 1987; Fine et al 1988). Given the above discussion, there are clear limits on the interpretability and generalizability of estimates of VE_S.

Problems

7.1. Models of vaccine action

(a) Show that under the leaky model in Smith, Rodriguez and Fine (1984) $E(c) = N(1 - e^{-\lambda RT})$, $E(Y) = N(\int_0^T e^{-\lambda RT})dt = N(1 - e^{-\lambda RT})/\lambda R$.

(b) Show that under the all-or-none model, $E(c) = N(1 - e^{-\lambda RT})$, $E(Y) =$

$N(\int_0^T [(1-R) + Re^{-\lambda T}]dt = N(T(1-R) + R(1-e^{-\lambda T})/\lambda$.
(c) Derive the results in equations 7.1 to 7.4 from (a) and (b).

7.2. Vaccine efficacy in cohort study
(a) Smith et al (1984) gave an example (Table A1 and A2)
(b) The second part of the problem is described here.

7.3. Vaccine efficacy in case-control study
(a) Smith et al (1984) gave an example (Table A3)
(b) The second part of the problem is described here.

7.4. The problem[1] is described here. The problem is described here. The problem is described here.

7.5. Problem Heading
(a) The first part of the problem is described here.
(b) The second part of the problem is described here.

[1] Footnote