Journal of Biopharmaceutical Statistics, 16: 453–462, 2006

ISSN: 1054-3406 print/1520-5711 online
DOI: 10.1080/10543400600719343

Taylor & Francis
Taylor & Francis Group

EVALUATION OF THE INDIRECT EFFECTS OF A PNEUMOCOCCAL VACCINE IN A COMMUNITY-RANDOMIZED STUDY

L. H. Moulton, K. L. O'Brien, R. Reid, R. Weatherholtz, and M. Santosham

Center for American Indian Health, Department of International Health, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland, USA

G. R. Siber

Wyeth Lederle Vaccines, Pearl River, New York, USA

When a sufficiently high proportion of a population is immunized with a vaccine, reduction in secondary transmission of disease can confer significant protection to unimmunized population members. We propose a straightforward method to estimate the degree of this indirect effect of vaccination in the context of a community-randomized vaccine trial. A conditional logistic regression model that accounts for within-randomization unit correlation over time is described, which models risk of disease as a function of community-level covariates. The approach is applied to an example data set from a pneumococcal conjugate vaccine study, with study arm and immunization levels forming the covariates of interest for the investigation of indirect effects.

Key Words: Community-randomized trial; Indirect effects; Pneumococcal vaccine.

INTRODUCTION

Vaccines protect individuals against many diseases by reducing individual susceptibility in the face of an exposure to pathogens. Bolstering a person's immune system, besides reducing susceptibility, can also have the effect of reducing the probability of transmission of the pathogen to another person. The efficacies of vaccines have been most commonly evaluated in individually randomized trials of participants who form a very small part of the population in which they mix on a daily basis. In these studies, any effect the vaccine may have in preventing secondary transmission from a case to a susceptible person, or through reduction in carriage of the pathogen, is minimal with respect to the overall community-level transmission dynamics. In a large population-based, yet individually randomized trial, up to half the population might receive the vaccine under study. This may lead to an

Received and Accepted March 13, 2006
Address correspondence to L. H. Moulton, Department of International Health, Center for American Indian Health, Room E5519, Johns Hopkins Bloomberg School of Public Health, 615 North Wolfe Street, Baltimore, MD 21205, USA; Fax: 530-325-6867; E-mail: LMoulton@jhsph.edu

appreciable lowering of secondary transmissions in the population, but the effect is not easily measured, as it affects both the immunized and the unimmunized in the same population.

With allocation of vaccine at the community level, i.e., with participating individuals in defined communities all receiving the vaccine and participating individuals in other communities all receiving placebo or control vaccine, the effect of reducing secondary transmission can be estimated. Using the nomenclature of Halloran et al. (1997), the *total effect* is the degree to which disease incidence is lowered in the study participants who received vaccine as compared to those who received, say, a placebo preparation. The total effect is comprised of the *direct effect* of protection of the individual against a direct challenge, and the *indirect effect* of being protected from the challenge in the first place, through reduction of transmission in the community receiving vaccine.

Different approaches have been used to estimate the indirect effects of a vaccine. For example, Ramsay et al. (2003) used surveillance data collected over a long period of time for all of England in before-after types of analyses that depend on assumptions of minimal secular trends. Haber et al. (1991) and Longini et al. (1998) developed transmission models, with varying levels of assumptions, that can be used to estimate indirect effects.

In this communication, we take a "black box" approach to determining the additional degree, beyond the direct effect, of protection conferred on people living in communities where a pneumococcal vaccine was administered. Longini et al. (1998) has proposed study designs in which the vaccine coverage levels in different communities are varied by the investigator. We use the variability that occurred during a community-randomized pneumococcal vaccine study both in time and in space, as communities geared up their recruitment efforts at varying paces, and directly estimate the indirect effects via regression models that have vaccine coverage level as the primary covariate of interest.

BACKGROUND

The vaccine efficacy study had 38 geographically defined randomization units, or communities, which had been formed by combining slightly larger administrative units so as to minimize the degree of social mixing, and hence contamination, between randomization units. Half the units were randomized to the study vaccine, a seven-valent conjugate pneumococcal vaccine (PnCRM7 vaccine), and the other half to an active control, a conjugate meningococcal group C vaccine (MnCC vaccine) that could have no effect on the pneumococcal disease outcomes of the trial. Randomization was stratified by reservation and estimated population of children under the age of 2, with vaccine masked via a system of six coded letters. These vaccine codes corresponded to subunits of the randomization units, so that field personnel were unaware of the randomization unit boundaries. There were 4164 infants enrolled in PnCRM7 communities, and 3926 in MnCC communities between April 1997 and December 1999. At the midpoint of the trial, December 1998, there were about 4,800 and 5,000 total children under the age of 2 years residing in PnCRM7 and MnCC communities, respectively.

All infants less than 2 years of age on two American Indian reservations were eligible to enroll in the study, with doses given at about 2, 4, and 6 months of age,

followed by a booster dose at 12 months; catch-up schedules were used for older infants and children up to 24 months of age. Details of the design may be found in Moulton et al. (2001). The primary invasive disease analyses, given in O'Brien et al. (2003), measured the total effect of the vaccine by comparing incidence rates between enrollees in vaccine and control communities. It was anticipated that immunization would reduce the amount of carriage of the pneumococcus in the PnCRM7 vaccine communities, resulting in reduced secondary transmission of disease.

Here, we concern ourselves with isolating the PnCRM7 vaccine's indirect effects through analysis of the rates of invasive pneumococcal disease among those who were not enrolled in the trial, i.e., those who never received any study vaccine during the trial. In particular, a rate comparison that can establish a direct estimate of indirect effects is incidence rate in vaccine (PnCRM7) units among non-enrolled children, versus the incidence rate in control (MnCC) units among non-enrolled children. Similarly, other rate estimates that could be used to estimate the indirect effects could include incidence rate measures among adults or other age groups in each of the two types of communities.

The above rate comparison can be refined by making a dose-response analysis in which the level of indirect effect is modified by the degree of immunization coverage at the community level, an approach used previously in a synthetic case-cohort design (Moulton et al., 2000). Thus, for example, a non-enrolled infant living in a community allocated to PnCRM7 where 50% of infants have received at least one vaccine dose would be expected to be at lower risk than a non-enrolled infant in a community with 25% PnCRM7 coverage. This would also hold for vaccinated infants, although their indirect benefit would be much smaller than their direct benefit. It should be noted, however, that comparing effects at different levels of coverage can be complicated by the fact that children who enroll in a study, or enroll early, may differ in exposure or suseceptibility to the non- or late-enrolled. As the proportion covered changes during a study, so do the relative characteristics of the enrolled and non-enrolled.

METHODS

Sources of Data

Detailed information on study vaccine administration, including dates of vaccination, among all the infants and children enrolled in the trial was collected during the course of the study. We also were able to obtain overall denominators for each of the 38 geographically defined randomization units from Indian Health Service (IHS) User Population data and birth logs. With these two data sources (i.e., study data and general demographic data) we were able to interpolate to estimate the size of the non-enrolled populations at any day in between April 1997 and October 2000, the period of the trial up to the point of unmasking. The number of non-enrolled children was obtained as the estimated total population minus the enrolled population on any given day in any given randomization unit.

The numerator data, cases of invasive pneumococcal disease in the whole population (i.e., enrolled and non-enrolled children) due to any of the seven serotypes contained in the PnCRM7 vaccine, were those gathered during the course

of the trial up to the day of study unblinding. A standard protocol was used for intensive population and laboratory-based surveillance for all cases of invasive pneumococcal disease, defined as isolation of pneumococcus from a normally sterile body sites in tribal members of all ages. This surveillance was conducted at all IHS facilities, as well as at surrounding and referral hospitals where tribal members might seek care.

Statistical Methods

The underlying model is a non-homogeneous Poisson process in time and space. Let λ_{it} be the rate of invasive disease among non-enrolled children under 2 years of age in randomization unit i on day t. A simple model for λ_{it} may be written as:

$$\lambda_{it} = n_{it} \exp(\alpha_t + \gamma z_i) \tag{1}$$

where n_{it} is the person-days of exposure in the ith unit on the tth day, α_t represents the effect of the tth day (it captures any day-specific secular trends, such as weekend or seasonal effects), and γ is the log rate ratio comparing those in the PnCRM7 units ($z_i = 1$) to those in the MnCC units ($z_i = 0$). If living in a PnCRM7 unit confers protection to non-enrolled children in that unit, i.e., those children receive protective indirect effects of the vaccine, γ will be negative, corresponding to a rate reduction. We fit a somewhat more complicated model that permits comparison across treatment arms of similarly-covered randomization units, and allows us to see at what vaccine coverage level the indirect effects (if any) begin to take effect. We construct dummy variables representing 0–24%, 25–49%, and $\geq$50% enrollment (defined as having received at least one vaccine dose) of the under-2 population on a given day in a given randomization unit, and cross them with the dummy variable for treatment arm, and reparameterize to get five dummy variables representing six conditions. Specifically, the model is:

$$\lambda_{it} = n_{it} \exp\left(\alpha_t + \beta_1 Mnc_{it}^{25-49} + \beta_2 Mnc_{it}^{50+} + \beta_3 Pnc_{it}^{0-24} + \beta_4 Pnc_{it}^{25-49} + \beta_5 Pnc_{it}^{50+}\right) \tag{2}$$

where, for example, Mnc_{it}^{25-49} is unity for the ith unit on the tth day if it is a unit randomized to MnCC vaccine, and if 25–49% of the children under age 2 on that day have received at least one immunization; else, it is zero. Of particular interest are the comparisons across treatment arms within coverage levels. Thus, if β_3, the log rate ratio comparing PnCRM7 units with low (0–24%) coverage to MnCC units with low coverage, is negative, it would indicate indirect effects at low coverage levels. Likewise, the differences $\beta_4 - \beta_1$ and $\beta_5 - \beta_2$ are also of interest, with negative differences indicating the presence of indirect effects at the corresponding levels of vaccine coverage. Specifically, $\exp(\beta_5 - \beta_2)$ represents the rate ratio comparing the incidence in PnCRM7 units with over 50% coverage to incidence in MnCC units with over 50% coverage. Of course, we can control for coverage level by fitting a model with one term for treatment arm and one for a continuous coverage variable; or add a third for their interaction which, if negative, indicates indirect effects increase with vaccine coverage.

The analytic strategy is to eliminate the nuisance parameters α_t from consideration through conditioning on each day of the study. Each day delineates a risk set, similar to a stratum in a case-control study, or the risk set in a Cox regression model. The characteristics of those randomization units that experienced a case on that day, defined as the date of culture, are compared to those who did not have any cases on that day. This produces a probability that the cases' units' characteristics are associated with the cases' units and the characteristics of the units without cases are associated with those units. This is done for each day, and then the probabilities are multiplied together to achieve the conditional likelihood function:

$$\prod_{t=1}^{t=T} \left[n_{it} \exp(\mathbf{x}_{it}\boldsymbol{\beta}) \Big/ \sum_{j \in R(t)} n_{jt} \exp(\mathbf{x}_{jt}\boldsymbol{\beta}) \right]^{\delta_t} \tag{3}$$

where T is the number of days in the study, δ_t is one if there is a case on the tth day and zero otherwise, $R(t)$ is the set of indices of those units "at risk" on day t, and $\mathbf{x}_{jt}$ is the row vector of dummy variables for the jth unit on day t, with $j = i$ representing the unit with a case on that day. This conditional likelihood function is maximized to produce estimates of $\boldsymbol{\beta}$. However, there may be correlation across the individual risk sets, engendered by the fact that the cases are coming from the same randomization units over time. To account for the correlation, we use two approaches. In both, seasonality and secular trends of invasive disease incidence are perfectly accounted for through elimination of each day's rate effect α_t.

In the first approach, we perform a complete-observation bootstrap at the randomization unit level, resampling the entire histories of each unit and then performing the estimation 2,000 times.

Bootstrap standard errors and percentile intervals are used for statistical inference. Computations are carried out in the Stata version 8 routine for conditional logistic regression, using an offset term of $\ln(n_{it})$, and wrapping a bootstrap around it.

The second approach uses a robust variance estimator to account for the correlation, employing a modified version of a SAS macro (Ishikawa and Barlow, 1998) that employs the SAS PHREG procedure. The data are still stratified by day, with the robust variance calculated by squaring each randomization unit's contribution to the score vector and sandwiching it between inverted Fisher information matrices (Royall, 1986). This is accomplished through taking advantage of deletion diagnostics already programmed in SAS PROC PHREG.

Study immunizations did not begin simultaneously throughout the reservations. To ensure that we only used data collected under study conditions, we began each randomization unit's entry into the analysis on the day the first study child in that unit was immunized. Thus a unit's numerator and denominator data were not included in any risk sets until its day of first enrolled (immunized) child.

RESULTS

Figure 1 displays, by randomized study arm and by calendar time, the medians of each trial arm's units' proportions of children under 2 years of age who had received at least one dose of study vaccine. Throughout the study time frame,

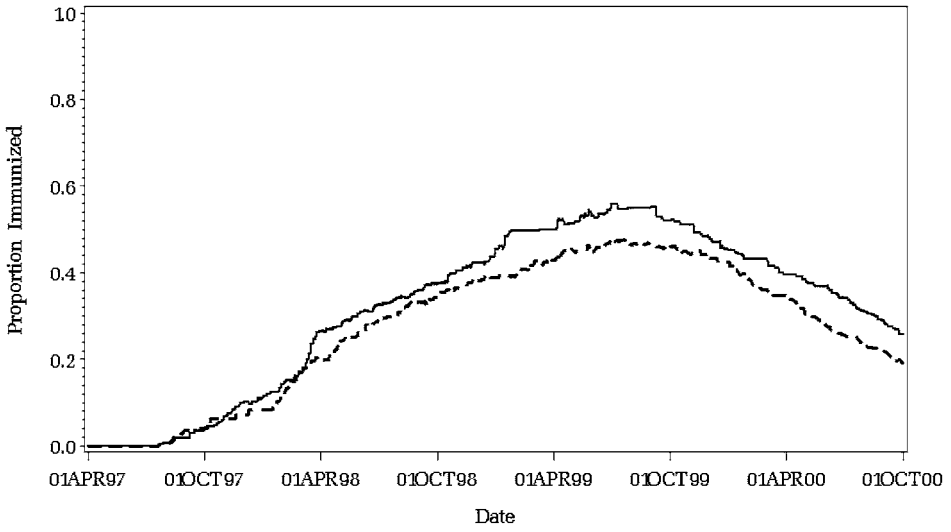

Figure 1 Medians over time of the proportions in randomization units of children under 2 years of age immunized with at least one PnCRM7 (solid line) or MnCC (dashed line) vaccine.

there were slightly higher proportions immunized in the PnCRM7 communities than in the MnCC communities. As of July 1, 1999, near the peak enrollment, the proportions enrolled in PnCRM7 units ranged from 0.06 to 0.71, with a mean of 0.50; in MnCC units, the range was 0.26 to 0.62, with mean 0.46. Enrollment stopped toward the end of the study period, leading to declining proportions.

There were 21 cases of invasive pneumococcal disease due to study vaccine serotypes among nonstudy children living in MnCC randomization units, and 27 cases among those in PnCRM7 units. The distribution of the numbers of cases in units is given in Table 1. Zero cases were observed in 18 of the units, and 2 units had six cases each.

In Table 2, the cases are broken down by the percentage of those under 2 years of age who had received at least one study dose, with rates based on the estimated child-years of exposure in each immunization coverage category. The overall, unadjusted estimate of indirect effects is given by the rate ratio 1.41 (95% CI: 0.77, 2.62), indicating higher incidence in the PnCRM7 communities than in the MnCC communities. In the lowest coverage category (0–24%), the incidence rate

Table 1 Numbers of randomization units in each study arm by number of cases of vaccine serotype invasive pneumococcal disease among non-enrolled children. A total of 48 cases observed in 38 (19 in each arm) randomization units

Vaccine arm	Number of randomization units with given number of cases							
	0 cases	1 cases	2 cases	3 cases	4 cases	5 cases	6 cases	Total
MnCC	8	6	3	1	0	0	1	19 units
PnCRM7	10	1	3	2	2	0	1	19 units

Table 2 Cases and rates of vaccine serotype invasive pneumococcal disease among non-enrolled children by study arm and immunization coverage level

Study Arm/ % coverage	No. cases	Child-years	Rate/ 100,000 child years	Rate ratio (and exact 95% confidence interval) of PnCRM7 to MnCC for each coverage category
MnCC Overall	**21**	**17922**	**117**	
MnCC 0–24%	4	11050	36	
MnCC 25–49%	13	5810	224	
MnCC 50+%	4	1062	377	
PnCRM7 Overall	**27**	**16374**	**165**	**1.41 (0.77, 2.62)**
PnCRM7 0–24%	11	10218	108	2.97 (0.88, 12.81)
PnCRM7 25–49%	8	4163	192	0.86 (0.31, 2.24)
PnCRM7 50+%	8	1992	402	1.07 (0.29, 4.84)

was higher in PnCRM7 communities than in MnCC communities, while the rates were similar across study arms in the higher coverage levels.

Table 3 has the results of fitting the conditional logistic regression model with linear predictor as specified in Eq. (2). The reference status, 0–24% coverage in MnCC units, had the lowest associated rate of invasive pneumococcal disease. The highest rates occurred in units at higher levels of vaccine coverage, with rate ratios of about 7 (exp[1.93], exp[1.96]) for the MnCC and PnCRM7 units with over 50% immunized as compared to the reference category. When units had less than 25% coverage, PnCRM7 units had higher rates than MnCC units (rate ratio = $\exp(1.09) = 3.0$). Comparing PnCRM7 to MnCC units within the coverage levels 25–49% or >50%, however, there was very little difference, indicating no observed indirect effects. As an example of how to make a direct comparison, $\exp(\hat{\beta}_5 - \hat{\beta}_2) = \exp(1.96 - 1.93) = 1.03$. Using the naïve covariance matrix for the parameter estimates, s.e.$(\hat{\beta}_5 - \hat{\beta}_2) = \sqrt{0.694 + 0.557 - 2(0.436)} = 0.616$, yielding a 95% Wald

Table 3 Analysis results from fitting conditional logistic models with five dummy variables to represent six vaccine arm/percentage vaccine coverage combinations. Conditional maximum likelihood estimates (CMLE), standard errors and corresponding 95% confidence intervals, and 95% bootstrap percentile intervals

Dummy variable (arm/% coverage)	CMLE	Naive SE	Bootstrap SE	Robust SE	Naive CI	Bootstrap SE-based CI	Robust CI	Bootstrap percentile interval
MnCC 0–24%	0*							
MnCC 25–49%	1.18	0.64	0.62	0.51	−0.08, 2.43	−0.04, 2.39	0.18, 2.17	0.12, 2.74
MnCC 50+%	1.93	0.83	0.81	0.66	0.30, 3.56	0.35, 3.52	0.64, 3.23	0.46, 4.25
PnCRM7 0–24%	1.09	0.59	0.60	0.49	−0.06, 2.24	−0.08, 2.26	0.14, 2.04	−0.07, 2.58
PnCRM7 25–49%	0.98	0.66	0.75	0.62	−0.32, 2.28	−0.49, 2.45	−0.24, 2.19	−1.05, 2.59
PnCRM7 50+%	1.96	0.75	0.85	0.71	0.50, 3.43	0.29, 3.64	0.56, 3.37	0.68, 4.37

*Reference category: Units that received MnCC vaccine which on a given day had less than 25% of children enrolled in the study. The CMLEs are the log rate ratios comparing incidence in non-enrolled children in the given category to the reference category.

interval for the ratio 1.03 of (0.31, 3.45). These results, which are completely adjusted for any secular or seasonal trends, are consistent with those of the crude rates given in Table 2: 1.07 (0.29, 4.84).

There was little difference between the three estimates of standard error: the naïve estimate, that does not adjust for within-unit correlation, and the robust and bootstrap estimates, that do. Of all the confidence intervals, those produced via the robust standard errors tended to be the shortest.

DISCUSSION

Although with the current analyses we have not found any evidence of indirect effects of PnCRM7, the confidence intervals surrounding the estimates were quite wide, as they were based on a total of only 48 events. Thus, we were not able to rule out the possibility that important indirect effects for children under 2 years of age existed at the levels of immunization coverage seen during the study, which peaked at 50–60%. It may be, however, that higher coverage is required to result in measurable indirect effects—relatively high levels of nasopharyngeal carriage of the pneumococcus in all age groups may mean that modest coverage levels, even with an effective vaccine, do not significantly reduce an infant's chances of having contact with a carrier.

The highest rate ratios were obtained for both MnCC and PnCRM7 units that had the highest proportions immunized, but this may reflect access-to-care issues: perhaps those areas that were closest to the immunization centers (generally in towns) had both higher immunization rates and greatest chance of detection of episodes. This is why it is important to compare rates within similar immunization coverage strata, to avoid the possibility of confounding by access-related variables such as distance and rural/urban mix. Alternatively, data could be collected on variables thought to be related to access and included in the regression analyses. It should be noted that in community randomized trials that are not masked, or only partially masked, during the study period immunization coverage rates could change differentially by trial arm as the study population perceives added advantage (or disadvantage) to receiving the study vaccine. In such situations, interpretation of coverage-related variables will be problematic.

Controlling for immunization levels and study arms, we did not see evidence of overdispersion; indeed, surprisingly, the naïve standard errors were larger than the robust standard errors. This may be due to the relatively small number of cases, and the fact that so many randomization units had only zero or one case. For this example data set, the added complexity of adjusting for correlation turned out to have little effect on the results. We would not have known this, however, if we had not gone through the adjustment exercise. As a demonstration of how additional within-unit correlation could be expressed, we added four additional (fake) cases to one of the MnCC units that already had a few cases and had low ($<25\%$) immunization coverage. The ratios of the robust standard errors to the naïve standard errors, which had been less than unity for the original data, became 1.4, 1.2, and 1.2 for the PnCRM7 terms (0–24%, 35–49%, and 50%+, respectively), and closer to unity for the two MnCC terms.

A full treatment of these data is beyond the scope of this paper, whose focus is on methodology for indirect effects estimation. For example, large decreases

have also been observed among Alaskan Native adults in communities with good pneumococcal vaccine coverage of infants and young children (Hennessy et al., 2005). It could be useful to use our approach to conduct further analyses to model the incidence of invasive pneumococcal disease in older children and adults as a function of vaccine coverage in infants and young children. Also, although receipt of at least one vaccine dose is highly correlated with receipt of further doses, other analyses would look at incidence rate reduction as a function of the proportion of the population who had received a full primary series (an age-specific definition). If pneumococcal carriage in individuals is not much reduced until two, three, or booster doses are given, then such analyses may be more powerful.

Our proposed conditional logistic regression approach, with either bootstrap or robust inference methods to account for correlation, is a straightforward way to approach the estimation of indirect effects in a group-randomized study setting. It may also be employed in phased-implementation one-way crossover studies, also called "stepped wedge" studies, when secular trends may be an issue. Such studies may be more acceptable to a population, especially when a licensed vaccine of known efficacy is being used as a probe to estimate the degree of disease burden in an area. For example, the order of introduction of a vaccine in 24 districts in a country may be randomized, with the supply of vaccine begun in the clinics of another district every month. After 2 years, the whole country has been covered. In any given month, the rates of the disease of interest in the districts that have not yet received vaccine are compared to the rates in those that have received vaccine. The analysis can proceed in the same manner as the ones proposed in this article for the non-enrolled groups, except that the main covariate of interest will be whether a person's district is in a vaccine or non-vaccine status on a given day or month. That would give an intent-to-treat effectiveness estimate; alternatively, the cases' individual status can be used, giving an asvaccinated estimate. Either way, within-district correlation would still need to be handled, as per the approaches we have examined.

ACKNOWLEDGMENT

This work has been supported in part by grants from Wyeth Lederle Vaccines and USAID (HRN-A-00-96-9006). The opinions of the authors are not necessarily those of the Indian Health Service.

REFERENCES

Haber, M., Longini, I. M., Halloran, M. E. (1991). Measures of the effects of vaccination in a randomly mixing population. *Int. J. Epidemiol.* 20:300–310.

Halloran, M. E., Struchiner, C. J., Longini, I. M. Jr. (1997). Study designs for evaluating different efficacy and effectiveness aspects of vaccines. *Am. J. Epidemiol.* 146:789–803.

Hennessy, T. W., Singleton, R. J., Bulkow, L. R., Bruden, D. L., Hurlburt, D. A., Parks, D., Moore, M., Parkinson, A. J., Schuchat, A., Butler, J. C. (2005). Impact of heptaalent pneumococcal conjugate vaccine on invasive disease, antimicrobial resistance and colonization in Alaska natives: progress towards elimination of a health disparity. *Vaccine* 23:5464–5463.

Ishikawa, L., Barlow, W. (1998). ROBPHREG—A SAS macro program using the PHREG procedure to fit case-cohort designs and other proportional hazards models with a robust covariance matrix. V.1.0.2, posted on statlib.

Longini, I. M. Jr., Sagatelian, K., Rida, W. N., Halloran, M. E. (1998). Optimal vaccine trial design when estimating vaccine efficacy for susceptibility and infectiousness from multiple populations. *Stat. Med.* 17:1121–1136.

Moulton, L. H., Chung, S., Croll, J., Reid, R., Weatherholtz, R. C., Santosham, M. (2000). Estimation of the indirect effect of Hib conjugate vaccine in an American Indian population. *Int. J. Epidemiol.* 29:753–756.

Moulton, L. H., O'Brien, K. L., Kohberger, R., Chang, I., Reid, R., Weatherholtz, R., Hackell, J. G., Siber, G. R., Santosham, M. (2001). Design of a group-randomized *Streptococcus pneumoniae* vaccine trial. *Controlled Clinical Trials* 22:438–452.

O'Brien, K. L., Moulton, L. H., Reid, R., Weatherholtz, R., Oski, J., Brown, L., Kumar, G., Parkinson, A., Hu, D., Hackell, J., Chang, I., Kohberger, R., Siber, G., Santosham, M. (2003). A group-randomized efficacy and safety trial of a seven valent conjugate pneumococcal vaccine against invasive pneumococcal disease among high risk American Indian children. *Lancet* 362:355–361.

Ramsay, M. E., Andrews, N. J., Trotter, C. L., Kaczmarski, E. B., Miller, E. (2003). Herd immunity from meningococcal serogroup C conjugate vaccination in England: database analysis. *BMJ* 326:365–366.

Royall, R. M. (1986). Model robust confidence intervals using maximum likelihood estimators. *Int. Stat. Rev.* 54:221–226.