Comparison of Two Smoothing Methods for Exploring Waning Vaccine Effects

L. Kathryn Durham; M. Elizabeth Halloran; Ira M. Longini, Jr.; Amita K. Manatunga

Applied Statistics, Vol. 48, No. 3 (1999), 395-407.

Stable URL:
http://links.jstor.org/sici?sici=0035-9254%281999%2948%3A3%3C395%3ACOTSMF%3E2.0.CO%3B2-L

Applied Statistics is currently published by Royal Statistical Society.

Appl. Statist. (1999)
48, *Part* 3, *pp.* 395–407

Comparison of two smoothing methods for exploring waning vaccine effects

L. Kathryn Durham

Pfizer Inc., Groton, USA

and M. Elizabeth Halloran, Ira M. Longini, Jr, and Amita K. Manatunga

Emory University, Atlanta, USA

[Received January 1998. Final revision September 1998]

Summary. We consider the statistical evaluation and estimation of vaccine efficacy when the protective effect wanes with time. We reanalyse data from a 5-year trial of two oral cholera vaccines in Matlab, Bangladesh. In this field trial, one vaccine appears to confer better initial protection than the other, but neither appears to offer protection for a period longer than about 3 years. Time-dependent vaccine effects are estimated by obtaining smooth estimates of a time-varying relative risk RR(t) using survival analysis. We compare two approaches based on the Cox model in terms of their strategies for detecting time-varying vaccine effects, and their estimation techniques for obtaining a time-dependent RR(t) estimate. These methods allow an exploration of time-varying vaccine effects while making minimal parametric assumptions about the functional form of RR(t) for vaccinated compared with unvaccinated subjects.

Keywords: Smoothing; Survival analysis; Time-dependent relative risk; Vaccine efficacy

1. Introduction

When evaluating the protective effect of a vaccine against infection or disease, it is important to consider whether the protection conferred on vaccinated individuals wanes with time. An initial vaccination can be followed with subsequent booster doses to maintain a protective level of immunity among susceptible individuals, but the nature of the protection over time must be understood so that an effective vaccination and boosting schedule can be implemented. One method that has been used to detect waning vaccine effects involves analysing periodic measurements of the level of various serologic indicators for the individuals under study (Gilks *et al.*, 1993). However, for many infections, the correlation between waning surrogate markers and waning protection against infection is not clearly understood. We consider the case where epidemiological data on the time to infection for both vaccinated and unvaccinated subjects are available, and we use survival analysis techniques to estimate the relative risk (RR) for vaccinated compared with unvaccinated individuals.

The goal of our analysis is to use a smooth estimate of RR(t) to examine the trend in the data while imposing a minimum of parametric restrictions on the functional form. Many different methods have been proposed to estimate the hazard ratio $\mathrm{RR}(t) = \exp\{\beta(t)\}$ as a nonparametric function of time using survival data, where $\beta(t)$ measures the time-dependent effect on survival (e.g. Hastie and Tibshirani (1993), McKeague and Sasieni (1994),

Address for correspondence: L. Kathryn Durham, Department of Biometrics, Clinical Research Division, Pfizer Inc., Eastern Point Road, Groton, CT 06340, USA.
E-mail: l_kathryn_durham@groton.pfizer.com

 0035–9254/99/48395

Kooperberg *et al.* (1995) and Abrahamowicz *et al.* (1996)). In this paper we compare two nonparametric approaches that obtain cubic instantaneous estimates of time-varying treatment effects in the Cox (1972) model framework. The first is based on the general class of varying-coefficients models discussed by Hastie and Tibshirani (1993), in which covariate effects are allowed to vary as functions of other variables. They developed a method for fitting the partial likelihood function while allowing the covariate effects to vary as a function of time, so that the coefficients $\beta(t)$ are estimated by using a ridge regression style method (see also Zucker and Karr (1990)). We take a computationally simpler approach, however, because our interest is primarily in estimating time-dependent effects for covariates that can be assigned the value 0 or 1 for vaccinated or control, or a series of indicator variables if the experiment involves multiple vaccines. For varying-coefficients models in the exponential family, if the covariates whose effects are modified by other variables are all factor variables, the model can be fitted by using the generalized additive model (GAM) framework (Hastie and Tibshirani, 1986, 1990, 1993). We extend the Poisson likelihood method of Whitehead (1980) to include arbitrary time-varying covariate effects, and we estimate smooth time-varying effects by using the GAM approach.

The second method is based on smoothed Schoenfeld partial residuals (Schoenfeld, 1982; Grambsch and Therneau, 1994) obtained from maximizing the log-partial-likelihood in the counting process formulation of the Cox model. By smoothing the Schoenfeld partial residuals against time, we can obtain estimates of the time-varying RR.

In Section 2 we describe a field trial of two oral cholera vaccines, where an estimation of the vaccine efficacy was the major objective. In Section 3 we describe how we obtained smooth $\beta(t) = \ln\{\mathrm{RR}(t)\}$ estimates in the GAM framework by using a Poisson likelihood. In Section 4, we briefly describe the method for obtaining nonparametric $\beta(t) = \ln\{\mathrm{RR}(t)\}$ estimates from the smoothed Schoenfeld partial residuals. In Section 5, we present a reanalysis of the cholera vaccine trial. In Section 6, we compare these two methods with respect to their power to detect a time-dependent vaccine effect.

2. Field trial of two oral cholera vaccines in Bangladesh

A field trial of two oral cholera vaccines was conducted in Matlab, Bangladesh, from May 1985 to November 1989 (see Clemens *et al.* (1990) for details). Of those initially enrolled in the study, 62285 subjects received three complete doses of either a placebo, whole cell or B-subunit whole cell vaccine, with 20837, 20743 and 20705 in each group respectively. 580 cases of cholera occurred, with 284, 150 and 146 cases in each group respectively. The outcome measured was the time until first cholera infection.

The methods used to analyse waning vaccine efficacy from this trial involve partitioning the study duration into discrete units and comparing piecewise constant RR estimates for successive time periods (Clemens *et al.*, 1990; van Loon *et al.*, 1996). For example, Table 1 gives the piecewise constant RR estimates for the whole cell and B-subunit whole cell vaccines. The RR for each year is calculated by using a ratio of incidence rates, where the incidence among those vaccinated (numerator) is compared with the incidence among those unvaccinated (denominator). Note that the time period which we are calling 'year 4' includes 19 months of follow-up, after which there were no observed cholera cases. The yearly estimates indicate that the efficacy for the whole cell vaccine may be fairly constant through the first three years and then may fall sharply in year 4. The efficacy for the B-subunit whole cell vaccine appears to fall at a more steady rate, with significant protection lost before some point during year 3.

Table 1. Piecewise constant RR estimates, with approximate 95% confidence intervals (CIs), for the whole cell and B-subunit whole cell vaccines, Matlab, Bangladesh, May 1st, 1985–November 30th, 1989

Year	*Dates*	*Whole cell vaccine*		*B-subunit 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)

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. Besides being sensitive to the choice of intervals used for comparing piecewise constant rates, this method does not evaluate whether the RR estimates from successive periods are significantly different. Our goal is to obtain smooth estimates for the RR as a function of time for the data from this vaccine trial, and to test whether the protection has waned with time.

3. Fitting the Poisson survival model in the generalized additive model framework

In this section we propose a method for estimating time-dependent effects from survival data by fitting a Poisson likelihood in the GAM framework. By showing that the partial likelihood function is proportional to the likelihood based on a Poisson formulation, Whitehead (1980) estimated the proportional hazards model parameters in the generalized linear models framework (see also McCullagh and Nelder (1989), pages 429–430). He showed that this model is strictly interpretable as a survival model for the case when the covariates are not time dependent and the survival functions are assumed to be continuous. However, the estimation technique still works if covariates are time varying. It is straightforward to extend Whitehead's derivation to include coefficients that are arbitrary functions of time suitable for estimation in the GAM framework. In the Poisson formulation, for each failure time t_h, $h = 1, \ldots, q$, let random variables $X_{h,r}$, where $r = 1, \ldots, k_h$, represent the number of failures in each failure time–covariate group. Then $X_{h,r}$ is Poisson distributed with mean parameter

$$\mu_{h,r} = N_{h,r} \exp\{\alpha_h + \mathbf{z}_{h,r}\, \boldsymbol{\beta}(h)\}, \tag{1}$$

where $N_{h,r}$ is the number at risk in covariate group r just before t_h, $\mathbf{z}_{h,r}$ are the covariates and $\boldsymbol{\beta}(h)$ are the time-varying effect estimates. The terms α_h in the Poisson model have the interpretation as the base-line hazard of failure in the short time interval $(t_h - \delta t_h, t_h]$.

In varying-coefficients models, the model is assumed to be linear in the predictors but the coefficients are allowed to change smoothly with time. We assume that the time-dependent effect $\beta(t)$ can be modelled by $\beta(t) = \beta + f(t)$, where β is the time invariant component and $f(t)$ is an arbitrary smooth function of time. Using the extension of Whitehead's (1980) Poisson likelihood approach, we can estimate the following varying-coefficients version of the proportional hazards model (Hastie and Tibshirani, 1990, 1993):

$$\lambda(t) = \lambda_0(t) \exp\left[\sum_{j=1}^{v} z_j\{\beta_j + f_j(t)\} + \sum_{s=v+1}^{p} z_s \psi_s\right]$$

where $z_j, j = 1, \ldots, v$, are 0–1 indicators of vaccine status whose effects may vary with time and z_s, $s = v + 1, \ldots, p$, are other general covariates whose effects on survival are time invariant. If the treatment effects do not vary with time, then the coefficients β_j will be adequate to estimate the effects on survival.

To estimate the parameters, we fit a Poisson GAM (Hastie and Tibshirani, 1990) using the S-PLUS function `gam()` (Statistical Sciences, 1995). To do this, the data set must be arranged so that it contains one observation for each failure time–covariate group combination, as indicated by equation (1). The number of events at each time in each covariate group is the response variable, a factor variable for time is used to estimate base-line terms α_h and the number at risk in each time by covariate group, $N_{h,r}$, is an offset term. We estimate smooth functions of time $f(t)$ by using smoothing splines with 4 degrees of freedom (see Chambers and Hastie (1992), p. 299). To test whether the nonparametric smooth terms $f(t)$ are needed to describe the treatment effects, the model is fitted both with and without these terms. The change in the Pearson χ^2-statistic is used to evaluate whether the smooth term contributes significantly to the model fit.

It is important to point out that the GAM model fitting algorithm centres the smooth at 0 at each step by subtracting the average of the smooth. This ensures that the constant in the additive model is identifiable (Hastie and Tibshirani (1990), p. 115). Although this allows the shape and size of the effect to be examined, it is not clear how to obtain a direct estimate of $\ln\{\mathrm{RR}(t)\}$. These aspects differ in the approach based on the Schoenfeld partial residuals that we present next.

4. Estimating smooth RR(t) functions by using scaled Schoenfeld residuals

Another technique for obtaining nonparametric estimates of $\mathrm{RR}(t)$ is to smooth scaled partial residuals from a proportional hazards model. This approach was introduced by Schoenfeld (1982) and later extended to the multivariate setting by Grambsch and Therneau (1994). Once again, the true time-varying coefficient $\beta(t)$ can be written as

$$\beta(t) = \beta + f(t),$$

where this now represents the sum of the coefficient from the standard proportional hazards model fit, β, plus a time-varying component $f(t)$. The idea behind this technique is that, if a model is fitted which ignores a time-dependent effect, the functional form for $f(t)$ will be captured in the Schoenfeld partial residuals from the misspecified model (for details, see Schoenfeld (1982) and Grambsch and Therneau (1994)). Estimates of $\mathrm{RR}(t)$ can be obtained from the smoothed plots of the residuals if there are enough events, and Grambsch and Therneau (1994) suggested testing for a linear association between the scaled Schoenfeld residuals and time to check whether an effect has a significant time-dependent component. The Schoenfeld residuals and the linear association test can be obtained by using S-PLUS functions `coxph()` and `cox.zph()` (Statistical Sciences, 1995).

5. Estimation methods for RR(t) by using data from a cholera vaccine trial

We applied the methods described above to estimate a smooth, time-dependent RR for vaccinated *versus* unvaccinated individuals using data from the cholera vaccine field trial presented in Section 2. All models contain effects for age less than or equal to 5 years and vaccine status. A more complex analysis by cholera biotype and age subgroups appears in Durham *et al.* (1998).

5.1. Exploratory graphical RR(t) analyses

We obtained the Schoenfeld residuals from a proportional hazards model fit assuming that the vaccine effect was constant over time. Fig. 1 displays the smooth estimates of $\ln\{\mathrm{RR}(t)\}$. The p-values for the test of linear association between the scaled residuals and time were 0.008 (whole cell) and 0.002 (B-subunit whole cell), indicating that the efficacy for both vaccines varies with time. Fig. 2 shows centred estimates of $\ln\{\mathrm{RR}(t)\}$ obtained by using the Poisson survival GAM approach. The χ^2-tests on 4 degrees of freedom had p-values of 0.021 (whole cell) and 0.010 (B-subunit whole cell), again indicating time-dependent effects for both vaccines. The plots obtained by using both techniques indicate that the protection associated with the B-subunit whole cell vaccine wanes at a steady rate throughout the study period. They also suggest that the effect of the whole cell vaccine may be roughly constant through the first 2 years of the study (to about day 730), after which the effect wanes linearly. The efficacy of both vaccines appears to wane such that, by the end of roughly 3 years, neither vaccine has a protective effect.

5.2. Smooth and parametric estimates of RR(t)

Because the $\mathrm{RR}(t)$ estimates obtained by using the Schoenfeld residual approach are on the original scale of the data, we can obtain a smoothed estimate of $\mathrm{RR}(t)$ for any day throughout the study. Table 2 contains estimates for RR to vaccinated individuals as a function of time at selected days during follow-up. Because the estimates from the GAM approach are not on the original scale for $\ln\{\mathrm{RR}(t)\}$, we cannot directly calculate smooth $\mathrm{RR}(t)$ estimates.

However, we can use either of these exploratory approaches to construct a parametric estimate of the functional form of log-RR to vaccinated individuals. The plots in both Fig. 1 and Fig. 2 for the B-subunit whole cell vaccine indicate that including a linear function of time should adequately model the time dependence. The plots of the functional form for the whole cell vaccine indicate that we can capture this effect in a parametric model by modelling the hazard ratio for the whole cell vaccine with no time dependence until day 730 (2 years), and including a linear effect which begins at day 730.

We also used the Poisson likelihood approach to fit a parametric model for time-dependent effects (Whitehead, 1980). We include two (0–1) indicator variables for either the whole cell (z_1) or B-subunit whole cell (z_2) vaccine, and we have two continuous time variables: t measures the survival time in days and t_{730} measures the number of days survived after day 730. For simplicity of presentation we report only the vaccine coefficients that are needed to calculate $\mathrm{RR}(t)$ from the parametric model in Table 3.

Table 2 contains estimates for $\mathrm{RR}(t)$ at some specific time points throughout the study, along with approximate 95% confidence intervals based on the parametric model. As in the $\mathrm{RR}(t)$ estimates obtained by using Schoenfeld residuals, the protective effect of both vaccines appears to have disappeared by the end of year 3. The parametric approach is helpful in providing a more stable estimate of $\mathrm{RR}(t)$ near the beginning of the study, but the confidence limits for the $\mathrm{RR}(t)$ estimates for day 365 and beyond are similar to those from the smoothed Schoenfeld residuals approach.

6. Simulation study: power to detect a waning vaccine effect

There are various mechanisms which can result in waning vaccine efficacy and therefore many possible functional forms for $\mathrm{RR}(t)$. We were also interested in using simulations to

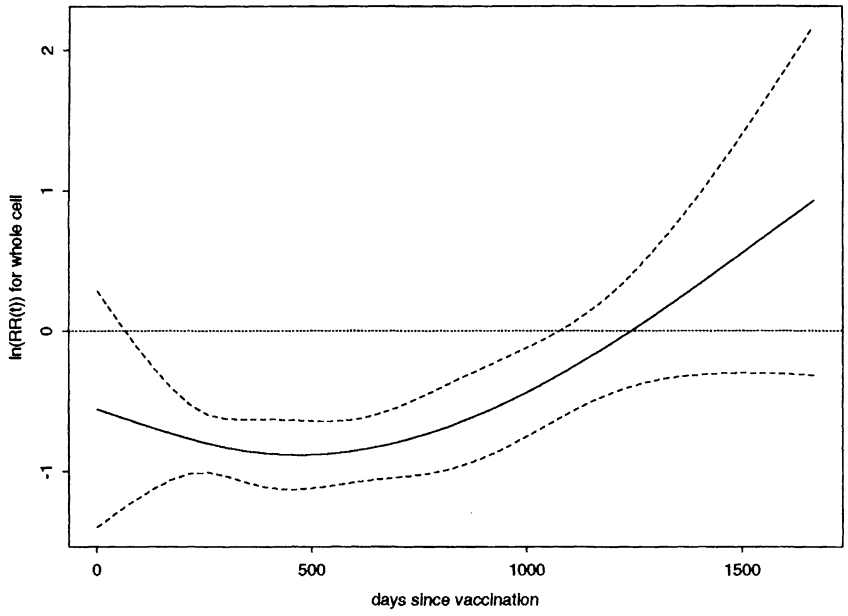

(a)

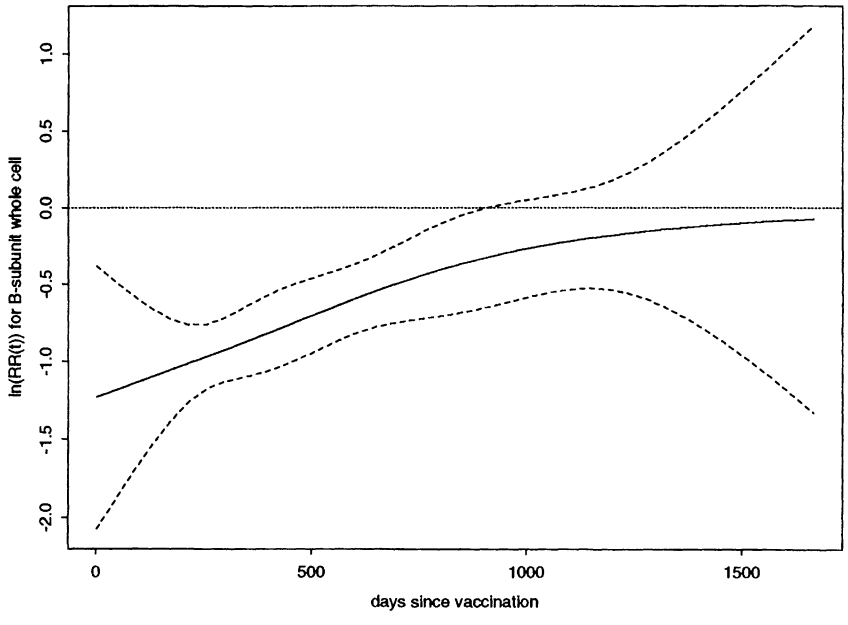

(b)

Fig. 1. Smooth $\ln\{\mathrm{RR}(t)\}$ functions from scaled Schoenfeld residuals: (a) whole cell vaccine; (b) B-subunit whole cell vaccine (- - - - -, approximate 95% pointwise standard error bands; ·····, $\ln\{\mathrm{RR}(t)\} = 0$ (corresponding to an RR of 1) for reference)

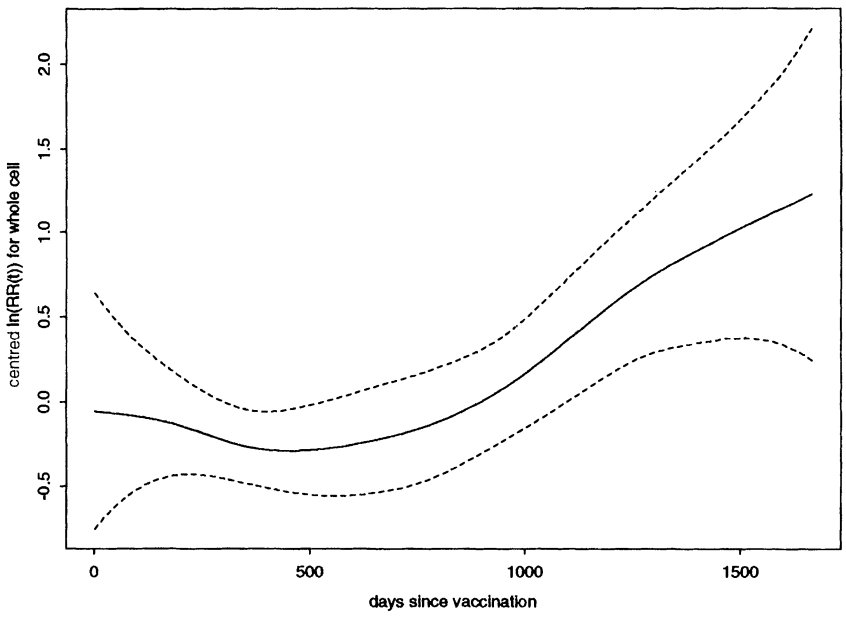

(a)

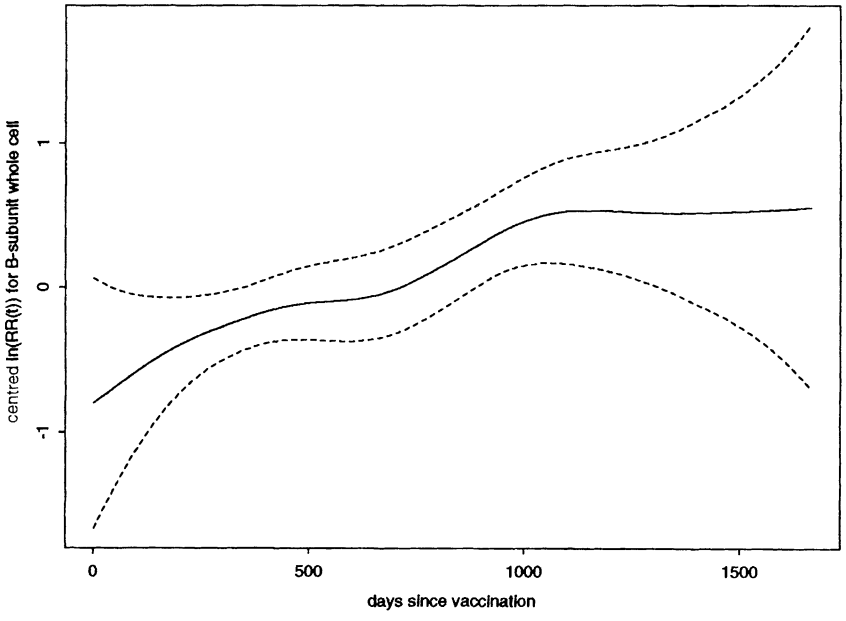

(b)

Fig. 2. Smooth estimate of the functional form of $\ln\{RR(t)\}$ by using the GAM approach: (a) whole cell vaccine; (b) B-subunit whole cell vaccine (- - - - -, approximate 95% pointwise standard error bands)

Table 2. Estimates of RR(t) from smooth scaled Schoenfeld residuals and from a parametric model

Day	*Whole cell vaccine*		*B-subunit whole cell vaccine*	
	RR(day)	*95% CI*	*RR(day)*	*95% CI*
Smoothed scaled Schoenfeld residuals				
0	0.570	(0.242, 1.342)	0.287	(0.121, 0.681)
365	0.422	(0.330, 0.534)	0.428	(0.338, 0.543)
730	0.462	(0.352, 0.606)	0.626	(0.476, 0.824)
1095	0.755	(0.555, 1.028)	0.798	(0.584, 1.089)
Parametric model				
0	0.436	(0.347, 0.547)	0.298	(0.203, 0.439)
365	0.436	(0.347, 0.547)	0.415	(0.330, 0.533)
730	0.436	(0.347, 0.547)	0.576	(0.477, 0.729)
1095	0.810	(0.597, 1.121)	0.800	(0.594, 1.156)

Table 3. Coefficient estimates for a parametric model for RR(t)

Effect	*Variable*	*Estimate*	*Standard error*
Whole cell	z_1	−0.8312	0.1159
Whole cell × days > 730	$z_1 \times t_{730}$	0.0017	0.0005
B-subunit whole cell	z_2	−1.2091	0.1973
B-subunit whole cell × days	$z_2 \times t$	0.0009	0.0003

compare the relative power of these two techniques for detecting different types of waning vaccine effects. An important difference between these two techniques is their strategy for testing whether a treatment effect should be modelled as a function of time. Both tests consider the null hypothesis that there is no time effect, i.e. that the proportional hazards assumption is valid. A possible advantage of the test for linear association used with the Schoenfeld technique is that, because it involves only the residuals, the degrees of freedom chosen to estimate the RR(t) function do not affect the test, but, because it is designed to diagnose linear departures, it may miss time-dependent treatment functions that do not exhibit a linear trend. However, although the test based on the change in model deviance used with the GAM approach may be affected by the degrees of freedom chosen to represent the time-varying effect, it should have more flexibility than the linear association test to identify a variety of RR(t) functions. We point out that, although there are several good tests for evaluating the proportional hazards assumption (e.g. Cox (1972) and Gill and Schumacher (1987)), we focus on the two tests associated with the smooth estimates presented in Section 5.

To explore possible strengths and weaknesses of the various strategies with respect to testing for time-varying effects, we generated two types of data set with known average RR(t). For the first type of data, we assumed that the vaccine efficacy wanes exponentially with time, so that

$$\mathrm{RR}(t) = 1 - (1 - \theta)\exp(-\omega t),$$

where θ is the initial RR to vaccinated individuals and ω is the rate at which protection is lost with time. In the second type of data, we assumed a piecewise constant form for RR(t). For

both types of data, we simulated a 3-year study where the number of days to infection was recorded for 1000 vaccinated and 1000 unvaccinated individuals. We used a base-line hazard function that was proportional to a prevalence

$$p(t) = 0.02 \exp\{\sin(6\pi t/1095)\},$$

to simulate the seasonal risk pattern associated with an infectious disease with environmental components such as cholera. This resulted in an average number of infections in the control group of about 750. The remaining uninfected individuals resulted from both a small constant probability of drop-outs during the study, accounting for about 5% loss due to type I censoring, and about 20% loss due to type II censoring. The number of events in the treatment group depended on the assumed initial treatment effect and rate of waning.

For the Poisson GAM model, we calculated the number of data sets out of 200 for which the χ^2-statistic was significant at the 0.05 level. A value of 4 degrees of freedom was used to fit all smooth functions. For the same data sets, we obtained the Schoenfeld residuals from the Cox model fit and calculated the number of data sets out of 200 for which the linear association between the smooth residuals and time was significant at the 0.05-level. We point out that here we are only comparing the *relative* power of the linear association test and the χ^2 change of deviance test for different RR(t) functions. The absolute values of power that are of interest will depend on the event rates for each particular study.

Fig. 3 shows the underlying shapes of RR(t) for some of the simulations presented in Table 4. Notice that values of $\omega = 0.02$ represent a rapid waning of vaccine protection, whereas values of $\omega = 0.0005$ represent a slow waning effect. When evaluating the efficacy of a vaccine, the case of a rapidly waning vaccine effect may not be very important for researchers, because the RR estimate should accurately reflect that the vaccine does not offer much protection. However, vaccine protection that wanes slowly over time could cause important problems. The duration of the study in such a case may not be adequate to detect these slow changes in efficacy, and vaccinated individuals who were initially protected could find themselves susceptible to infection at a much later time. Such issues need to be considered in the design phase of a vaccine study.

For the scaled Schoenfeld residuals and the Poisson survival GAM, Table 4 shows that the simulated power values were quite similar for the two different strategies, with a few exceptions. Both tests are most effective for intermediate values of ω, or when vaccine protection is lost at a moderate rate during the study. The χ^2-test on the change of deviance for the GAM appears to have somewhat better power in the case of an initially protective treatment effect that wanes quickly ($\omega = 0.02$), but the linear association test performed as well or better in the case of an effect that wanes slowly ($\omega = 0.0005$). This is not surprising, as the RR(t) functions with $\omega = 0.0005$ display a linear shape, whereas the RR(t) functions with $\omega = 0.02$ are more complex (see Fig. 3). However, simulated power values were identical for both tests with $\omega = 0.005$, in which the RR(t) function still displays a large degree of curvature. Thus for these simulations it appears that the increased flexibility of the GAM test does not necessarily result in increased power to detect exponentially waning RR(t) effects.

Table 5 gives simulation results for piecewise constant waning, in which the RR(t) function is constant except for a shift in the RR at a designated 'changepoint' of the simulated 3-year study. These shifts were simulated to occur after the first 6 months (day 183), the midpoint of the study (day 548) and after the first two and a half years (day 913). In Table 5, θ_1 represents the RR at the beginning of the study, and θ_2 represents the RR after the changepoint. The simulations indicate that the two tests have very comparable power to detect both large and

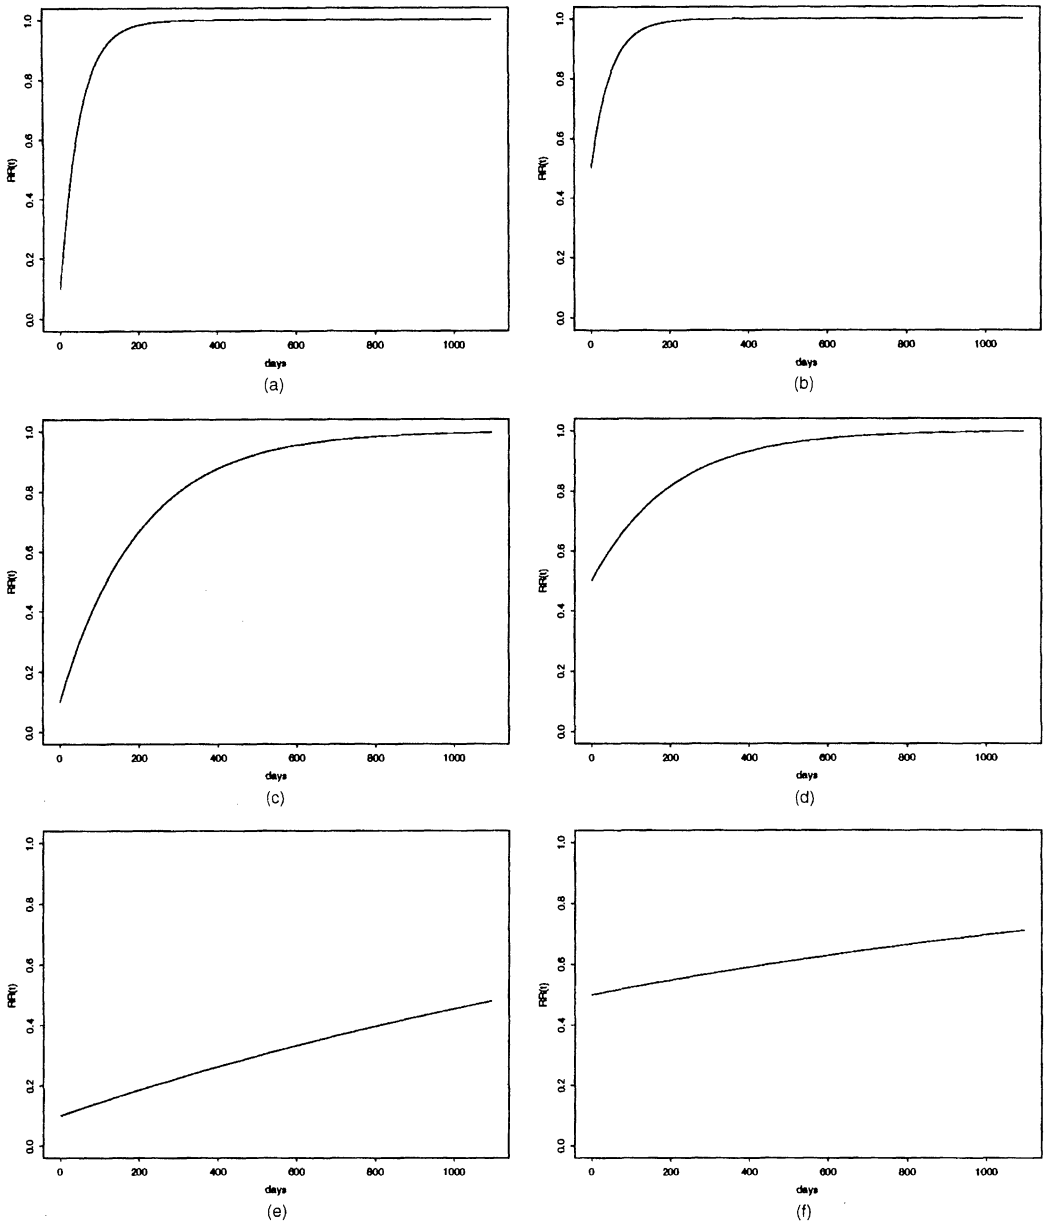

Fig. 3. Different shapes for exponentially increasing RR(t): (a) $\theta = 0.1$, $\omega = 0.02$; (b) $\theta = 0.5$, $\omega = 0.02$; (c) $\theta = 0.1$, $\omega = 0.005$; (d) $\theta = 0.5$, $\omega = 0.005$; (e) $\theta = 0.1$, $\omega = 0.0005$; (f) $\theta = 0.5$, $\omega = 0.0005$

small jumps in the RR at (any) point throughout the study. As with an exponentially decreasing effect, the change of deviance test performed using the GAM model does not appear to yield increased power compared with the linear association test for detecting piecewise constant RR(t) functions. For both tests the power to detect a smaller deviation from proportional hazards is still quite good for jumps at days 183 and 548 but falls sharply

Table 4. Simulations for an exponentially decreasing effect†

Initial effect θ	*Rate of waning* ω	*RR(0)*	*RR(1095)*	*Proportions for the following models:*	
				Schoenfeld residuals: test of linear association	*Poisson GAM: test of change in deviance*
0.1	0.02	0.1	1.0	0.550	0.855
0.1	0.005	0.1	0.996	1.0	1.0
0.1	0.0005	0.1	0.479	1.0	1.0
0.1	0	0.1	0.1	0.055	0.065
0.5	0.02	0.5	1.0	0.205	0.240
0.5	0.005	0.5	0.998	0.860	0.860
0.5	0.0005	0.5	0.7110	0.460	0.280
0.5	0	0.5	0.5	0.060	0.060
1.0	0	1.0	1.0	0.060	0.040

†200 simulations, proportion of samples rejected at $\alpha = 0.05$. RR(0) is the RR at day 0; RR(1095) is that at day 1095. Details are given in the text.

Table 5. Simulations for a piecewise constant effect†

Day of changepoint	*Initial effect* θ_1	*Later effect* θ_2	*Proportions for the following models:*	
			Schoenfeld residuals: test of linear association	*Poisson GAM: test of change in deviance*
183	0.1	0.9	1.0	1.0
183	0.4	0.6	0.780	0.780
548	0.1	0.9	1.0	1.0
548	0.4	0.6	0.785	0.680
913	0.1	0.9	1.0	1.0
913	0.4	0.6	0.115	0.125

†200 simulations, proportion of samples rejected at $\alpha = 0.05$. Details are given in the text.

for day 913. It is possible that the initial sample size and failure rate used resulted in a risk set size this late in the study that was insufficient to detect such a small change in RR(t).

7. Discussion

We have presented two methods for estimating vaccine efficacy using time-to-infection data in the case where the protective effect wanes with time. We have compared the Poisson GAM approach, which estimates a time-varying RR function directly from the likelihood, with the use of smooth scaled Schoenfeld (1982) residuals to obtain nonparametric estimates of RR(t). We applied these methods to the analysis of a 5-year trial of two oral cholera vaccines in Bangladesh and compared this approach with current methods for evaluating vaccine effects when the efficacy appears to wane with time. For this data set, the smooth RR(t) estimates are consistent with earlier results using relative Poisson rates (Clemens *et al.*, 1990; van Loon *et al.*, 1996), while also providing additional insights into the nature of the waning vaccine protection. In general, the B-subunit whole cell vaccine appears to confer better initial protection than does the whole cell vaccine, and neither vaccine appears to offer protection for a period longer than about 3 years.

We have focused on what may be the most likely scenario for a time-dependent vaccine

effect, the case where the efficacy decreases or wanes with time. It is also possible that exposure to natural infection may increase or boost protection, leading to a population RR(t) function that increases with time (Halloran and Struchiner, 1991). The methods presented in this paper can also be used to estimate RR(t) in this case. However, it is important to consider that unmeasured heterogeneity of protective effects of the vaccine can also result in an estimate of RR(t) that appears to increase with time. The most susceptible individuals will probably be infected early in the study, leaving individuals who have a higher level of protection in the risk sets for later time points. Thus the population average RR(t) estimate will increase with time, although protection at the individual level is constant. In this case, frailty models can be used to estimate vaccine efficacy (Halloran *et al.*, 1996; Longini and Halloran, 1996). Understanding the biological processes that result in time-varying efficacy, such as natural boosting of protection or the evolution over time of vaccine-resistant strains, is an important area of future work.

In this paper, for our underlying model, we have assumed that all subjects have the same time origin for follow-up, which can be synchronized with a common base-line hazard for all individuals. In the cholera vaccine trial that we analysed here, the three-dose regime was given over a short period of time in the winter of 1985, before the start of the cholera season. Follow-up for all individuals who received the complete doses began on May 1st, 1985. However, in some field trials, a more complex analysis may be warranted that allows the time origin of the base-line hazard to be different from the time origin of vaccination. This would be especially important for a disease with a highly seasonal risk component, where vaccination occurred over a long period of time through part of the duration of the trial.

Both techniques presented in this paper for evaluating vaccine efficacy in the presence of waning have advantages over current methods that are used by researchers in vaccine field trials and other methods proposed to evaluate waning vaccine protection. They do not require the use of surrogate markers, such as antibody levels, as often the relationship between the levels of serological data and the amount of protection conferred by the vaccine is unknown. The methods presented here also do not involve arbitrarily grouping data to calculate estimates under an imposed piecewise constant model, but they provide a smooth continuous estimate of the RR over time.

However, the estimation of time-varying covariate effects by using smoothed scaled Schoenfeld residuals has several advantages over the approach in the GAM framework for detecting waning vaccine effects. The smoothed estimates from the Schoenfeld technique are obtained on the natural scale of interest and thus can be interpreted more easily. The model also allows for time-dependent covariate values as well as time-varying covariate effects. Thus the approach is very flexible. Also, the Schoenfeld residual approach may be slightly more powerful for detecting vaccine effects that wane slowly over time, which can be a critical effect to diagnose when evaluating vaccine trial data. Finally, the method is easy to implement using standard statistical software (S-PLUS). Because of its flexibility for use as either a diagnostic or an estimation technique, along with its adequate power compared with the GAM procedure for detecting time-dependent effects for many different RR functions, we recommend the smoothed scaled Schoenfeld residuals method as an important tool for evaluating protection against infection and disease in vaccine field trials.

Acknowledgements

This research was partially supported by grants from the National Institute for Allergy and Infectious Diseases R01-AI32042 and T32-AI07442. We thank Rob Tibshirani for bringing

the paper on varying-coefficients models to our attention. We thank John D. Clemens and Malla Rao for providing the cholera vaccine trial data from Bangladesh.

This research was done while LKD was a graduate student at Emory University.

References

Abrahamowicz, M., MacKenzie, T. and Esdaile, J. M. (1996) Time-dependent hazard ratio: modeling and hypothesis testing with application in lupus nephritis. *J. Am. Statist. Ass.*, **91**, 1432–1439.

Chambers, J. M. and Hastie, T. J. (eds) (1992) *Statistical Models in S*. Pacific Grove: Wadsworth and Brooks/Cole.

Clemens, J. D., Sack, D. A., Harris, J. R., van Loon, F., Chakraborty, J., Ahmed, F., Rao, M. R., Khan, M. R., Yunis, M., Huda, N., Stanton, B. F., Kay, B. A., Walter, S., Eeckles, R., Svennerholm, A. and Holmgren, J. (1990) Field trial of oral cholera vaccines in Bangladesh: results from three-year follow-up. *Lancet*, **335**, 270–273.

Cox, D. R. (1972) Regression models and life-tables (with discussion). *J. R. Statist. Soc.* B, **34**, 187–220.

Durham, L. K., Longini, I. M., Halloran, M. E., Clemens, J. D., Nizam, A. and Rao, M. (1998) Estimation of vaccine efficacy in the presence of waning: application to cholera vaccines. *Am. J. Epidem.*, **147**, 948–959.

Gilks, W., Wang, C., Yvonnet, B. and Coursaget, P. (1993) Random-effects models for longitudinal data using Gibbs sampling. *Biometrics*, **49**, 441–454.

Gill, R. and Schumacher, M. (1987) A simple test of the proportional hazards assumption. *Biometrika*, **74**, 289–300.

Grambsch, P. and Therneau, T. M. (1994) Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika*, **81**, 515–526.

Halloran, M. E., Longini, I. M. and Struchiner, C. J. (1996) Estimation and interpretation of vaccine efficacy using frailty mixing models. *Am. J. Epidem.*, **144**, 83–97.

Halloran, M. E. and Struchiner, C. J. (1991) Study designs for dependent happenings. *Epidemiology*, **2**, 331–338.

Hastie, T. J. and Tibshirani, R. (1986) Generalized additive models (with discussion). *Statist. Sci.*, **1**, 297–318.

——— (1990) *Generalized Additive Models*. London: Chapman and Hall.

——— (1993) Varying-coefficient models (with discussion). *J. R. Statist. Soc.* B, **55**, 757–796.

Kooperberg, C., Stone, C. J. and Truong, Y. K. (1995) Hazard regression. *J. Am. Statist. Ass.*, **90**, 78–94.

Longini, Jr, I. M. and Halloran, M. E. (1996) A frailty mixture model for estimating vaccine efficacy. *Appl. Statist.*, **45**, 165–173.

van Loon, F., Clemens, J. D., Chakraborty, J., Rao, M. R., Kay, B. A., Sack, M., Yunis, D. A., Ali, M., Svennerholm, A. and Holmgren, J. (1996) Field trial of inactivated oral cholera vaccines in Bangladesh: results from five years of follow-up. *Vaccine*, **14**, 162–166.

McCullagh, P. and Nelder, J. A. (1989) *Generalized Linear Models*, 2nd edn. London: Chapman and Hall.

McKeague, I. W. and Sasieni, P. D. (1994) A partly parametric additive risk model. *Biometrika*, **81**, 501–514.

Schoenfeld, D. (1982) Partial residuals for the proportional hazards regression model. *Biometrika*, **69**, 239–241.

Statistical Sciences (1995) *S-PLUS Guide to Statistical and Mathematical Analysis, Version 3.3*. Seattle: StatSci.

Whitehead, J. (1980) Fitting Cox's regression model to survival data using GLIM. *Appl. Statist.*, **29**, 268–275.

Zucker, D. M. and Karr, A. F. (1990) Non-parametric survival analysis with time-dependent covariate effects: a penalized likelihood approach. *Ann. Statist.*, **18**, 329–353.