



Published in final edited form as:

J Am Stat Assoc. 2010 June 1; 105(490): 612–620. doi:10.1198/jasa.2010.tm09386.

Test of Association Between Two Ordinal Variables While Adjusting for Covariates

Chun Li [Associate Professor] and Bryan E. Shepherd [Assistant Professor]

Department of Biostatistics, Vanderbilt University, Nashville, TN 37232.

Abstract

We propose a new set of test statistics to examine the association between two ordinal categorical variables X and Y after adjusting for continuous and/or categorical covariates \mathbf{Z} . Our approach first fits multinomial (e.g., proportional odds) models of X and Y , separately, on \mathbf{Z} . For each subject, we then compute the conditional distributions of X and Y given \mathbf{Z} . If there is no relationship between X and Y after adjusting for \mathbf{Z} , then these conditional distributions will be independent, and the observed value of (X, Y) for a subject is expected to follow the product distribution of these conditional distributions. We consider two simple ways of testing the null of conditional independence, both of which treat X and Y equally, in the sense that they do not require specifying an outcome and a predictor variable. The first approach adds these product distributions across all subjects to obtain the expected distribution of (X, Y) under the null and then contrasts it with the observed unconditional distribution of (X, Y) . Our second approach computes “residuals” from the two multinomial models and then tests for correlation between these residuals; we define a new individual-level residual for models with ordinal outcomes. We present methods for computing p -values using either the empirical or asymptotic distributions of our test statistics. Through simulations, we demonstrate that our test statistics perform well in terms of power and Type I error rate when compared to proportional odds models which treat X as either a continuous or categorical predictor. We apply our methods to data from a study of visual impairment in children and to a study of cervical abnormalities in human immunodeficiency virus (HIV)-infected women. Supplemental materials for the article are available online.

Keywords

Categorical data analysis; Concordance–discordance; Regression for ordinal outcome; Residual for ordinal outcome

© 2010 American Statistical Association

chun.li@vanderbilt.edu

bryan.shepherd@vanderbilt.edu

AUTHORS’ NOTE DURING PROOF In our current follow-up work, we found connections between our statistics and traditional ones. When there are no covariates to adjust for, our statistic T_1 is Goodman and Kruskal’s gamma and our statistic T_2 is Spearman’s rank correlation. Proofs will be published with other properties in a separate paper.

SUPPLEMENTAL MATERIALS

R-functions for data analysis: Contains R-functions for analyzing data using proportional odds models for $P(Y|\mathbf{Z})$ and $P(X|\mathbf{Z})$. COBOT stands for conditional ordinal-by-ordinal test. (COBOT-analysis.r)

R-functions for simulations: Contains R code for simulations. (COBOT-simulation.r)

Other simulation results: Provides other results not included in article. (ordinal-otherresults.pdf)

1. INTRODUCTION

Consider the situation where there are two ordinal categorical random variables X and Y , and we want to examine the association between X and Y after adjusting for continuous and/or categorical covariates \mathbf{Z} . This is a common scenario in medical research and social sciences. For example, a recent study collected data on stages of cervical squamous intraepithelial lesions and condom use among human immunodeficiency virus (HIV)-infected women in Zambia (Parham et al. 2006). Cervical dysplastic stage is an ordered categorical variable with four levels according to the commonly used revised Bethesda system: normal, low-grade squamous intraepithelial lesions, high-grade squamous intraepithelial lesions, and squamous cell carcinoma. Condom use was categorized as never, rarely, almost always, and always. Researchers are interested in testing the association between cervix health and condom use after controlling for other variables such as age and CD4 count.

Proportional odds models (also known as ordinal logistic regression) (McCullagh 1980), other cumulative link models (Aitchison and Silvey 1957; Farewell 1982), and continuation ratio models (Läärä and Matthews 1985) are commonly used to examine the association between an ordinal response variable and continuous or categorical predictors. These regression models are useful because they account for the natural ordering of the outcome, but do not treat the outcome as a continuous variable. However, when one of the predictors is ordered categorical, all traditional regression approaches including these for ordered categorical outcomes have to treat the ordinal predictor as either numerical or categorical. The former enforces a linearity assumption and the latter ignores the order information.

When the ordinal predictor is treated as continuous, we assume the effect of moving from level 1 to level 2 is the same as that from level 2 to level 3. Often this assumption is unreasonable. In the previous example, there is little reason to suppose that the effect difference on cervical neoplastic stage between no condom use and rare condom use is the same as that between rare and almost always use. Alternatively, one can assign numbers to the categories or transform the predictor in some manner so that the assigned values reflect a linear relationship with the appropriately transformed expected outcome. The problem with this approach is that such a transformation of the predictor is difficult to choose and may lead to data dredging, ignoring uncertainty in model selection. Splines, a special type of transformation (e.g., Ramsay 1988), have other drawbacks: uncertainty in number and locations of knots, dependence of results on how the categories are coded, nonmonotonic results when a monotonic relationship is expected, and difficulty when there are only three categories.

When the ordinal predictor is treated as discrete, the order information is ignored. When there are many categories, this approach may have low power due to high degrees of freedom. In addition, the effect estimates may be nonmonotonic. Isotonic regression (Barlow et al. 1972) addresses the latter problem by grouping adjacent categories if their relative effect is in reverse of the general trend. However, the categorization is data driven and the results need appropriate adjustment for this source of model selection variability. In addition, the degrees of freedom may still be high, particularly if the original categories already manifest a monotonic relationship.

Another, possibly Bayesian, option can be to perform an analysis using latent variables; the ordinal variables X and Y can be thought of as some categorization of continuous latent variables, say U and V . A test of the association between X and Y conditional on \mathbf{Z} can then be some test of the correlation between U and V conditional on \mathbf{Z} . The primary problem we see with latent variable approaches of this nature is that they require assuming a distribution for the latent variable, which essentially forces one to assign a scale to the ordinal categorical variables.

Some approaches model X and Y jointly conditional on \mathbf{Z} . Examples include log-linear, especially linear-by-linear association, models (Goodman 1979; Agresti 2002) and generalized Cochran–Mantel–Haenszel test (Mantel 1963). However, these methods require assigning numeric values to the ordinal categories, and their implementation also requires grouping continuous or multivariable \mathbf{Z} into strata, generating arbitrary cutoffs on \mathbf{Z} and losing information.

Other approaches for measuring conditional associations of ordinal categorical data have built on classic two-variable statistics such as Kendall’s concordance–discordance statistic tau, Goodman and Kruskal’s gamma (Goodman and Kruskal 1954), and Spearman’s rank-based correlation (Agresti 2002). Kendall’s partial tau (Kendall 1948) and an extension of Kendall’s partial tau to multivariable \mathbf{Z} (Hawkes 1971) are a few examples, although the usefulness of these approaches is questionable given that the expectation of these statistics is generally not zero under the null hypothesis that X and Y are independent conditional on \mathbf{Z} (Goodman 1959; Agresti 1977; Schemper 1991). Other proposed methods involve stratifying data according to \mathbf{Z} and computing weighted averages of stratum-specific measures of association between X and Y (Torgerson 1956; Davis 1967; Agresti 1977). Because these techniques need multiple observations per stratum, these approaches again require collapsing continuous or multivariable \mathbf{Z} into strata.

Our method is motivated by the following observation. If X and Y are continuous variables, we can carry out linear regression $Y = \beta_0 + \beta_1 X + \gamma_1 Z_1 + \dots + \gamma_k Z_k + e$, and test for significance of the coefficient β_1 . Alternatively, we can carry out the following procedure: (a) fit a linear regression of Y on the covariates \mathbf{Z} and obtain the residual Y_{res} for each subject, (b) fit a linear regression of X on \mathbf{Z} , and obtain the residual X_{res} for each subject, and (c) perform a simple linear regression $Y_{res} = \alpha_0 + \alpha_1 X_{res} + e$, and test for significance of the coefficient α_1 . It is well known that the coefficient estimates for β_1 and α_1 are the same (Rao 1973; Mosteller and Tukey 1977). Their corresponding significance levels are similar if the number of subjects is much larger than the number of covariates. Note that for each subject, given the covariate values $\mathbf{Z} = \mathbf{z}$, the linear regression in (a) will effectively yield a distribution of possible realizations of $Y|\mathbf{z}$ for the subject and the regression in (b) will yield a distribution of possible realizations of $X|\mathbf{z}$ for the subject. These two distributions are expected to be independent if there is no association between Y and X after adjusting for \mathbf{Z} .

Similar to the linear regression setting, our approach is to fit Y and X on \mathbf{Z} separately using multinomial models (e.g., proportional odds models), obtain the conditional distributions $Y|\mathbf{Z}$ and $X|\mathbf{Z}$ for each subject, and use this information to construct test statistics. In Sections 2 and 3 we describe our method. In Section 4 we investigate the performance of our method through simulations. In Section 5 we apply our method to two datasets: one on the association between anisometropia and amblyopia and the other on the association between cervical neoplastic stage and condom use. We discuss our results in Section 6.

2. METHOD

Let Y and X be two ordinal variables with s and t categories, respectively. The categories are denoted as $1^Y, 2^Y, \dots, s^Y$, and $1^X, 2^X, \dots, t^X$. Note that the numbers (e.g., 1^Y) are simply symbols for categories and should not be interpreted as quantities. Without loss of generality, suppose the order of categories are $1^Y < 2^Y < \dots < s^Y$ and $1^X < 2^X < \dots < t^X$. We will omit the superscript when it is clear from the context which variable it belongs to. Our goal is to examine the relationship between Y and X after adjusting for k covariates, $\mathbf{Z} = (Z_1, \dots, Z_k)$. We will test the null hypothesis

$$H_0: P(Y, X|Z) = P(Y|Z)P(X|Z),$$

that is, conditional on \mathbf{Z} , Y and X are independent.

As described in the Introduction, we first carry out a multinomial regression analysis of Y on \mathbf{Z} , and a multinomial regression analysis of X on \mathbf{Z} . A commonly used multinomial regression model for ordinal outcomes is the proportional odds model (McCullagh 1980). Other possible models include cumulative link models (Aitchison and Silvey 1957; Farewell 1982) and continuation ratio models for ordinal outcomes (Läärä and Matthews 1985), and multinomial logit models for categorical outcomes (Mantel 1966). One can choose different multinomial models for the two ordinal variables.

For each subject, the regression analyses of Y and X on \mathbf{Z} will provide conditional distributions of Y and X given \mathbf{Z} . If there is no relationship between Y and X after adjusting for the covariates (i.e., under the null), these two conditional distributions will be independent, and the observed value of (Y, X) for the subject is expected to follow the product distribution of these two conditional distributions. We consider two ways of testing the null. The first approach adds these product distributions across all subjects to obtain the expected marginal distribution of (Y, X) under the null of conditional independence and then contrasts this distribution with the observed marginal distribution for (Y, X) (i.e., not conditional on \mathbf{Z}). Our second approach computes “residuals” from the two multinomial regression models and then tests for correlation between these residuals. We define a new individual-level residual for models with ordinal outcomes. We will describe these two approaches separately.

2.1 Observed versus Expected Distributions

In this approach, we compare the observed joint distribution between Y and X with their expected distribution under the null. In general, the joint distribution between Y and X , $P = P(Y, X) = \{\pi_{jl}\}$, can be written as

$$P(Y, X) = \int_{\mathbf{z}} P(Y, X|Z) dP(\mathbf{Z}). \quad (1)$$

Under the null of conditional independence, the joint distribution between Y and X can also be written as

$$P_0(Y, X) = \int_{\mathbf{z}} P(Y|Z)P(X|Z) dP(\mathbf{Z}). \quad (2)$$

Let $P_0 = P_0(Y, X) = \{\pi_{jl}^0\}$. Then under the null, $P = P_0$.

Assume (X_i, Y_i, \mathbf{Z}_i) , $i = 1, \dots, n$, are iid copies from the random vector (X, Y, \mathbf{Z}) . [Note that although we assume (X_i, Y_i, \mathbf{Z}_i) are iid, our methods hold if \mathbf{Z}_i is set to specific levels.] For subject i ($i = 1, \dots, n$), let $p_i^j = P(Y_i = j | \mathbf{Z}_i = \mathbf{z}_i)$ for $j = 1, \dots, s$, and $q_i^l = P(X_i = l | \mathbf{Z}_i = \mathbf{z}_i)$ for $l = 1, \dots, t$. Let $\gamma_i^j = P(Y_i \leq j | \mathbf{Z}_i = \mathbf{z}_i)$; for convenience, let $\gamma_i^0 = 0$ and $\gamma_i^s = 1$. Then $p_i^j = \gamma_i^j - \gamma_i^{j-1}$. The probabilities q_i^l can be similarly written as differences between cumulative probabilities.

To estimate P_0 , we fit separate multinomial models for $P(Y|Z)$ and $P(X|Z)$ to estimate the probabilities p_i^j and q_i^l , denoted as \widehat{p}_i^j and \widehat{q}_i^l . The distribution of \mathbf{Z} is estimated by its empirical distribution. Plugging these estimates into Equation (2), we obtain the estimate $\widehat{P}_0 = \{\widehat{\pi}_{jl}^0\}$, where

$\widehat{\pi}_{jl}^0 = \frac{1}{n} \sum_i \widehat{p}_i^j \widehat{q}_i^l$. Without assuming the null hypothesis, P can be estimated empirically as $\widehat{P} = \{\widehat{\pi}_{jl}\}$, where $\{\widehat{\pi}_{jl}\} = n_{jl}/n$ and n_{jl} is the number of subjects with $Y = j$ and $X = l$.

We then summarize the observed and expected distributions separately. This can be achieved by calculating Goodman and Kruskal's gamma (Goodman and Kruskal 1954), which for a two-

way probability distribution $P = \{\pi_{jl}\}$ is $\Gamma(P) = \frac{C - D}{C + D}$, where $C = \sum_{j_1 < j_2, l_1 < l_2} \pi_{j_1 l_1} \pi_{j_2 l_2}$ and $D = \sum_{j_1 < j_2, l_1 > l_2} \pi_{j_1 l_1} \times \pi_{j_2 l_2}$. Let $\Gamma_1 = \Gamma(\widehat{P})$ and $\Gamma_0 = \Gamma(\widehat{P}_0)$ be the gamma for the observed and expected distributions for (Y, X) . Our test statistic will be $T_1 = \Gamma_1 - \Gamma_0$. When the multinomial models are correct, under the null, $\widehat{P}_0 \rightarrow P$ and $\widehat{P} \rightarrow P$ as $n \rightarrow \infty$, and thus $T_1 = \Gamma(\widehat{P}) - \Gamma(\widehat{P}_0) \rightarrow 0$.

Note that our test statistic accounts for the order information in Y and X , whereas a direct goodness-of-fit approach comparing the observed and expected counts using a statistic in the form of $\sum_{j,l} (\text{Observed} - \text{Expected})^2 / \text{Expected}$, ignores the order information in Y and X .

2.2 Residual-Based Test Statistics

Another way of constructing test statistics is to mimic the residual-based linear regression analysis described in the Introduction. We wish to calculate “residuals” for multinomial models of ordinal outcomes and test if they correlate. For ordered categorical outcome variables, however, we are not aware of a standard way of calculating individual-level residuals. We first define individual-level residuals for ordinal outcome variables and then present test statistics that are based on these residuals.

In linear regression, to calculate the residual for a subject with outcome $Y = y$ and input $\mathbf{Z} = \mathbf{z}$, we first obtain the fitted value of the outcome variable given \mathbf{z} , $\widehat{y} = E(Y|\mathbf{z})$, and then calculate the residual as $y - \widehat{y}$. However, for an ordered categorical outcome variable, we cannot calculate its “fitted” value and need to re-think the derivation of residuals. It should be noted that in addition to providing the fitted value described previously, a linear regression model also gives an estimated distribution of possible outcome values given \mathbf{z} , say $Y_{fit} \sim Y|\mathbf{z}$. The residual can be written as $y - \widehat{y} = y - E(Y_{fit}) = E(y - Y_{fit})$. In other words, we may think of a random variable $y - Y_{fit}$, which is the difference between the observed outcome value y and a random outcome value under the model given \mathbf{z} ; the residual is the expectation of this random variable. This motivates us to define residuals for ordered categorical outcome variables.

Assume a multinomial model for $P(Y|\mathbf{Z})$ with model parameters θ^Y . Let $Y_i = y_i$ be the observed outcome level for subject i . The corresponding distribution of possible outcome levels $Y_{i,fit} \sim Y_i|\mathbf{z}_i$ given covariate \mathbf{z}_i is multinomial with probability distribution p_i^j . Since the outcome variable is ordered categorical, we cannot calculate the difference between y_i and $Y_{i,fit}$, but we can compare them with respect to whether y_i is at a higher or lower level than $Y_{i,fit}$. The

probability for y_i to be higher than $Y_{i,fit}$ is $p_{i,high} = P(y_i > Y_{i,fit}) = \gamma_i^{y_i-1}$, where $\gamma_i^j = P(Y_i \leq j | \mathbf{Z} = \mathbf{z}_i)$ as defined earlier; similarly, the probability for y_i to be lower than $Y_{i,fit}$ is

$p_{i,low} = P(y_i < Y_{i,fit}) = 1 - \gamma_i^{y_i}$; the probability for y_i to tie with $Y_{i,fit}$ is $P(y_i = Y_{i,fit}) = p_i^{y_i}$. We then assign scores to these three types of events: 1, -1, and 0, corresponding to higher, lower, and tie, respectively. The expected score $p_{i,high} - p_{i,low}$ is denoted as $Y_{i,res}$, which is a function of data (Y_i, \mathbf{Z}_i) and model parameters θ^Y . For a fitted model with parameter estimates $\widehat{\theta}^Y$, the residual for subject i is defined as $y_{i,res} = Y_{i,res}(\widehat{\theta}^Y)$. For multinomial models for $P(X|\mathbf{Z})$, the

probabilities $q_{i,high}$ and $q_{i,low}$ can be similarly defined and so are $X_{i,res} = q_{i,high} - q_{i,low}$ and $X_{i,res}$.

When the models for $P(Y|\mathbf{Z})$ and $P(X|\mathbf{Z})$ are correct, we have

$$E(p_{i,high}|\mathbf{Z}_i) = E(\gamma_i^{Y_i-1}|\mathbf{Z}_i) = \sum_{j=1}^s p_i^j \gamma_i^{j-1} = \sum_{j_1 < j_2} p_i^{j_1} p_i^{j_2}, \text{ and similarly}$$

$$E(p_{i,low}|\mathbf{Z}_i) = \sum_{j_1 < j_2} p_i^{j_1} p_i^{j_2}. \text{ Therefore, } E(Y_{i,res}|\mathbf{Z}_i) = E(p_{i,high} - p_{i,low}|\mathbf{Z}_i) = 0 \text{ and } E(Y_{i,res}) = E[E(Y_{i,res}|\mathbf{Z}_i)] = 0. \text{ Similarly, } E(X_{i,res}|\mathbf{Z}_i) = 0 \text{ and } E(X_{i,res}) = 0. \text{ Under the null, since}$$

$$E(Y_{i,res}X_{i,res}) = E[E(Y_{i,res}X_{i,res}|\mathbf{Z}_i)] = E[E(Y_{i,res}|\mathbf{Z}_i)E(X_{i,res}|\mathbf{Z}_i)] = 0,$$

we have $cov(Y_{i,res}, X_{i,res}) = E(Y_{i,res}X_{i,res}) - E(Y_{i,res}) \times E(X_{i,res}) = 0$, and thus the correlation between $Y_{i,res}$ and $X_{i,res}$ is zero.

Once the residuals for models for $P(Y|\mathbf{Z})$ and $P(X|\mathbf{Z})$ are calculated, the sample correlation coefficient between $Y_{i,res}$ and $X_{i,res}$ across all subjects, T_2 , can be used as a test statistic. Under the null, T_2 converges to zero as $n \rightarrow \infty$.

A variation of the previous approach is to compare the observed value of (Y_i, X_i) for subject i with the distribution of possible values of (Y, X) given covariate \mathbf{z}_i . Under the null, subject i 's observed value (Y_i, X_i) should follow the product distribution $\{p_i^j q_i^l\}$. Consider drawing a

random value from the subject's product distribution, (Y'_i, X'_i) , and comparing it with the observed value (Y_i, X_i) . If $Y'_i > Y_i$ and $X'_i > X_i$ (or $Y'_i < Y_i$ and $X'_i < X_i$), that is, both variables are in the same direction, then we record "concordance"; if the two variables are in opposite directions, then we record "discordance"; otherwise we record a tie. Under the null, since both (Y'_i, X'_i) and (Y_i, X_i) follow the same product distribution, the probability of concordance is

equal to the probability of discordance. In fact, there is no need to draw (Y'_i, X'_i) , as under the null, we can derive the probability of concordance as $C_i = p_{i,high}q_{i,high} + p_{i,low}q_{i,low}$, and the probability of discordance as $D_i = p_{i,high}q_{i,low} + p_{i,low}q_{i,high}$. Since $C_i - D_i = (p_{i,high} - p_{i,low})(q_{i,high} - q_{i,low}) = Y_{i,res}X_{i,res}$, we have $E(C_i - D_i) = E(Y_{i,res}X_{i,res}) = 0$. Our third test statistic will, therefore, be the average difference of these probabilities across all subjects,

$$T_3 = \frac{1}{n} \sum_i (\widehat{C}_i - \widehat{D}_i) = \frac{1}{n} \sum_i Y_{i,res}X_{i,res}.$$

3. DISTRIBUTION OF TEST STATISTICS UNDER THE NULL

We present two approaches to obtaining p -values for our test statistics. One is based on empirical distributions generated under the null, the other is based on asymptotic distributions derived from estimating equations.

3.1 Empirical Distribution

Let T be one of the three test statistics described in the last section. To generate an empirical distribution of T , we simulate replicate datasets under the null. To simulate a replicate dataset, we randomly generate one observation from the product distribution $\{p_i^j q_i^l\}_{j,l}$; this is done for $i = 1, \dots, n$. Then we carry out the entire estimating procedure for the replicate dataset to obtain the corresponding statistic, denoted as T^* (i.e., fit separate multinomial models, obtain

predicted probabilities, and calculate test statistic). This is then repeated many, say N^{emp} , times, to get an empirical distribution of T under the null. The two-sided p -value is then computed as either

$$\#(|T^*| \geq |T|) / N^{emp}$$

or

$$2 \times \min \{ \#(T^* \geq T), \#(T^* \leq T) \} / N^{emp}.$$

From our simulations the results are almost the same for these two p -values, so we will present only the first. This procedure is essentially a parametric bootstrap procedure (Efron and Tibshirani 1993).

3.2 Asymptotic Distribution

An alternative approach to computing the p -value is to use the asymptotic distributions of our test statistics under the null hypothesis. In general, we define a vector of parameters θ of length p , whose estimate $\widehat{\theta}$ can be obtained by solving the equation $\sum_{i=1}^n \Psi_i(\theta) = 0$, where $\Psi_i(\theta) = \Psi(Y_i, X_i, \mathbf{Z}_i; \theta)$ is a p -variate function that does not depend on i or n and satisfies $E_{\theta}[\Psi_i(\theta)] = \mathbf{0}$. From M -estimation theory (Stefanski and Boos 2002), if Ψ is suitably smooth, then as $n \rightarrow \infty$,

$$\sqrt{n}(\widehat{\theta} - \theta) \rightarrow_d N(0, V(\theta)),$$

where $V(\theta) = A(\theta)^{-1}B(\theta)[A(\theta)^{-1}]'$, $A(\theta) = E\left[-\frac{\partial}{\partial\theta}\Psi_i(\theta)\right]$ and $B(\theta) = E[\Psi_i(\theta)\Psi_i(\theta)']$. If a test statistic T is smooth function of $\widehat{\theta}$, $T = g(\widehat{\theta})$, then from the delta method,

$$\sqrt{n}[g(\widehat{\theta}) - g(\theta)] \rightarrow_d N(0, \sigma^2),$$

where $\sigma^2 = \left[\frac{\partial}{\partial\theta}g(\theta)\right]V(\theta)\left[\frac{\partial}{\partial\theta}g(\theta)\right]'$. To estimate σ^2 , we estimate $A(\theta)$ as $\frac{1}{n}\sum_i\left[-\frac{\partial}{\partial\theta}\Psi_i(\widehat{\theta})\right]$, $B(\theta)$ as $\frac{1}{n}\sum_i\Psi_i(\widehat{\theta})\Psi_i(\widehat{\theta})'$, and $\frac{\partial}{\partial\theta}g(\theta)$ as $\frac{\partial}{\partial\theta}g(\widehat{\theta})$. If $g(\theta) = 0$ under the null, then the p -value can be computed approximately as $2\Phi\left(\frac{-|T|}{\widehat{\sigma}/\sqrt{n}}\right)$, where Φ is the cumulative distribution function of the standard normal distribution.

We now define θ , $\Psi_i(\theta)$, and $g(\theta)$ for the three test statistics introduced in Section 2. For all three statistics, the parameter vector will have the form $\theta = (\theta^Y, \theta^X, \theta^T)$, where θ^T is different for each statistic. The corresponding estimating function $\Psi_i(\theta)$ will have the form

$$\Phi_i(\theta) = \begin{cases} \frac{\partial}{\partial\theta^Y}l_Y(Y_i, \mathbf{Z}_i; \theta^Y), \\ \frac{\partial}{\partial\theta^X}l_X(X_i, \mathbf{Z}_i; \theta^X), \end{cases}$$

where l_Y and l_X are the log-likelihood functions of the multinomial models that are used to model $P(Y|\mathbf{Z})$ and $P(X|\mathbf{Z})$, with parameters θ^Y and θ^X , respectively. They are score functions and thus $E\left[\frac{\partial}{\partial\theta^Y}l_Y(Y_i, \mathbf{Z}_i; \theta^Y)\right]=0$ and $E\left[\frac{\partial}{\partial\theta^X}l_X(X_i, \mathbf{Z}_i; \theta^X)\right]=0$. The function $\psi(Y_i, X_i, \mathbf{Z}_i; \theta)$ will be different for each statistic.

For $T_1 = \Gamma(\widehat{P}) - \Gamma(\widehat{P}_0)$, we define $\theta^T = (\pi_{11}, \dots, \pi_{1t}, \pi_{21}, \dots, \pi_{2t}, \dots, \pi_{s1}, \dots, \pi_{s,t-1})$, where $\pi_{jl} = P(Y = j, X = l)$ is not conditional on any covariates. Note that θ^T does not contain π_{st} , which is not an independent parameter because $\sum_{j,l} \pi_{jl} = 1$. The corresponding function ψ only depends on θ^T :

$$\psi(Y_i, X_i, \mathbf{Z}_i; \theta^T) = \begin{cases} I_{\{Y_i=1, X_i=1\}} - \pi_{11}, \\ \vdots \\ I_{\{Y_i=s, X_i=t-1\}} - \pi_{s,t-1}, \end{cases}$$

where I_a is the indicator function of event a . By definition, $E[\psi(Y_i, X_i, \mathbf{Z}_i; \theta^T)] = \mathbf{0}$. Let $g(\theta) = \Gamma(P) - \Gamma(P_0)$. Then $g(\theta) = 0$ under the null, and $T_1 = g(\widehat{\theta})$.

For T_2 , the sample correlation between residuals, $\theta^T = (w_1, w_2, w_3, w_4, w_5)$, where $w_1 = E(Y_{i,res})$, $w_2 = E(X_{i,res})$, $w_3 = E(Y_{i,res}X_{i,res})$, $w_4 = E(Y_{i,res}^2)$, and $w_5 = E(X_{i,res}^2)$. The corresponding function ψ is

$$\psi(Y_i, X_i, \mathbf{Z}_i; \theta) = \begin{cases} Y_{i,res} - w_1, \\ X_{i,res} - w_2, \\ Y_{i,res}X_{i,res} - w_3, \\ Y_{i,res}^2 - w_4, \\ X_{i,res}^2 - w_5. \end{cases}$$

By definition, $E[\psi(Y_i, X_i, \mathbf{Z}_i; \theta)] = 0$. Solving the equation $\sum_i \Psi(Y_i, X_i, \mathbf{Z}_i; \theta) = \mathbf{0}$, we have

$$\widehat{w}_1 = \frac{1}{n} \sum_i Y_{i,res}, \widehat{w}_2 = \frac{1}{n} \sum_i X_{i,res}, \widehat{w}_3 = \frac{1}{n} \sum_i Y_{i,res}X_{i,res}, \widehat{w}_4 = \frac{1}{n} \sum_i Y_{i,res}^2, \text{ and } \widehat{w}_5 = \frac{1}{n} \sum_i X_{i,res}^2. \text{ Let } g(\theta) = (w_3 - w_1w_2) / \sqrt{(w_4 - w_1^2)(w_5 - w_2^2)} = \text{cor}(Y_{i,res}, X_{i,res}). \text{ Then } g(\theta) = 0 \text{ under the null, and } T_2 = g(\widehat{\theta}).$$

For $T_3 = \frac{1}{n} \sum_i (\widehat{C}_i - \widehat{D}_i)$, θ^T is a single parameter $\theta^T = E(C_i - D_i)$, which is zero under the null. The corresponding function ψ is

$$\psi(Y_i, X_i, \mathbf{Z}_i; \theta) = (C_i - D_i) - \theta^T.$$

By definition, $E[\psi(Y_i, X_i, \mathbf{Z}_i; \theta)] = 0$. Let $g(\theta) = \theta^T$. Then $g(\theta) = 0$ under the null, and $T_3 = g(\widehat{\theta})$. Infact, the delta method is not needed here as $\widehat{\sigma}^2$ is simply the last element of the diagonal of $\widehat{V}(\theta)$.

4. SIMULATIONS

We carried out simulations to investigate the performance of our method and to compare it with four other approaches: (a) proportional odds model with X coded as a continuous variable,

$$P(Y \leq j|Z) = \left[1 + \exp\left(-\left(\alpha_j^Y + Z\beta^Y + \eta X\right)\right) \right]^{-1},$$

and testing if $\eta = 0$; (b) proportional odds model with X coded as a categorical variable (using dummy variables),

$$P(Y \leq j|Z) = \left[1 + \exp\left(-\left(\alpha_j^Y + Z\beta^Y + \eta_2 I_{(X=2)} + \dots + \eta_t I_{(X=t)}\right)\right) \right]^{-1},$$

and testing if $\eta_2 = \dots = \eta_t = 0$; (c) isotonic proportional odds which X was treated as a categorical variable and adjacent categories were combined to enforce monotonicity if necessary; and (d) proportional odds model with X transformed using restricted cubic splines with three preselected knots. For our method, we computed the p -value using both the empirical and asymptotic distribution approaches.

We investigated the performance of these approaches under multiple simulation scenarios. The first scenario was under the null to investigate Type I error rate. Next, we constructed various alternatives in manners so that different modeling assumptions would be favored. In our first alternative scenario, we generated data such that the effect of the ordinal categories ($1^X, \dots, t^X$) was linear, in the sense that had we simply done an analysis treating X as a continuous variable we would have gotten the correct answer. For the second alternative scenario, we generated data such that the effect of the levels of X was monotonic in a nonlinear fashion. Finally, we considered a simulation scenario with a nonmonotonic relationship between Y and X , a scenario that favors modeling X as categorical.

The specifics of our four data generating scenarios are as follows: We first generated a covariate Z using the standard normal distribution. Then we generated X with five categories using the proportional odds model

$$P(X \leq l|Z) = \left[1 + \exp\left(-\left(\alpha_l^X + \beta^X Z\right)\right) \right]^{-1},$$

with $\alpha^X = (\alpha_1^X, \dots, \alpha_4^X) = (-1, 0, 1, 2)$ and $\beta^X = 1$. The outcome variable Y was generated with four levels using the proportional odds model

$$P(Y \leq j|Z) = \left[1 + \exp\left(-\left(\alpha_j^Y + \beta^Y Z + \eta_1 I_{(X=1)} + \dots + \eta_5 I_{(X=5)}\right)\right) \right]^{-1},$$

with $\alpha^Y = (\alpha_1^Y, \alpha_2^Y, \alpha_3^Y) = (-1, 0, 1)$, $\beta^Y = -0.5$, and $\eta = (\eta_1, \dots, \eta_5)$ specified as

1. $\eta = (0, 0, 0, 0, 0)$ (the null).
2. $\eta = (-0.4, -0.2, 0, 0.2, 0.4)$ (linear effect).
3. $\eta = (-0.30, 0.18, 0.20, 0.22, 0.24)$ (monotonic nonlinear effect).
4. $\eta = (-0.2, 0, 0.2, 0, -0.2)$ (nonmonotonic effect).

For each simulation scenario, we generated 10,000 datasets, each consisting of 500 subjects. To obtain p -values using the empirical distribution of our test statistics, for each dataset we generated 1000 replicates. For all simulation scenarios, the null was rejected if the two-sided p -value was less than 0.05.

Simulation results are summarized in Table 1. Under the null hypothesis, the Type I error rate was at the nominal 5% level for all analysis methods except isotonic regression; to achieve the nominal level using isotonic regression we would need to account for model-selection using some sort of re-sampling procedure. Under the linear scenario, the highest power was obtained under the properly specified model with X treated as a continuous variable. Also as expected, power was lower when X was treated as a continuous variable using splines, or as a categorical variable with or without isotonic regression; this lower power was expected because of the additional degrees of freedom employed by each of these methods. Somewhat surprising was the minimal loss in power seen using any of our new test statistics, ranging around 85%–86%, just below the 87% power seen in the properly specified model.

Under the nonlinear monotonic simulations, our methods had better power than simple analyses treating X as continuous or categorical. The most powerful approach expanded X using restricted cubic splines; power was similar between our method and isotonic regression (although isotonic regression had inflated Type I error rate, so presumably its power would be smaller had we accounted for model-selection uncertainty). Finally, when the true relationship between X and Y conditional on Z was nonmonotonic, our test statistics had poor power. This was expected, as our methods assume monotonicity. Splines and treating X as categorical had higher power under this simulation scenario. Isotonic regression also had inflated power, although given that isotonic regression assumes monotonicity and the U-shaped effect of our simulation scenario, we believe many of these are false positives.

It is worth noting that power was comparable between all of our test statistics (T_1 , T_2 , and T_3) regardless of how we computed the p -value (empirically or asymptotically). Figure 1 shows the distribution of asymptotic p -values under the null for T_1 , T_2 , and T_3 ; all three are highly correlated, the residual-based test-statistics (T_2 and T_3) particularly so.

We also evaluated the Type I error rate of our method at smaller sample sizes ($n = 50, 100$). As expected, empirical-based p -values were more accurate than their asymptotic counterparts, although both approaches yielded Type I error rates close to the nominal 5% level (Table 2).

We also performed a simulation where the model for $P(X|Z)$ was incorrectly specified. Specifically, we generated data as before ($n = 500$) except that X was generated with $\beta^X = 0, 1, 2, 3$ for $l = 1, 2, 3, 4$, respectively. We then performed the analyses assuming a proportional odds model (i.e., constant β^X) for $P(X|Z)$. The Type I error rates were under control (Table 2), and the power under the alternative simulation scenarios was also similar to that reported in Table 1 (data not shown).

Finally, we also performed additional simulations with different numbers of categories for Y and X , and saw results similar to those reported here.

5. EXAMPLES

5.1 Anisometric Amblyopia

Amblyopia is a leading cause of acquired monocular visual impairment. Anisometric amblyopia typically presents later than other types of amblyopia. Because improvement in visual acuity with amblyopia treatment depends on the age at which treatment begins, earlier detection of children with anisometric amblyopia is desired. In a photoscreening program

of preschool children, anisometropia ($\geq 1D$ difference in refractive power between the eyes in any meridian) was detected on 974 preschool children (Leon et al. 2008). Anisometropia magnitude (difference in spherical equivalent $< 1D$, 1 to $< 2D$, 2 to $< 4D$, and $\geq 4D$) was measured along with age and visual acuity, which is used to define amblyopia levels (severe, moderate, mild, and no amblyopia). There is interest in testing the association between anisometropia and amblyopia while adjusting for the effect of age.

Investigators carried out ordinal logistic regression with amblyopia as outcome (Y) and anisometropia (X) and age (Z) as continuous input variables (Leon et al. 2008) and found highly significant association between amblyopia and anisometropia after adjustment for age ($p < 10^{-20}$). Applying the methods described in this article with ordinal logistic regression as models for $P(Y|Z)$ and $P(X|Z)$ led to an even smaller p -value. However, p -value comparisons at this magnitude are essentially meaningless as they lead to the same conclusion. Such significant results are partly due to the large sample size of the study. The comparison between statistical methods will be more meaningful if the sample size was smaller. Thus, we generated 50 datasets of 50 subjects randomly selected from the 974 children, and compared the results of our method and the method used by the investigators (Figure 2). Computing p -values based on the asymptotic distribution of the test statistic, T_1 yielded smaller p -values than the method used in the original study analysis. However, when p -values were calculated based on empirical distributions, they were similar to the p -values of the method used in the original analysis. This is consistent with our simulation results with $n = 50$ where the asymptotic p -values for T_1 had slightly inflated Type I error rates. Results were similar for T_2 . In contrast, p -values based on the asymptotic distribution of T_3 tended to be slightly larger than their empirical counterparts, also consistent with our simulation results for $n = 50$ (Table 2).

The residuals of our method using all subjects are plotted in Figure 3. Note that the four levels of anisometropia are still well separated in the residual plot, while the four levels of amblyopia are overlapping. This indicates that age has a relatively weaker association with anisometropia than with amblyopia.

5.2 Cervical Neoplastic Stage and Condom Use

Cervical specimens for 150 nonpregnant HIV-infected women in Lusaka, Zambia were collected (Parham et al. 2006). Based on cytological analysis, 36 specimens were categorized as normal, 35 as low-grade squamous intraepithelial lesions, 49 as high-grade squamous intraepithelial lesions, and 30 as squamous cell carcinoma. The women also reported condom use: 53 women reported never using condoms, 60 rarely, 13 almost always, and 24 always. Researchers were interested in testing for an association between cervical stage and condom use. Kendall's rank correlation tau was estimated as 0.03 (using R function `cor.test` with argument `method = "kendall"`), which was not statistically different from zero (p -value = 0.64). However, CD4 T-cell count, age, education, and marital status have been linked to cervical abnormalities and may be associated with condom use, so researchers wanted to test for an association between cervical stage and reported condom use after adjusting for these variables. Our method was applied resulting in p -values ranging from 0.76 to 0.91, depending on the test statistic employed, indicating that there was insufficient evidence to conclude that cytological abnormalities were associated with condom use after adjusting for CD4 T-cell count, age, education, and marital status. For the sake of comparison, treating condom use as a continuous variable (1, 2, 3, or 4) and putting it in an ordinal logistic model assuming linearity and expanding using restricted cubic splines with 3 knots yielded p -values of 0.52 and 0.48, respectively. When condom use was treated as a categorical variable, the p -value was 0.66. The residuals of our method are plotted in Figure 3.

6. DISCUSSION

We have developed a new method for testing for associations between two ordinal categorical variables while adjusting for other continuous or categorical variables. In our approach, we separately fit the two ordinal variables on the other covariates using multinomial models such as ordinal logistic regression and then built test statistics based on the predicted probability distributions for these two variables. For our three test statistics, we described approaches to calculate p -values based on either empirical or asymptotic distributions. Our methods are simple to implement and simulations showed our new tests are powerful to detect monotonic associations between two ordinal variables while appropriately adjusting for the effects of other covariates.

In the process of constructing test statistics, we defined a new concept of residual for ordinal outcome variables. This residual can be calculated for any multinomial regression model as long as the outcome variable is ordinal. In our definition of residuals, we assigned scores +1, -1, and 0, reflecting the direction of comparison between the observed and expected outcomes; positive (negative) residuals imply the observed level is higher (lower) than the expected. Our residuals are consistent with concordance–discordance statistics such as Kendall’s tau and Goodman and Kruskal’s gamma, which similarly compare the direction between observations, but make no assumption regarding the magnitude of the distance between ordered categories. Our definition results in one residual per subject, which is, therefore, useful for constructing test statistics. Our residuals may also be useful in other ways (e.g., diagnostics), which we are currently studying.

Other types of residuals, such as deviance and Pearson residuals, were defined for logistic regression (Agresti 2002), and they can be extended to multinomial outcomes. However, deviance residuals ignore the order information, and Pearson residuals result in multiple values for each subject when there are more than two levels. Hence, the utility of these residuals for our purposes is not readily apparent. For proportional odds models, McCullagh (1980) described a different concept of residual. However, this residual is defined for each level of the multinomial outcome variable and is “always positive and thus does not indicate the direction of departure of the observed values from the fitted values.”

Our method has some features which may be undesirable in certain scenarios. First, if one of the two ordinal variables can be designated as the outcome variable, a traditional regression analysis with the other variable as a predictor can model interactive effects between the predictor variable and the covariates. In our method, both ordinal variables are treated equally, avoiding the need to pick one as the outcome, but we therefore assume no interaction effects exist between the predictor and other covariates. Second, our method requires explicit modeling of the relationship between Y and Z and between X and Z . The consistency of our results, therefore, depends on correct specification of these models. It should be recognized that our method is applicable for any (even different) multinomial regression models of Y on Z and X on Z . And our limited simulations suggest results are fairly robust to misspecifications of these models, although this warrants further investigation. Note that a single regression analysis of Y on X and Z does not require modeling the relationship between X and Z , but requires explicit specification of the pattern of effects of X and Z on Y . Third, our approach is designed for testing the relationship between two ordinal variables while adjusting for other covariates. When the primary interest is the relationship between an ordinal outcome and a continuous or categorical variable while adjusting for an ordinal covariate, our methods may not be useful.

Although our presentation focused on hypothesis testing, our test statistics are to some extent interpretable and thus may be used to measure the magnitude of association between Y and X conditional on Z . The test statistic T_1 captures the discrepancy in gamma (the difference in

probability of concordant and discordant random pairs) between the observed distribution for (Y, X) and the expected distribution under the null. The test statistic T_2 is the correlation coefficient between the residuals of models $Y|Z$ and $X|Z$.

When an ordinal variable has only two levels, it is a binary variable. For a binary outcome variable, logistic regression is often used for association testing. For a binary input variable, treating it as continuous or categorical will result in the same model. It is interesting to evaluate how our method performs when Y or X or both are binary variables. If X is binary, our simulations showed that our approach yields similar results to ordinal logistic regression with X treated as a categorical variable (data not shown). If Y is binary, our approach yields results consistent with ordinal logistic regression treating X as the outcome variable and Y as a categorical predictor. Finally, if both Y and X are binary, our approach yields results consistent with logistic regression with X as a dichotomous predictor.

Another possible direction of research is to extend the weighted average of stratum-specific association measures described in the Introduction to continuous or multivariable Z . The weighted average approach currently requires grouping a continuous or multivariable Z into discrete categories, and computing concordance and discordance only for subject pairs falling in the same category. An alternative is to score concordance and discordance for all subject pairs, but to weight the scores according to the similarity in Z between subjects. We are working to evaluate the performance of this approach.

Finally, our methods suggest a potential solution to the general problem in regression when the input variable of interest is ordinal. As we stated in the Introduction, in any regression analysis, ordinal predictor variables have to be treated as continuous or categorical variables, imposing a linearity assumption in the former and ignoring order information in the latter. We defined individual-level residuals for ordinal variables and developed a residual-based method for testing correlation between ordinal Y and X . We believe a similar approach can be developed when Y is another type of variable as long as its individual-level residuals can be calculated. We are also working to evaluate this approach for various variable types for Y .

Regression analysis with ordinal input variables has been difficult to deal with appropriately. Our method will be useful for testing for association between two ordinal variables while adjusting for other covariates. Our simulation and data analysis code is available as supplemental materials and also at <http://biostat.mc.vanderbilt.edu/OrdinalRegression>.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was partially supported by the National Institutes of Health grants R01HG004517 (C. L.) and P30AI54999 (B. E. S.). The authors thank Dr. Sean Donahue for providing data on anisometric amblyopia and Dr. Vikrant Sahasrabudhe for providing data on cervical neoplastic stage and condom use.

REFERENCES

- Agresti A. Considerations in Measuring Partial Association for Ordinal Categorical Data. *Journal of the American Statistical Association* 1977;72:37–45. [613].
- — — . *Categorical Data Analysis*. 2nd ed. Wiley; Hoboken, NJ: 2002. [613,619]
- Aitchison J, Silvey SD. The Generalization of Probit Analysis to the Case of Multiple Responses. *Biometrika* 1957;44:131–140. [612,613].

- Barlow, RE.; Bartholomew, DJ.; Bremner, JM.; Brunk, HD. *Statistical Inference Under Order Restrictions*. Wiley; London: 1972. [612]
- Davis JA. A Partial Coefficient for Goodman and Kruskal's Gamma. *Journal of the American Statistical Association* 1967;62:189–193. [613].
- Efron, B.; Tibshirani, RJ. *An Introduction to the Bootstrap*. Chapman & Hall; London: 1993. [615]
- Farewell VT. A Note on Regression Analysis of Ordinal Data With Variability of Classification. *Biometrika* 1982;69:533–538. [612,613].
- Goodman LA. Partial Tests for Partial Taus. *Biometrika* 1959;46:425–432. [613].
- — — . Simple Models for the Analysis of Association in Cross-Classifications Having Ordered Categories. *Journal of the American Statistical Association* 1979;74:537–552. [613].
- Goodman LA, Kruskal WH. Measures of Association for Cross Classifications. *Journal of the American Statistical Association* 1954;49:732–764. [613,614].
- Hawkes RK. The Multivariate Analysis of Ordinal Measures. *American Journal of Sociology* 1971;76:908–926. [613].
- Kendall, MG. *Rank Correlation Methods*. Hafner; New York: 1948. [613]
- Läärä E, Matthews JNS. The Equivalence of Two Models for Ordinal Data. *Biometrika* 1985;72:206–207. [612,613].
- Leon A, Donahue SP, Morrison DG, Estes RL, Li C. The Age-Dependent Effect of Anisotropia Magnitude on Anisotropic Amblyopia Severity. *Journal of AAPOS* 2008;12:150–156. [617]. [PubMed: 18155938]
- Mantel N. Chi-Square Tests With One Degree of Freedom: Extensions of the Mantel-Haenszel Procedure. *Journal of the American Statistical Association* 1963;58:690–700. [613].
- — — . Models for Complex Contingency Tables and Polychotomous Dosage Response Curve. *Biometrics* 1966;22:83–95. [613]. [PubMed: 5954728]
- McCullagh P. Regression Models for Ordinal Data. *Journal of the Royal Statistical Society, Ser. B* 1980;42:109–142. [612,613,619].
- Mosteller, F.; Tukey, JW. *Data Analysis and Regression*. Addition-Wesley; Reading, MA: 1977. [613]
- Parham GP, Sahasrabudhe VV, Mwanahamuntu MH, Shepherd BE, Hicks ML, Stringer EM, Vermund SH. Prevalence and Predictors of Squamous Intraepithelial Lesions of the Cervix in HIV-Infected Women in Lusaka, Zambia. *Gynecologic Oncology* 2006;103:1017–1022. [612,619]. [PubMed: 16875716]
- Ramsay JO. Monotone Regression Splines in Action. *Statistical Science* 1988;3:425–441. [612].
- Rao, CR. *Linear Statistical Inference and Its Applications*. 2nd ed. Wiley; New York: 1973. [613]
- Schemper M. Non-Parametric Partial Association Revisited. *The Statistician* 1991;40:73–76. [613].
- Stefanski LA, Boos DD. The Calculus of M-Estimation. *The American Statistician* 2002;56:29–38. [615].
- Torgerson WS. A Non-Parametric Test of Correlation Using Rank Orders Within Subgroups. *Psychometrika* 1956;21:145–152. [613].




Figure 1. Comparison of p -values of T_1 , T_2 , and T_3 under the null. Top row are all p -values on log-scale. Bottom row contains p -values between 0 and 0.1. Sample size is 500.




Figure 2. Results of amblyopia data analysis. The x -axis is the p -value based on ordinal logistic regression with anisometropia treated as a continuous variable. The y -axis is the p -value using T_1 . The left and right plots contain p -values based on the asymptotic and empirical distributions of T_1 , respectively. Each point represents results using a random sample of 50 children.

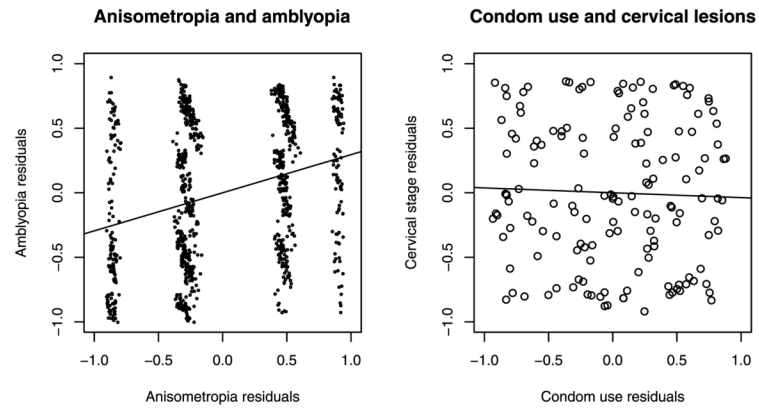


Figure 3. Residuals for the two datasets. The lines are fitted linear regression lines. Some jittering is applied so that the dots are not on top of others. Note the correlation between residuals in the amblyopia example and the lack of correlation in the cervical lesions example.

Table 1

Type I error rate and power (%)

Analysis method	Simulation scenarios			
	Null	Linear	Nonlinear	Nonmonotonic
T_1 empirical	5.0	86.0	57.5	7.4
T_1 asymptotic	4.8	85.4	56.4	7.0
T_2 empirical	5.0	85.8	57.7	6.9
T_2 asymptotic	4.6	85.9	57.8	7.0
T_3 empirical	5.0	85.9	58.0	7.0
T_3 asymptotic	4.9	85.2	57.0	6.6
X linear	4.9	87.4	52.4	5.7
X categorical	5.1	70.3	52.5	28.5
Isotonic	7.1	77.5	57.4	21.9
Splines	4.8	79.9	60.0	34.8

NOTE: $n = 500$.

Table 2

Type I error rate for small sample sizes and violation of model assumptions (%)

Analysis method	$n = 50$	$n = 100$	Misspecified $P(X Z)$
T_1 empirical	4.7	4.4	5.2
T_1 asymptotic	6.0	4.8	4.9
T_2 empirical	5.1	4.7	5.2
T_2 asymptotic	7.0	5.6	5.3
T_3 empirical	5.6	5.0	5.0
T_3 asymptotic	4.0	4.1	5.2
X linear	5.3	4.8	5.1
X categorical	4.2	4.5	5.2
Isotonic	7.1	7.0	7.1
Splines	5.2	4.6	5.2