

Research Article

Marginal and Mixed-Effects Models in the Analysis of Human Papillomavirus Natural History Data

Xiaonan Xue¹, Stephen J. Gange², Ye Zhong¹, Robert D. Burk¹, Howard Minkoff³, L. Stewart Massad⁴, D. Heather Watts⁵, Mark H. Kuniholm¹, Kathryn Anastos¹, Alexandra M. Levine⁶, Melissa Fazzari¹, Gypsyamber D'Souza², Michael Plankey⁷, Joel M. Palefsky⁸, and Howard D. Strickler¹

Abstract

Human papillomavirus (HPV) natural history has several characteristics that, at least from a statistical perspective, are not often encountered elsewhere in infectious disease and cancer research. There are, for example, multiple HPV types, and infection by each HPV type may be considered separate events. Although concurrent infections are common, the prevalence, incidence, and duration/persistence of each individual HPV can be separately measured. However, repeated measures involving the same subject tend to be correlated. The probability of detecting any given HPV type, for example, is greater among individuals who are currently positive for at least one other HPV type. Serial testing for HPV over time represents a second form of repeated measures. Statistical inferences that fail to take these correlations into account would be invalid. However, methods that do not use all the data would be inefficient. Marginal and mixed-effects models can address these issues but are not frequently used in HPV research. The current study provides an overview of these methods and then uses HPV data from a cohort of HIV-positive women to illustrate how they may be applied, and compare their results. The findings show the greater efficiency of these models compared with standard logistic regression and Cox models. Because mixed-effects models estimate subject-specific associations, they sometimes gave much higher effect estimates than marginal models, which estimate population-averaged associations. Overall, the results show that marginal and mixed-effects models are efficient for studying HPV natural history, but also highlight the importance of understanding how these models differ. *Cancer Epidemiol Biomarkers Prev*; 19(1); 159–69. ©2010 AACR.

Introduction

Human papillomavirus (HPV), a sexually transmitted virus, is the central etiologic agent in the development of cervical cancer and precancerous cervical lesions (1). There are more than 40 types of HPV, which commonly infect the cervical epithelium as well as other anogenital tissues. At least 13 of these anogenital HPV types are considered oncogenic (2). The incidence of HPV infection is very high, with 3-year cumulative incidence rates reported to be as much as 43% among sexually active college-aged women, and the oncogenic HPV types cause the majority of these infections (3). Most cervical HPV in-

fections, however, including those with oncogenic HPV types, resolve spontaneously within 2 years (4). Only a minority of oncogenic HPV infections persists and leads to clinically significant cervical disease.

Although much has been learned about the natural history of HPV, there is currently a growing focus on type-specific HPV associations rather than data grouped by broad categories (e.g., any oncogenic HPV type), which was common in the past. Randomized clinical trials, for example, are using the incident development of a persistent infection with a vaccine-related HPV type as an intermediate end point (5). Observational studies are also increasingly studying the natural history of HPV on a type-specific basis, whether to compare vaccinated versus unvaccinated women (in nonclinical trial settings) or to investigate other exposures while accounting for possible type-specific differences in effect estimates. The standard logistic regression and Cox model approaches commonly used in HPV research up until now may not always be optimal for such analyses. More efficient statistical methods may need to be adopted.

The choice of statistical method(s) must address several aspects of HPV natural history that are not often encountered elsewhere in infectious disease and cancer research. For example, although coinfection by more than one HPV type is common (3, 6, 7, 8, 9), prior studies suggest that they infect separate cells, causing separate foci

Authors' Affiliations: ¹Albert Einstein College of Medicine, Bronx, New York; ²Johns Hopkins University, Baltimore, Maryland; ³Maimonides Medical Center and SUNY Downstate, Brooklyn, New York; ⁴Southern Illinois University School of Medicine, Springfield, Illinois; ⁵National Institute of Child Health and Human Development, NIH, Bethesda, Maryland; ⁶University of Southern California, Los Angeles, California; ⁷Georgetown University Medical Center, Washington, District of Columbia; and ⁸University of California at San Francisco, San Francisco, California

Corresponding Author: Xiaonan Xue, Department of Epidemiology and Population Health, Albert Einstein College of Medicine, 1300 Morris Park Avenue, Belfer 1308, Bronx, NY 10461. Phone: 718-430-2431; Fax: 718-430-8780. E-mail: xiaonan.xue@einstein.yu.edu

doi: 10.1158/1055-9965.EPI-09-0546

©2010 American Association for Cancer Research.

of infection (8). Therefore, the prevalence, incidence, clearance/persistence, and progression of each HPV type-specific infection can be examined as distinct events/outcomes. On the other hand, because of shared risk factors, these separate HPV infections may be correlated. The probability of detecting any given HPV type, for example, is greater among individuals who are currently positive for at least one other HPV type (3, 6–9). Another type of correlation in HPV natural history data relates to repeated (serial) testing of the same women over time. If these two kinds of correlations are ignored, it can affect P values and confidence intervals (CI), leading to an incorrect statistical inference. On the other hand, statistical models that do not use all the available data (e.g., across all the visits and HPV types examined) may be inefficient.

The current study provides an overview of selected statistical methods that can be used to address these issues. We then illustrate how these methods may be applied to HPV natural history data and compare their results. The example data set is from the Women's Interagency HIV Study (WIHS), a large prospective cohort of HIV-infected and HIV-uninfected women. It is worth noting that these analyses involve more than twice the total person-visits of data available at the time of an earlier report about HPV type-specific infection in HIV-positive women based in the WIHS (8) and represent an update of those earlier findings.

Materials and Methods

Subjects

Data were obtained from the WIHS cohort. Details of this cohort have been previously reported (6, 7, 8). Briefly, between October 1994 and November 1995, 2,058 HIV-seropositive and 568 HIV-seronegative women older than 13 y were enrolled in the WIHS from similar clinical and outreach sources in Brooklyn/Manhattan/Bronx, NY; Chicago, IL; Los Angeles and San Francisco, CA; and Washington, DC. In 2002, an additional 738 HIV-seropositive and 406 HIV-seronegative women were similarly enrolled. On a semiannual basis in this ongoing cohort, subjects undergo a structured interview and a physical examination that includes a gynecologic examination. Following a Pap smear, a cervicovaginal lavage is collected for HPV DNA testing, as are blood samples. At the time of this analysis, there was a median follow-up of 14 visits among all subjects, with 9,475 person-years of observation among HIV-positive women and 2,708 person-years among HIV-negative women.

Laboratory Testing/Data

As described in detail elsewhere (6, 7, 8, 9), HPV testing was conducted at every visit using an MY09/MY11/HMB01 PCR assay. Amplification of a 268-bp cellular β -globin DNA fragment was included in each assay as an internal control. Following amplification, the presence of HPV DNA was assessed using filters hybridized with biotinylated oligonucleotide probes for specific HPV

types, as well as a general probe mixture able to detect most anogenital HPV. For this analysis, oncogenic HPV types were defined as types 16/18/31/33/35/39/45/51/52/56/58/59/66.

Blood specimens were tested for CD4⁺ T-cell count and HIV RNA levels in laboratories participating in the AIDS Clinical Trials Quality Assurance Program. CD4⁺ count was categorized into three strata (i.e., >500, 200–500, and <200 CD4⁺ cells/mm³), and plasma HIV RNA level into four strata (i.e., <4,000; 4,000–20,000; 20,001–100,000; and >100,000 copies/mL), a total of 13 combined CD4⁺/HIV RNA strata that were found to be relevant in prior studies (4, 6), with HIV-negative women as the reference group.

Overview of Statistical Methods

Analysis of HPV Prevalence

Standard Logistic Regression. It is common in HPV studies to use multivariable logistic regression to cross-sectionally analyze the factors associated with “any HPV,” “any oncogenic HPV type,” or an individual HPV type at a given study visit. In such analyses, each woman contributes a single outcome. Notationally, let $Y_i = 1$ represent the detection of HPV16 infection for the i th person at the selected study visit, and 0 otherwise. We model $P_i = P(Y_i = 1)$ by

$$\text{logit } P_i = \beta_0 + \beta Z_i + \gamma W_i \quad (1.1)$$

where Z_i is our exposure variable of interest (e.g., a vector of indicators of 12 separate CD4⁺/HIV RNA strata) and W_i is the vector of other adjusted variables, including age, race, and lifetime of number of male sex for the i th person. The primary parameter of interest is β , defined as the log odds ratio (OR) of HPV16 infection between women in a certain CD4⁺/HIV RNA stratum and HIV-negative women.

Repeated Observations over Time. However, if HPV16 infection were assessed repeatedly over several visits, it would not be appropriate to use model (1.1) to incorporate all the available data: model (1.1) would naively treat repeated observations of each individual over time as if they are independent (i.e., as if they were from different subjects), in essence amplifying the sample size. This can result in underestimation of the variation in the estimate of β and its P value, as well as a CI that is too narrow (10). To avoid these problems, a generalized estimating equation (GEE; ref. 11) can be incorporated in the logistic regression model.

GEE Models. Let Y_{ij} represent the HPV16 infection status (0/1) for the i th person at j th semiannual visit, we model $P_{ij} = P(Y_{ij} = 1)$ by

$$\text{logit } P_{ij} = \beta_0 + \beta Z_{ij} + \gamma W_{ij} \quad (1.2)$$

where Z_{ij} and W_{ij} are defined similarly as in model (1.1). The difference with standard logistic regression is in how the parameters are estimated. Specifically, GEE uses an estimating equation to estimate β in which a particular correlation structure (referred to as the working correlation) is

assumed between repeated measures of Y . Several choices of working correlations are available in most statistical packages, including “independent” (suitable when a small correlation between data at different visits is assumed) and “exchangeable correlation” (suitable when the correlation between all pairs of observations can be assumed to be similar; e.g., visits 1 and 2 have the same correlation as visits 3 and 14), as well as other working correlations. Often, we do not know how the repeated HPV outcomes are correlated with each other, but a useful feature of the GEE approach is that even if the working correlation is misspecified, the estimate of β is still unbiased. Further, because a robust variance is used, the CI for β is correct for large sample size. The main benefit of properly specifying the correlation between repeated observations is that if the working correlation chosen is close to the actual one, the model has greater efficiency (11, 12). In practice, it is recommended that several working correlations be assessed. The one that produces the smallest SE for the estimate of β is generally accepted to be the one that is most consistent to the true correlation. In many cases, though, the results will be similar across different correlation assumptions. Time-dependent covariates can be incorporated.

In model (1.2), β has the same interpretation as that in the logistic regression model [model (1.1)]. Because β describes the average risk difference between two groups/populations, it is often referred to as the “population-averaged effect.”

Multiple HPV types. GEE models, furthermore, can be used to incorporate data across multiple types of HPV. In a single GEE logistic regression model, we can treat each of the several different oncogenic HPV types as separate end points and then estimate the exposure effect on oncogenic HPV as a whole. We note that this common effect can also be examined with any oncogenic HPV as a single binary outcome. However, greater efficiency is obtained for grouped HPV data because each HPV type separately contributes data to the analysis, and the final result is a weighted average across HPV types. This source of efficiency is in addition to that discussed above, of being able to assess repeated visits over time.

We can also assess HPV type-specific relationships with the exposure variables using GEE models. Compared with conducting separate standard logistic regression models, GEE models still obtain greater efficiency because we reduce the number of covariate coefficients by assuming a common effect for “nuisance variables.” For example, only one “lifetime number of sexual partners” coefficient is required if all individual oncogenic HPVs are studied in the same GEE model, whereas the sexual partners variable would require a separate coefficient in each separate standard logistic regression for each HPV type. Further, GEE models allow us to statistically evaluate if the exposure-disease association is common across types.

To illustrate how GEE models can be used to model multiple types of HPV measured across multiple visits, let Y_{ijk} represent the infection of HPV type k for the i th

person at j th semiannual visit. To assess an overall effect estimate across oncogenic HPV types, a type-specific baseline prevalence is assumed for each type of infection but a common OR is assumed across the types, the infection rate of Y_{ijk} , P_{ijk} , is modeled by

$$\text{logit } P_{ijk} = \beta_0 + \theta_k + \beta Z_{ij} + \gamma W_{ij}. \quad (1.3)$$

When sufficient events are available to study each HPV type individually, we can further allow the association with HIV and $CD4^+$ to be different across HPV types. In this case, model (1.3) becomes

$$\text{logit } P_{ijk} = \beta_0 + \theta_k + \beta_k Z_{ij} + \gamma W_{ij}. \quad (1.4)$$

In this model, the association with $CD4^+$ /HIV RNA stratum is modeled by β_k so that it can vary by HPV type k . Model (1.4) also allows us to test if β_k s are the same across HPV types.

Note that two types of correlations exist in the above GEE models: correlation between different types of HPV infection and correlation between the same types of HPV infections at different visits. Standard GEE methods do not easily enable construction of different correlation structures for these two types of correlation (13, 14): only one working correlation structure can usually be used in a single model. Based on prior experience (8), we recommend using either an independent or exchangeable working correlation as an approximation for both correlations.

ORs are a good approximation of relative risk when prevalence rates are low (e.g., <10%; ref. 15), but when prevalence rates are high, ORs overestimate relative risk. Therefore, an additional advantage of the GEE logistic regression approach proposed under model (1.3) is that because it estimates associations “across oncogenic HPV types” (the weighted average of the effects for individual types), the OR is likely to be a good estimate of relative risk (i.e., because the prevalence of oncogenic HPV on a type-specific basis is usually <10% even in high-risk populations). In contrast, any oncogenic HPV as defined under model (1.1) is binary (detection of one or more oncogenic HPV types), which can involve high prevalence in high-risk populations. When a situation arises in which the prevalence of HPV is substantially >10%, direct estimation of prevalence ratios may be preferred to the use of an OR. This can be accomplished in GEE (or mixed-effects models; see below) using log link instead of the logit link used in logistic regression (i.e., ORs), but log link is more likely than logit link to have problems with nonconvergence (16, 17).

Note that there are several links available for studying binary HPV outcomes, such as probit, complementary log-log link, or identity links. Alternative links provide different measures of exposure-disease associations; for example, an identity link provides an estimate of the difference in prevalence (i.e., attributable risk). These other links can sometimes give different results than logit or log links. The decision about which link functions are

correct depends on the underlying biological relationship and usually cannot be determined solely on a statistical basis. However, the logit link (standard logistic regression) is the most common link function used in HPV research.

All the above GEE models can be implemented using PROC GENMOD in SAS. Details to implementing model (1.4) are given in Appendix.

Mixed-Effects Models. A mixed-effects model (18, 19) is an alternative approach for incorporating repeated HPV prevalence data across HPV types and across visits involving the same woman. A mixed-effects model for repeated assessments of HPV16 infection can be expressed as follows:

$$\text{logit } P_{ij} = \mu + \alpha_i + \delta Z_{ij} + \eta W_{ij}, \quad (1.5)$$

which allows each subject to have her own baseline rate of infection modeled by $\mu + \alpha_i$, where μ is the population average baseline rate of infection and α_i represents the i th subject's departure from the population average and is modeled as a random effect. The random effect is unobserved and characterized by a distribution function (typically a normal distribution centered at 0 with variation σ^2), whereas Z and W are observed covariates and referred to as fixed effects. Model (1.5) is therefore called a mixed-effects model. The parameter δ has a different interpretation from β in GEE: it is the log OR of HPV16 infection in a woman in a certain CD4⁺/HIV RNA stratum compared with her risk if she were HIV negative. Because δ describes the change in risk within a person were her risk factor level to change, it is often referred to as the "subject-specific effect" (20, 21). In addition to δ and η , the variation in the random effect σ can also be estimated from the model. The variation in the random effect is a measure of the amount of heterogeneity (i.e., the level of intrasubject correlation) in the population; if $\sigma = 0$, the model reduces to a fixed-effects model. Under a normal assumption, the variation in the random effect can be interpreted as follows: 95% of the subjects in the population will not depart from the population mean μ by $>1.96\sigma$.

Because repeated observations of HPV16 infection share a common α_i , they are correlated. The simple mixed-effects model (1.5) above specifically assumes an exchangeable correlation structure, analogous to that described for GEE (above); that is, it assumes that any pair of repeated observations of HPV within the same woman (e.g., at visits 1 and 5 versus at visits 3 and 14) has the same correlation. However, other correlation structures can be assumed in mixed-effects models, and in general, mixed-effects models attempt to more accurately model the correlations rather than rely on the use of robust variance as in the GEE models. If the correlation is correctly specified, a mixed-effects model is more efficient than a GEE model. However, if it is not correctly specified, the results can be biased (18).

To incorporate repeated observations across multiple HPV types in addition to repeated observations across multiple visits, a mixed-effects model of P_{ijk} (which assumes a common OR across the types) can be defined as

$$\text{logit } P_{ijk} = \mu + \alpha_i + \gamma_k + \delta Z_{ij} + \eta W_{ij} \quad (1.6)$$

where the type-specific baseline prevalence rate represented by γ_k , like α_i , is assumed to be a random effect. This model assumes that correlations among repeated infections of the same type are higher than correlations between different types of HPV, consistent with empirical observations. We can further expand model (1.6) to allow HPV type-specific ORs (details are not provided here). Mixed-effects models for binary outcome are in general more computationally intensive than GEE and can sometimes have difficulty in convergence (i.e., the models cannot be fit and no result is obtained) especially with small data sets.

The mixed-effects model can be implemented using PROC NL MIXED in SAS.

Comparison between GEE and mixed-effects models. Numerically, compared with β in the GEE model, δ in the mixed-effects model is in general larger in magnitude (20), but because the SE for the estimate of δ is also larger, their significance levels (P value) are similar. The interpretations of these estimators also differ, and mixed-effects models may be of greater clinical relevance. For example, if a doctor wants to describe to an individual patient how much her risk of HPV16 infection will change if her CD4⁺ count changes, the subject-specific estimator δ more directly addresses this. The population-averaged estimator given by the GEE is of greater interest from a public health perspective because it describes on average how two groups of women with different CD4⁺ levels will differ in prevalence of HPV infection. For a more detailed discussion on the comparison of β and δ , see Hu et al. (22).

Furthermore, on a practical basis, certain types of analyses that can be conducted using mixed-effects models cannot be conducted using GEE. For example, a mixed-effects model but not GEE can be used to contrast rates of HPV infection before compared with after the initiation of highly active antiretroviral therapy using women as their own comparison group, an instance when it is the within-individual comparison that is of interest.

Power and sample size. Because GEE and mixed-effects models make use of all available data, they generally produce a more efficient estimate and have greater statistical power than standard logistic regression, which involves a single study visit and a single HPV endpoint. It is, however, often difficult to determine an exact power estimate for GEE and mixed-effects models. Statistical simulations are sometimes necessary to estimate power for such studies. Here, we discuss a simplified approach, which may be used as a first-order approximation (23, 24). To obtain a "ball park" estimate of power for correlated binary outcomes, we first assume an average correlation (ρ) among repeated

observations within a subject. If on average each subject contributes J observations and there are a total of n subjects, then the “effective size” (i.e., the equivalent size if each subject only contributes one observation) is $nJ/[1 + (J - 1)\rho]$. For example, if $\rho = 0.3$ and there are $J = 5$ multiple observations per subject and 100 subjects, the power based on these 100 subjects is equivalent to the power based on $100 * 5/[1 + (5 - 1) * 0.3] = 227$ subjects with only one observation per subject. Standard software packages for power/sample size calculation, for example, PASS (25), can then be used to calculate the power for the latter case (see Results).

Analysis of Incident Detection and/or Persistence of HPV

Standard Cox Models. A standard approach for analyzing time to first incident detection of HPV of a particular type, such as HPV16, is to use a Cox proportional hazards model. The Cox model is defined as follows:

$$\lambda(t_i) = \lambda_0(t_i)e^{\beta Z_i(t_i) + \gamma W_i(t_i)} \quad (1.7)$$

where $\lambda_0(t)$ is the baseline hazard function and $\exp(\beta)$ is the hazard ratio (HR) of HPV16 infection between women in a certain CD4⁺/HIV RNA stratum and HIV-negative women while holding all other factors constant. Both the primary covariate Z and adjusting covariates W can be time dependent (i.e., the values of Z and W are allowed to change during the follow-up).

However, model (1.7) cannot simultaneously analyze time to first incident detection of several different types of HPV because it does not address possible correlations between incident HPV infections. Instead, it is often more efficient to use methods that can simultaneously assess multiple types of HPV as separate end points in a single model. There are two main statistical methods appropriate for this purpose, as follows.

Wei-Lin-Weissfeld Method. Wei, Lin, and Weissfeld (26) developed an approach that can be used to simultaneously analyze time to first incident detection of several types of HPV either in the same or different clinical visits, taking into account possible correlations between the types. Although each HPV type is allowed to have its own baseline hazard function, a common (overall) exposure effect can be assumed in the following Wei-Lin-Weissfeld (WLW) model:

$$\lambda_k(t_{ik}) = \lambda_{0k}(t_{ik})e^{\beta Z_i(t_{ik}) + \gamma W_i(t_{ik})} \quad (1.8)$$

where $\lambda_{0k}(t)$ is the baseline hazard function for the k th type of event and $\exp(\beta)$ is the common HR across oncogenic HPV types. Similar to the parameter estimate in GEE, β has a population-averaged interpretation. Model (1.8) is in essence a Cox model stratified by HPV type and it is estimated under an independent working correlation, with a variance estimator that is robust to possible correlations between events.

Alternatively, if there are an adequate number of end points to study each HPV type separately, we can calculate type-specific HRs using the following WLW model:

$$\lambda_k(t_{ik}) = \lambda_{0k}(t_{ik})e^{\beta_k Z_i(t_{ik}) + \gamma W_i(t_{ik})} \quad (1.9)$$

where $\exp(\beta_k)$ is the HR for incident detection of HPV type k for subjects in a certain CD4⁺/HIV RNA stratum relative to HIV negatives.

Model (1.8) has greater power than a standard Cox model because in WLW women contribute multiple time to event end points (one for each HPV type) that are incorporated into the overall effect estimate. Greater efficiency is also obtained in the type-specific analyses using WLW in model (1.9), relative to using separate Cox regression models for each type, because we reduce the number of covariate coefficients by assuming a common effect (i.e., γ) for nuisance variables. Further, model (1.9) allows testing of equality in exposure effects across HPV types.

The WLW model can be implemented using the SAS PHREG procedure, selecting the STRATA option to allow different baseline hazards function for each HPV type. A robust variance is requested by choosing the option COVS(AGGREGATE).

Frailty Models. An alternative to the WLW method is to use frailty models to analyze multiple incident HPV infections (27). A frailty model introduces a frailty term, v_i , to the Cox model as follows:

$$\lambda_k(t_{ik}) = \lambda_0(t_{ik})v_i e^{\theta_k + \delta z_i(t_{ik}) + \gamma W_i(t_{ik})} \quad (1.10)$$

where v_i , treated as a random effect, represents the i th woman's unobserved individual susceptibility to develop an oncogenic HPV infection relative to the average level in the population. The coefficient δ is interpreted as the log of the HR for any incident oncogenic HPV detection in a woman at a certain CD4⁺/HIV RNA stratum compared with her risk if she were HIV negative while holding all other factors constant. Frailty models are a form of mixed-effects model, and the parameter estimate δ has a subject-specific interpretation. Model (1.10) can also be extended to assess HPV type-specific associations (data not shown).

By incorporating e^{θ_k} in the model, frailty models assume that the baseline hazard functions are proportional to one another (i.e., the hazard functions have similar shape). If instead the baseline hazard function for one HPV type increases with age, whereas the baseline hazard function for another decreases with age, this assumption would be violated. WLW, in contrast, is more flexible in that the baseline hazard function for each infection can be entirely different from each other. In addition, as with other mixed-effects models, frailty models (a) explicitly model the correlation between multiple events and, therefore, can be more efficient than WLW (a marginal model) if the correlation is correctly specified but (b) tend to be more computationally intensive than marginal models, and obtaining model convergence can be problematic, especially with small data sets.

Table 1. The weak association between the prevalence of HPV16 and host immune status relative to that for other oncogenic HPV types

HPV types	Ratio of the OR (95% CI)*
HPV16 (reference)	1.00
HPV18	2.05 (0.97-4.34)
HPV31	2.62 (1.02-6.71)
HPV33	2.97 (1.09-8.13)
HPV35	2.66 (1.08-6.55)
HPV39	5.24 (1.39-19.7)
HPV45	3.49 (1.39-8.80)
HPV51	2.56 (1.06-6.19)
HPV52	3.03 (1.35-6.80)
HPV56	2.55 (1.14-5.72)
HPV58	2.70 (1.09-6.68)
HPV59	4.42 (1.72-11.4)
HPV66	5.85 (2.55-13.5)

*The ratio of ORs was estimated using a GEE logistic regression model [model (1.4)] examining the associations of host immune status with the prevalent detection of HPV by type, with HPV16 as the reference type. For convenience, the table shows only the contrast between HIV-positive women with CD4⁺ count <200 cells/mm³ and HIV viral load >100,000 copies/mL versus HIV-negative women. The adjustment variables included age at visit (<30, 30-34, 35-39, 40-44, ≥45 y), race (White, Black, Hispanic, and other), sexual partners in the past 6 mo [none, 1 (married), 1 (single), 2, and >2], and smoking status (none, former smoker, current smoker <10 packs per year, and current smoker ≥10 packs per year).

Frailty models can be implemented using a SAS macro %gamfrail.⁹

Clearance/persistence of HPV. HPV persistence can be measured as the time to clearance of a HPV type following its initial detection, with clearance defined as the first subsequent negative result (or preferably two sequential negative results) for that HPV type (6, 8). The WLW and frailty models described above can be used to assess time to clearance of HPV on a type-specific basis (6, 8).

Power and sample size. Power and sample size can be estimated for WLW and frailty models using the formula described for GEE and mixed-effects models [i.e., $n/[1 + (J - 1)\rho]$, where J is the number of HPV types; ref. 28].

SAS programs for implementing some of the statistical models (above) as well as cautions needed to use these programs are provided on the web.¹⁰

⁹ Available from: <http://www.biostat.mcw.edu/software/SoftMenu.html>.

¹⁰ <http://eph.aecom.yu.edu/xue>

Results

WIHS women have been followed on a semiannual basis. Midinterval (i.e., the midpoint calendar data between two consecutive visits) was used to estimate the time of each incident event. Women were censored for missed visits, and we assumed missing at random. Because our statistical models adjusted for CD4⁺ count and HIV viral load as well as other risk factors, the analyses addressed the main factors that might cause informative censoring in the WIHS cohort. Incident detection of HPV incident detection of type-specific HPV was defined as a positive HPV DNA test result for a specific type in a participant who was negative for that HPV type in all earlier visits (prevalent infections are excluded).

Below we discussed the results for the prevalence data and incidence data.

HPV Prevalence

We assessed the prevalent detection of HPV16 and other oncogenic HPV types (across all visits) and their relationships with host immune status (i.e., CD4/HIV RNA stratum, defined as a time-dependent covariate) using GEE. The prevalence of HPV16 infection varied by host immune status. For example, HPV16 prevalence [model (1.2)] was significantly greater among HIV-positive women with CD4⁺ count <200 cells/mm³ and HIV viral load >100,000 copies/mL compared with HIV-negative women (OR = $e^\beta = 4.66$; 95% CI, 2.60-8.35). However, the same effect estimate measured across oncogenic types [model (1.3)] was much stronger (OR = $e^\beta = 10.82$; 95% CI, 8.32-14.08) than that for HPV16 alone, suggesting that there may be type-specific differences in these associations.

GEE model (1.4) was, therefore, used to assess HPV type-specific differences in the effects of host immune status. HPV16 was the referent HPV type, and the ratio of the type-specific ORs was estimated. As shown in Table 1, each of the other oncogenic HPV types had stronger associations with host immune status than did HPV16. For example, the OR for HPV31 was 2.62-fold ($e^{\hat{\beta}_{31}}/e^{\hat{\beta}_{16}} = e^{\hat{\phi}_{31}}$) greater (95% CI, 1.02-6.71) and the OR for HPV45 was 3.49-fold ($e^{\hat{\beta}_{45}}/e^{\hat{\beta}_{16}} = e^{\hat{\phi}_{45}}$) greater (95% CI, 1.39-8.80) than the OR for HPV16.

We then compared these GEE results with those obtained using a mixed-effects model [model (1.5)]. As expected, the mixed-effects model gave much higher estimates of association, with wider CIs than did GEE (Fig. 1) but similar P values ($P < 0.001$). For example, the OR [i.e., e^δ in model (1.5)] of HPV16 detection associated with having a CD4⁺ count <200 cells/mm³ and HIV viral load >100,000 copies/mL (the most immunosuppressed stratum) relative to being HIV negative was 9.22 (95% CI, 4.84-17.6). This estimate indicates that if an individual HIV-negative woman were to become HIV positive and her immune status dropped to the most immunosuppressed stratum, her risk of having HPV16 infection would increase ~9-fold. In contrast, the OR [i.e.,

$e^{\hat{\beta}}$ in model (1.3)] from the above GEE model would be interpreted as that on average a HIV-positive woman in the most immunosuppressed stratum has a 4-fold greater risk of HPV16 than a HIV-negative woman. A common limitation of mixed-effects models was also revealed. Specifically, mixed-effects models to assess HPV type-specific associations with host immune status failed to converge, consistent with the greater data requirement to run these models compared with GEE.

Statistical power. We use data from Strickler et al. (8) to illustrate the gain in power with GEE compared with standard logistic regression. Based on 535 women with a CD4⁺ count <200 cells/mm³ and 458 women with ≥500 cells/mm³, Fig. 2 shows the power to detect OR = 2.0 for the association of HPV prevalence and CD4⁺ count with (a) only one observation (e.g., one visit with one HPV type), $J = 1$ (i.e., standard logistic regression), versus (b) $J = 5$ or $J = 10$ (based on GEE) per subject under various assumptions about the prevalence rate of HPV infection. In our example, the correlation between repeated observations across HPV types and visits was assumed to be either 0.3 (moderate) or 0.6 (high). As shown in Fig. 2, the power when $J = 5$ is substantially larger than the power without repeated observations ($J = 1$). The power when $J = 10$ is only slightly higher than the power with $J = 5$, indicating that although there is a benefit to statistical power with a certain number of repeated observations, there is a diminishing effect on power as the number of repeats increases. Further, the increase in power is greater when the correlation is lower (e.g., $\rho = 0.3$ versus $\rho = 0.6$). This is as expected because the level of correlation is inversely related to effective sample size. Intuitively, when the correlation between repeated observations is low, having additional repeats should provide substantial new information, whereas having a high correlation implies that the repeated observations are not greatly independent of one another and therefore the repeats provide limited additional information. If the correlation is unknown, a conservative estimate should be used to calculate power/sample size.

Incident HPV Detection

A WLW model of incident oncogenic HPV detection by type [model (1.8)] showed that women with a CD4⁺ count <200 cells/mm³ and HIV RNA >100,000 copies/mL had significantly increased risk of incident oncogenic HPV detection ($HR = e^{\hat{\beta}} = 4.87$; 95% CI, 3.75-6.34) relative to HIV-negative women (Table 2). Further analysis involving measurement of type-specific associations [model (1.9)], using HPV16 as the referent HPV type, showed that most oncogenic HPV types, with the exception of HPV18 and HPV56, had much stronger associations with host immune status than HPV16 (Table 3). Although on an individual basis these 2-fold or greater differences did not reach significance, the power to study incident HPV detection is less than to study HPV prevalence; in our WLW model of incident HPV detection, each person contributed one time-to-event result per

HPV type, whereas in our GEE logistic regression model on HPV prevalence the data set involved repeated observations of each HPV type at each visit. That said, when assessed as a group, the incident detection of all other oncogenic HPV types excluding HPV18 and HPV56 [model (1.9)] had a 3-fold ($HR = e^{\hat{\beta}_{\text{atherone}}} / e^{\hat{\beta}_{16}} = 3.1$) greater (95% CI, 1.4-7.2) association with CD4⁺ count than HPV16.

Lastly, we compared WLW and frailty model [model (1.10)] results. Figure 3 shows the findings of an analysis involving the incident detection of seven common oncogenic HPV (i.e., HPV16/18/31/33/35/45/58) selected because the frailty model had problem converging when less common HPV types were analyzed (most likely because of insufficient occurrences to define their baseline hazards). Interestingly, unlike the comparison of mixed-effects and GEE models, effect estimates obtained using the frailty model were only slightly larger than those obtained using WLW. Frailty models to assess type-specific associations failed to converge.

Discussion

This study shows that marginal and mixed-effects models are appropriate, efficient methods for the analysis of HPV natural history data. Although it has been common to use standard logistic regression for the analysis of the prevalent (cross-sectional) detection of HPV, and standard Cox models for time to incident detection (or

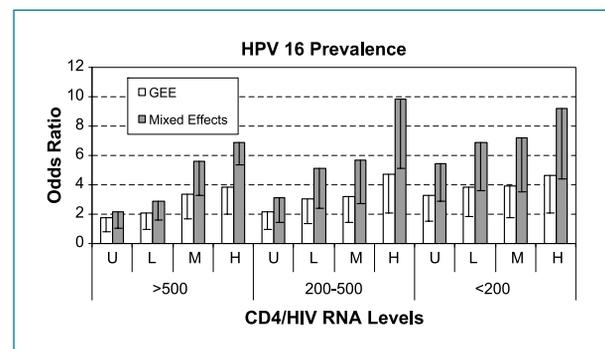


Figure 1. A comparison of the ORs calculated using mixed-effects models [model (1.5)] versus logistic regression models that incorporated a GEE [model (1.2)]. The specific example given is the relationship between HPV16 prevalence and 12 separate CD4⁺/HIV RNA strata, with HIV-negative women used as the reference group. The vertical bars represent the ORs, and the line within each bar represents the lower 95% confidence limit of those ORs. The CD4⁺/HIV RNA strata are defined as follows. Reference group = HIV-negative women. CD4⁺ count: >500, 200 to 500, and <200 cells/mm³. HIV RNA level: undetectable (U), <4,000; low (L), 4,000 to 20,000; moderate (M), 20,000 to 100,000; and high (H), >100,000 copies/mL. Both models adjusted for age (<30, 30-34, 35-39, 40-44, ≥45 y), race (White, Black, Hispanic, and other), sexual partners in the past 6 mo [none, 1 (married), 1 (single), 2, and >2], and smoking status (none, former smoker, current smoker <10 packs per year, and current smoker ≥10 packs per year). The results show that mixed-effects models provide much higher OR estimates than GEE (a form of marginal models) using the same data set.

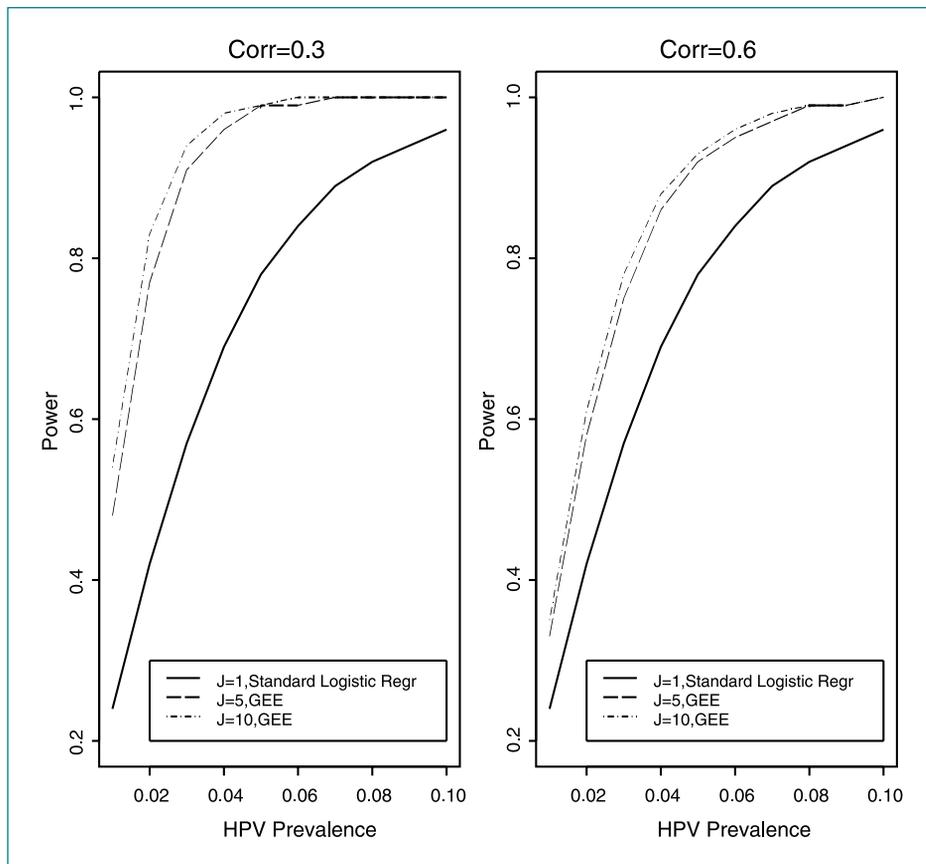


Figure 2. The power to detect an OR of 2.0 for the association between prevalent HPV detection and host immune status, comparing HIV-positive women with CD4⁺ count <200 cells/mm³ (n = 535) versus >500 cells/mm³ (n = 458) under various conditions: (a) using data from only a single observation (J = 1; equivalent to standard logistic regression), (b) using data from 5 or 10 repeated study visits per observations (J = 5, J = 10), and (c) assuming the correlation between repeated observations is moderate (correlation = 0.3) or very high (correlation = 0.6). The figure shows greater statistical power for models that incorporate a GEE, and use data from multiple observations, compared with standard logistic regression models, which can only use data from a single observation.

clearance) of HPV, these common approaches do not make use of all the available data. That is, standard logistic and Cox regression models can only consider a single HPV type, and logistic regression can also only consider a single visit or a single time period. Marginal and mixed-effects models, in contrast, can simultaneously ad-

dress multiple oncogenic HPV types as well as multiple visits over time. As shown in study, they consequently have greater statistical power in the analysis of HPV natural history data. In addition, these models allow comparison of exposure-disease associations between HPV types. Thus, marginal and mixed-effects models could

Table 2. HRs for the incident detection of oncogenic HPV infection as estimated using a WLW model

CD4 T-cell count (cells per mm ³)	HIV RNA (copies per mL)	HR (95% CI)
>500	<4,000	1.61 (1.23-2.10)
	4,001-20,000	2.03 (1.39-2.97)
	20,001-100,000	3.06 (1.96-4.76)
	>100,000	3.40 (1.82-6.32)
200-500	<4,000	2.58 (2.02-3.30)
	4,001-20,000	3.15 (2.37-4.18)
	20,001-100,000	4.02 (3.02-5.35)
	>100,000	4.77 (3.47-6.56)
<200	<4,000	4.55 (3.33-6.21)
	4,001-20,000	5.10 (3.70-7.03)
	20,001-100,000	4.55 (3.38-6.12)
	>100,000	4.87 (3.75-6.34)

NOTE: HIV-negative women were the reference group. The WLW model [model (1.10)] adjusted for the same variables reported in Table 1.

Table 3. The weak association between the incident detection of HPV16 and host immune status relative to that for other oncogenic HPV types

HPV types	Ratio of the HR (95% CI)*
HPV16 (reference)	1.00
HPV18	1.12 (0.37-3.39)
HPV31	2.02 (0.68-6.03)
HPV33	5.28 (1.53-18.2)
HPV35	1.71 (0.43-6.72)
HPV39	3.65 (0.84-15.9)
HPV45	2.05 (0.63-6.70)
HPV51	2.15 (0.72-6.41)
HPV52	5.54 (1.81-17.0)
HPV56	0.97 (0.32-2.93)
HPV58	2.42 (0.82-7.11)
HPV59	3.69 (1.12-12.1)
HPV66	6.29 (2.07-19.1)

*The ratio of HRs was estimated using a WLW model [model (1.9)] examining the associations of host immune status with the incident detection of HPV by type. The adjustment variables included all those shown in Table 1. For convenience, the table shows only the contrast between HIV-positive women with CD4⁺ count <200 cells/mm³ and HIV viral load >100,000 copies/mL versus HIV-negative women. Although most of the oncogenic types did not, on an individual basis, have a significantly different HR than that for HPV16, all other oncogenic types when assessed as a group (excluding HPV18 and HPV56) did have a significantly higher HR than HPV16 ($P = 0.008$). The findings were also significant if we included HPV18 and HPV56 in the analysis.

have many useful applications to HPV research both in randomized clinical trials and in epidemiologic studies.

There are, however, important differences in the estimates obtained from marginal and mixed-effects models. Our study examined the major current forms of these models relevant to the analysis of HPV data. The marginal models were GEE for the analysis of HPV prevalence and WLW for the incident detection or clearance of HPV. Marginal models estimate the average relative difference in risk between two groups of subjects with different levels of risk factor exposure. This is often referred to as the population-averaged effect. The mixed-effects models examined were frailty models for analyzing the incident detection (or clearance) of HPV and mixed-effects models for analyzing HPV prevalence. Mixed-effects models estimate the change in risk that would occur within a person were risk factor exposure to change. This is often referred to as the subject-specific effect.

The effect estimates obtained using mixed-effects models tend to be considerably larger than those

obtained using marginal models but, because the SEs in mixed-effect models are also larger, the significance levels obtained are often similar in both types of models. Interestingly, in our data set, we found that the results of GEE and mixed-effects models had these expected differences, but the WLW and frailty models provided very similar effect estimates. Close agreement between WLW and frailty models in our data may reflect small correlation between the incident detections of different types of HPV, whereas the large differences observed between the GEE and mixed-effects model estimates likely reflect strong correlations between the cross-sectional detection of the same HPV types over serial semiannual visits (i.e., the high frequency of persistent and recurrent infections).

As an important practical matter, mixed-effects models are much more computationally intensive than the marginal models and are more likely to have problems with convergence (i.e., mixed-effects models require more extensive data). For this reason, marginal models may be seen as having an advantage over mixed-effects models in many clinical trials and epidemiologic studies, or whenever data are limited. As mentioned, though, certain types of analysis require mixed-effects models. In general, we recommend a biostatistician be involved when implementing and interpreting results from mixed-effects models.

Although not our main focus, this study also served to further show that the prevalent and incident detection of HPV16 (the HPV type that accounts for half of all cervical cancers) is more weakly associated with host immune status than other oncogenic HPV types, as we and others have previously reported (8). The relative independence of HPV16 infection from host immune status has been interpreted as evidence that HPV16 may be better able to avoid the effects of immune surveillance than other HPV types and that this ability might help partly explain the predominance of HPV16 in cervical disease. The current study represents an update of the prior data involving more than double the person-years of observation available at the time of our earlier report, and approximately a third of the women were new to the analysis.

In summary, although the statistical methods discussed in this study are well established and highly efficient, their application to HPV research has been limited (29). To our knowledge, this is the first formal discussion of these methods in the context of HPV natural history data. However, even if the greater use of these methods were to represent an advance, the analysis of the natural history of HPV will continue to be a challenge. It is hoped, therefore, that this study will not only spur interest in current marginal and mixed-effects models but will additionally interest biostatisticians in the development of new and better methods for the analysis of HPV natural history data.

Appendix

Model (1.4)

$$\text{logit } P_{ijk} = \beta_0 + \theta_k + \beta_k Z_{ij} + \gamma W_{ij}$$

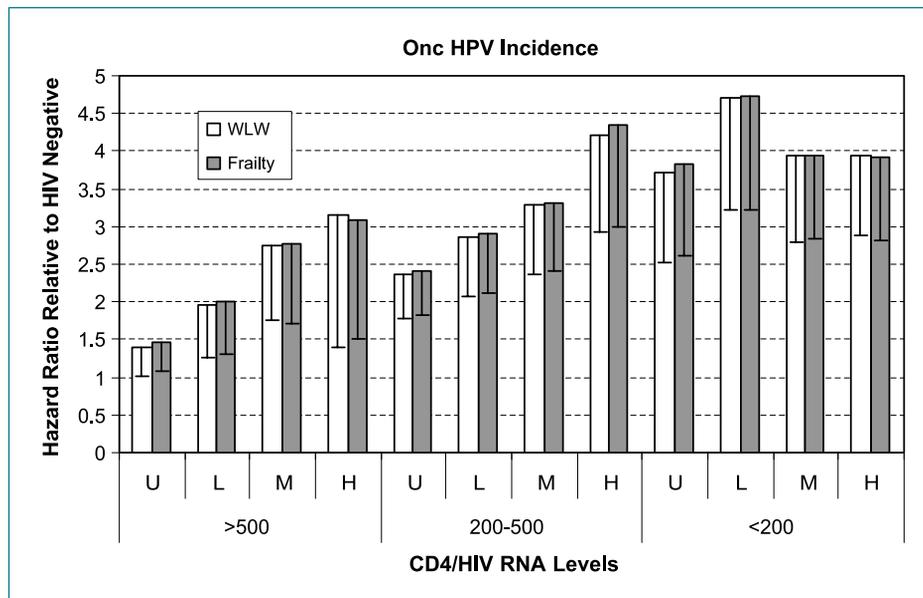


Figure 3. A comparison of the HRs calculated using frailty models [model (1.10)] versus WLW Cox models [model (1.8)]. The specific example given is the relation between the incident detection of seven common oncogenic HPV (i.e., 16/18/31/33/35/45/58) and 12 separate CD4⁺/HIV RNA strata, with HIV-negative women the reference group. The vertical bars represent the HRs, and the line within each bar represents the lower 95% confidence limit of those HRs. The CD4⁺/HIV RNA strata are defined as follows. Reference group = HIV-negative women. CD4⁺ count: >500, 200 to 500, and <200 cells/mm³. HIV RNA level: undetectable (U), <4,000; low (L), 4,000 to 20,000; moderate (M), 20,001 to 100,000; and high (H), >100,000 copies/mL. The adjustment variables are the same as reported in Fig. 1. Both the frailty and WLW models incorporated HPV type-specific baseline hazard functions while assuming a common association with CD4⁺ count/HIV RNA stratum. Overall, the results show that frailty models (a form of mixed-effects models) provide similar HR estimates as WLW (a form of marginal models) using the same HPV data set. As discussed in the text, the close agreement between WLW and frailty models in our data may reflect small correlation between the incident detections of different types of HPV, whereas the large differences observed between the GEE and mixed-effects model estimates (Fig. 1) likely reflect strong correlations between the cross-sectional detection of the same HPV types over serial semiannual visits (i.e., the high frequency of persistent and recurrent infections).

is implemented by including indicator variables ($Type_{ijk}$ for $k = 1, \dots, K$) for each of the other oncogenic HPV types (with HPV16 as a reference) and their interactions with CD4⁺/HIV RNA stratum in the model. The model becomes

$$\text{logit } P_{ijk} = \beta_0 + \theta_1 Type_{ij1} + \dots + \theta_K Type_{ijk} + \beta Z_{ij} + \phi_1 Type_{ij1} Z_{ij} + \dots + \phi_K Type_{ijk} Z_{ij} + \gamma W_{ij}$$

so that β represents the relationship of HPV16 detection with host immune status and ϕ_k represents the difference between HPV type k and HPV16 in relation to host immune status (i.e., the log ratio of the ORs). Type-specific analyses for other models (mixed-effects, WLW, frailty), such as model (1.9), are implemented using the same approach.

Disclosure of Potential Conflicts of Interest

The authors do not have a commercial or other association that might pose a conflict of interest.

Grant Support

HPV DNA testing is funded through R01-CA-085178. All specimens and other data in this study were col-

lected by the WIHS Collaborative Study Group with centers (principal investigators) at New York City/Bronx Consortium (Kathryn Anastos, M.D.); Brooklyn, NY (Howard Minkoff, M.D.); Washington, DC, Metropolitan Consortium (Mary Young, M.D.); The Connie Wofsy Study Consortium of Northern California (Ruth Greenblatt, M.D.); Los Angeles County/Southern California Consortium (Alexandra M. Levine, M.D.); Chicago Consortium (Mardge Cohen, M.D.); and Data Coordinating Center (Stephen J. Gange, Ph.D.). The WIHS is funded by the National Institute of Allergy and Infectious Diseases with supplemental funding from the National Cancer Institute and the National Institute on Drug Abuse (U01-AI-35004, U01-AI-31834, U01-AI-34994, U01-AI-34989, U01-AI-34993, and U01-AI-42590). Funding is also provided by the National Institute of Child Health and Human Development (U01-HD-32632) and the National Center for Research Resources (M01-RR-00071, M01-RR-00079, and M01-RR-00083).

The costs of publication of this article were defrayed in part by the payment of page charges. This article must therefore be hereby marked *advertisement* in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

Received 6/10/09; revised 10/27/09; accepted 11/3/09; published online 1/7/10.

References

1. IARC. IARC monographs on the evaluation of carcinogenic risks to human papillomaviruses, vol. 90. Lyon: IARC; 2007.
2. IARC. IARC monographs on carcinogenicity of human papillomaviruses, vol. 6. Lyon: IARC; 2005.
3. Ho GY, Bierman R, Beardsley L, Chang CJ, Burk RD. Natural history of cervicovaginal papillomavirus infection in young women. *N Engl J Med* 1998;338:423–8.
4. Rodríguez AC, Schiffman M, Herrero R, et al. Rapid clearance of human papillomavirus and implications for clinical focus on persistent infections. *J Natl Cancer Inst* 2008;100:513–7.
5. Jenkins D. A review of cross-protection against oncogenic HPV by an HPV-16/18 AS04-adjuvanted cervical cancer vaccine: importance of virological and clinical endpoints and implications for mass vaccination in cervical cancer prevention. *Gynecol Oncol* 2008;110:S18–25.
6. Strickler HD, Burk RD, Fazzari M, et al. HPV natural history and possible HPV reactivation in HIV-positive women. *J Natl Cancer Inst* 2005;97:577–86.
7. Barkan SE, Melnick SL, Preston-Martin S, et al. The Women's Inter-agency HIV Study. WIHS Collaborative Study Group. *Epidemiology* 1998;9:117–25.
8. Strickler HD, Palefsky JM, Shah KV, et al. Human papillomavirus 16 and immune status in human immunodeficiency virus-seropositive women. *J Natl Cancer Inst* 2003;95:1062–71.
9. Palefsky JM, Minkoff H, Kalish LA, et al. Cervicovaginal human papillomavirus infection in human immunodeficiency virus-1 (HIV)-positive and high-risk HIV-negative women. *J Natl Cancer Inst* 1999;91:226–36.
10. Diggle PJ, Heagerty P, Liang KY, Zeger SL. Analysis of longitudinal data. 2nd ed. Oxford: Oxford University Press; 2002.
11. Liang KY, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika* 1986;73:13–22.
12. Wang YG, Carey VJ. Working correlation misspecification, estimation and covariate design: implications for generalized estimating equation performance. *Biometrika* 2003;90:29–41.
13. Galecki AT. General class of covariance structures for two or more repeated factors in longitudinal data analysis. *Commun Stat Theory Methods* 1994;23:3105–20.
14. Shults J, Whitt MC, Kumanyika S. Analysis of data with multiple sources of correlation in the framework of generalized estimating equations. *Stat Med* 2004;23:3209–26.
15. Zhang J, Yu KF. What's the relative risk? A method of correcting the odds ratio in cohort studies of common outcomes. *JAMA* 1998;280:1690–1.
16. Wacholder S. Binomial regression in GLIM: estimating risk ratios and risk differences. *Am J Epidemiol* 1986;123:174–84.
17. McNutt LA, Wu C, Xue X, Hafner JP. Estimating the relative risk in cohort studies and clinical trials of common outcomes. *Am J Epidemiol* 2003;157:940–3.
18. Goldstein H. Multilevel statistical models. London: Halstead Press; 1995.
19. Fitzmaurice GM, Laird NM, Ware JH. Applied longitudinal analysis. 1st ed.: John Wiley and Sons; 2004.
20. Neuhaus JM, Kalbfleisch JD, Hauck WW. A comparison of cluster-specific and population-averaged approaches for analyzing correlated binary data. *Int Stat Rev* 1991;59:25–35.
21. Pendergast JF, Gange SJ, Newton MA, Lindstrom MJ, Palta M, Fisher MR. A survey of methods for analyzing clustered binary response data. *Int Stat Rev* 1996;64:89–118.
22. Hu FB, Goldberg J, Hedeker D, Flay BR, Pentz MA. Comparison of population-averaged and subject-specific approaches for analyzing repeated binary outcomes. *Am J Epidemiol* 1998;147:694–703.
23. Pan W. Sample size and power calculations with correlated binary data. *Control Clin Trials* 2001;22:211–27.
24. Xue X, Hoover D. Statistical methods in cancer epidemiologic studies. *Methods Mol Biol* 2009;471:239–72.
25. Hintze JL. PASS 2008. NCSS, LLC; Kaysville (UT): Available from www.ncss.com.
26. Wei LJ, Lin DY, Weissfeld L. Regression analysis of multivariate incomplete failure time data by modeling marginal distribution. *J Am Stat Assoc* 1989;84:1065–73.
27. Hougaard P. Analysis of multivariate survival data. New York: Springer; 2000.
28. Xie T, Waksman J. Design and sample size estimation in clinical trials with clustered survival times as the primary endpoint. *Stat Med* 2003; 133:2835–46.
29. Duffy SW, Rohan TE, McLaughlin JR. Design and analysis considerations in a cohort study involving repeated measurements of both exposure and outcome: the association between genital papillomavirus and risk of cervical intraepithelial neoplasia. *Stat Med* 1994; 13:379–90.

Marginal and Mixed-Effects Models in the Analysis of Human Papillomavirus Natural History Data

Xiaonan Xue, Stephen J. Gange, Ye Zhong, et al.

Cancer Epidemiol Biomarkers Prev 2010;19:159-169.

Updated version Access the most recent version of this article at:
<http://cebp.aacrjournals.org/content/19/1/159>

E-mail alerts [Sign up to receive free email-alerts](#) related to this article or journal.

Reprints and Subscriptions To order reprints of this article or to subscribe to the journal, contact the AACR Publications Department at pubs@aacr.org.

Permissions To request permission to re-use all or part of this article, use this link <http://cebp.aacrjournals.org/content/19/1/159>. Click on "Request Permissions" which will take you to the Copyright Clearance Center's (CCC) Rightslink site.