



Published in final edited form as:

Comput Stat Data Anal. 2008 March 15; 52(7): 3697–3708. doi:10.1016/j.csda.2007.12.012.

Bayesian Analysis of Multivariate Nominal Measures Using Multivariate Multinomial Probit Models

Xiao Zhang, W. John Boscardin, and Thomas R. Belin*

Department of Biostatistics University of California, Los Angeles Los Angeles CA, 90095-1772

Abstract

The multinomial probit model has emerged as a useful framework for modeling nominal categorical data, but extending such models to multivariate measures presents computational challenges. Following a Bayesian paradigm, we use a Markov chain Monte Carlo (MCMC) method to analyze multivariate nominal measures through multivariate multinomial probit models. As with a univariate version of the model, identification of model parameters requires restrictions on the covariance matrix of the latent variables that are introduced to define the probit specification. To sample the covariance matrix with restrictions within the MCMC procedure, we use a parameter-extended Metropolis-Hastings algorithm that incorporates artificial variance parameters to transform the problem into a set of simpler tasks including sampling an unrestricted covariance matrix. The parameter-extended algorithm also allows for flexible prior distributions on covariance matrices. The prior specification in the method described here generalizes earlier approaches to analyzing univariate nominal data, and the multivariate correlation structure in the method described here generalizes the autoregressive structure proposed in previous multiperiod multinomial probit models. Our methodology is illustrated through a simulated example and an application to a cancer-control study aiming to achieve early detection of breast cancer.

Keywords

multinomial multiperiod probit model; MCMC; Metropolis-Hastings; covariance matrix; breast cancer

1 Introduction

The past three decades have seen a great deal of research into discrete choice models, especially in the econometrics and transportation systems literature (Hausman & Wise, 1978; Daganzo, 1980; Bierlaire, 1998). The multinomial logit (MNL) model and the multinomial probit (MNP) model are among the most popular discrete choice models. The MNL model assumes independence of irrelevant alternatives (IIA), which limits its use for data with correlated categorical levels. In contrast, the MNP model provides a framework for representing association between levels of a multinomial outcome. However, MNP models are challenging to analyze due to the computational complexity of algorithms to fit them.

*Zhang is Assistant Professor, Department of Biostatistics, University of Alabama at Birmingham School of Public Health; Boscardin is Associate Professor, and Belin is Professor, Department of Biostatistics, UCLA School of Public Health.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Non-Bayesian methods for analyzing MNP models have focused on using simulated maximum likelihood estimation and the method of simulated moments to avoid direct numerical evaluation of the multidimensional probability integrals involved in maximum likelihood estimation (McFadden, 1989; Geweke, Keane, & Runkle, 1994a; McFadden & Ruud, 1994; Börsch-Supan & Hajivassiliou, 1993). These approaches have been seen to be sensitive to the method of estimating the choice probabilities (McCulloch & Rossi, 1994).

Statistical computing innovations including Markov chain Monte Carlo (MCMC) estimation (Tanner & Wong, 1987; Gelfand & Smith, 1990; Gilks et al., 1996; Gamerman, 1998; Carlin & Louis, 2000; M.-H. Chen et al., 2000; Liu, 2001; Congdon, 2001; Gelman et al., 2003; Robert & Casella, 2004) have allowed for the development of Bayesian methods for complex models. Bayesian approaches for fitting MNP models have included Albert and Chib (1993), McCulloch and Rossi (1994), Nobile (1998), Chib, Greenberg, and Chen (1998), Nobile (2000), McCulloch, Polson, and Rossi (2000), Z. Chen and Kuo (2002), and Imai and van Dyk (2005a).

Multivariate discrete choice models generalize these models to permit analysis of multivariate categorical data. Kim, Allenby, and Rossi (2002) proposed an additive random utility model for modeling consumer demand for more than one variety (alternative). Bhat (2005, 2006) developed a new random utility model. These authors employed simulated maximum likelihood estimation. Specialized multivariate discrete choice models have also been considered. The mixed multinomial logit (MMNL) model (Hensher & Greene, 2003) extends the MNL model to relax the IIA assumption and to allow analysis of multivariate nominal correlated measures. Maximum simulated likelihood estimation, the method of simulated moments, and Bayesian methods have been used for inference in the MMNL model (McFadden & Train, 2000; Train, 2001; Sivakumar, Bhat, & Ökten, 2005; Train & Sonnier, 2005).

Extensions of MNP models to nominal data at multiple time points have also been proposed (McCulloch & Rossi, 1994; Geweke, Keane, & Runkle, 1994b, 1997; Z. Chen & Kuo, 2002; Ziegler, 2002; Rendtel & Kaltenborn, 2004). Computational issues in these multiperiod multinomial probit models have confined the covariance matrix of the latent variables to have either a low-order factor structure (Ziegler, 2002), a first-order autoregressive structure (McCulloch & Rossi, 1994; Geweke et al., 1994b, 1997; Rendtel & Kaltenborn, 2004), or a covariance matrix with a random scale (Z. Chen & Kuo, 2002).

In the case where no covariates are used in the model, the data can be represented as a multiway contingency table. A traditional approach to analyzing such data is log-linear modeling (Bishop, Fienberg, & Holland, 1975). For multi-way contingency table data, the model we propose here has some similarities to certain types of log-linear models; distinctions relate to relaxation of assumptions similar to IIA in MNL models.

An alternative approach for analyzing non-normal repeated measures, such as repeated categorical data and repeated ordinal data, is to use the generalized estimating equation (GEE) method (Liang & Zeger, 1986; Zeger & Liang, 1986; Zeger, 1987; Liang, Zeger, & Qaqish, 1992). However, the GEE method is not indicated when inference about the correlation parameters, in addition to inference for regression parameters, is of interest.

Golob and Regan (2002) extend the MNP model for univariate nominal measure to the multivariate MNP (MVMNP) model to analyze multivariate nominal data, using the generalized least-squares approach proposed by Muthen (1983) for structural equation models. To solve the parameter identification, they standardized the covariance matrix, i.e., used a correlation matrix instead of a covariance matrix. However, this technique forces the magnitudes of each latent variables to be equal and this may not be appropriate without any knowledge of the latent variables.

In this paper, we consider the MVMNP model from a Bayesian perspective. We develop an MCMC algorithm allowing general specification for covariance matrices of the latent variables (i.e., the utilities in choice models). Our key statistical computing innovation is to use the parameter-extended Metropolis-Hastings (PX-MH) algorithm to sample the covariance matrix with restrictions on the diagonal elements. The PX-MH algorithm was proposed by Zhang, Boscardin, and Belin (2006) to sample a correlation matrix in the setting of the multivariate probit model. It was used by Boscardin and Zhang (2004) in a mixture of continuous and ordinal repeated measures. The Bayesian methods we propose both avoid parameter identification problems and allow flexible prior distributions on the covariance matrix of the latent variables.

Our paper proceeds as follows. Section 2 describes the MNP and MVMNP models, including some discussion about the model identification problem. Section 3 presents the MCMC sampling algorithm for the MVMNP model and describes the PX-MH algorithm for sampling the restricted covariance matrix. To illustrate our method, we use a simulated example in Section 4 and an example involving incomplete contingency table data from a cancer-control study in Section 5. Finally, we draw some conclusions and discuss other issues relative to MVMNP models in Section 6.

2 Multivariate multinomial probit (MVMNP) models

To describe the MVMNP model, we start with a description and a review of estimation methods for the MNP model.

2.1 MNP models for univariate nominal measures

Letting $i = 1, 2, \dots, n$ index subjects and $j = 1, 2, \dots, p$ index levels of a multinomial outcome having p levels, we let $y_{ij} = 1$ if subject i has outcome j and $y_{ij} = 0$ otherwise, with $y_i = (y_{i1}, \dots, y_{ip})$ representing a multinomial $1 \times p$ vector. More compactly, we define $d = (d_1, \dots, d_n)^T$, where d_i contains the index of the chosen alternative, i.e., $d_i = j$ if $y_{ij} = 1$. Following the notation in economics settings where utilities underlie choices, the MNP model assumes that there is a latent $1 \times p$ vector $u_i = (u_{i1}, \dots, u_{ip})$ underlying each multinomial vector y_i , such that the multinomial outcome is determined by the maximum u_{ij} , as would happen if the subject chooses the alternative with maximum utility score. That is

$$d_i = j \iff u_{ij} \geq \max_{1 \leq l \leq p} u_{il} \quad (1)$$

Although a simpler description of the latent-variable mean structure is possible, we proceed with a formulation that allows for covariates. In this framework, the MNP model further assumes that the vector u_i follows a multivariate normal distribution with mean equal to $A_i\beta$ and covariance matrix equal to V , where A_i is a $p \times k$ covariate matrix for subject i and β is a $k \times 1$ regression parameter vector. With this notation,

$$u_i = A_i\beta + \delta_i \quad (2)$$

where $\delta_i \sim N(0, V)$. The elements of A_i might reflect subject-specific covariates, in which case all of the elements in a row of A_i would be the same, or outcome-level-specific covariates (the classic case being the cost associated with each of the choices), in which case row elements would differ in general.

There are two identification problems in the above MNP model specification. Based on equation (1), we can see that the model will not be changed if a constant is added to both sides of equation (2). This first identification problem is known as *additive redundancy*. This

problem is usually solved by subtracting the p -th row of equation (2) from the first $(p - 1)$ rows. The model becomes

$$z_i = X_i \beta + \epsilon_i \tag{3}$$

where $\epsilon_i \sim N(0, \Sigma)$ independently, $z_{ij} = u_{ij} - u_{ip}$, $X_{ij} = A_{ij} - A_{ip}$, $\epsilon_{ij} = \mu_{ij} - \mu_{ip}$ and $\Sigma = [I_{p-1}, -1_{p-1}] V [I_{p-1}, -1_{p-1}]^T$, with I_s denoting the $s \times s$ identity matrix and 1_s a vector of length s comprised of 1's. Then the model can be described as:

$$d_i = \begin{cases} 0 & \text{if } \max_{1 \leq l \leq p-1} z_{il} < 0 \\ j & \text{if } \max_{1 \leq l \leq p-1} z_{il} = z_{ij} > 0 \end{cases} \tag{4}$$

We notice that the model defined by equations (3) and (4) is not substantively changed if both sides of equation (3) are multiplied by a non-zero constant. This second identification problem is known as *multiplicative redundance*. We solve this problem by restricting the first element of Σ , σ_{11} , to be equal 1. This strategy is also used by McCulloch et al. (2000). Thus, accommodating both additive and multiplying redundance, we describe the fully identifiable MNP model as follows:

$$d_i = \begin{cases} 0 & \text{if } \max_{1 \leq l \leq p-1} z_{il} < 0 \\ j & \text{if } \max_{1 \leq l \leq p-1} z_{il} = z_{ij} > 0 \end{cases} \tag{5}$$

where $z_i \sim N(X_i \beta, \Sigma)$ and $\sigma_{11} = 1$.

2.2 MVMNP models for repeated nominal measures

Buliding on the notation of Section 2.1, we now extend the MNP model to multivariate nominal measures.

Suppose for each subject i , there are g nominal measures, the first with p_1 levels, the next with p_2 levels, and so on up to the last with p_g levels. Let $d_i = (d_{i1}, \dots, d_{ig})$ denote the index vector of the alternatives the i -th subject chooses for the g measures. Assume each of these g nominal measures follows an MNP model. Therefore, for the q -th measure, $q = 1, \dots, g$, there is a $(p_q - 1)$ -dimensional underlying utility vector z_{iq} satisfying equation (5) with mean equal to $X_{iq} \beta$ and covariance matrix equal to Σ_q with the upper left element $\{\Sigma_q\}_{11} = 1$.

We describe the MVMNP model for the g measures as follows:

$$z_i = X_i \beta + \epsilon_i$$

where $z_i^T = (z_{i1}, \dots, z_{ig})$ with $z_{iq} = (z_{iq1}, \dots, z_{iq(p_q-1)})$, $X_i = (X_{i1}^T, \dots, X_{ig}^T)^T$ and $\epsilon_i \sim N(0, \Sigma)$ with $\sigma_{qq} = 1$, where $q = 1, (p_1 + p_2 - 1), (p_1 + p_2 + p_3 - 2), \dots, (p_1 + p_2 + \dots + p_{g-1} - g + 1)$. We then specify

$$d_{iq} = \begin{cases} 0 & \text{if } \max_{1 \leq l \leq p_q-1} z_{iq l} < 0 \\ j & \text{if } \max_{1 \leq l \leq p_q-1} z_{iq l} = z_{iq j} > 0 \end{cases}$$

for $i = 1, \dots, n$ and $q = 1, \dots, g$.

Contingency table data can be accommodated as a special case of this model by setting each of the X_i matrices to an identity matrix of dimension g . The correlation structure of the MVMNP model allows associations between pairs of frequencies, much as would occur with a log-linear model with all two-way interactions. We explore this idea further in section 5.

The following section presents a Bayesian sampling algorithm for MVMNP models.

3 Bayesian sampling algorithm for MVMNP models

3.1 MCMC framework

The joint posterior density of β , Σ , and $Z = (z_1, \dots, z_n)$ given $d = (d_1, \dots, d_n)$ is characterized as

$$p(\beta, \Sigma, Z|d) \propto p(\beta) \times p(\Sigma) \times \prod_{i=1}^n [I_i \times \phi(z_i; X_i \beta, \Sigma)]. \tag{6}$$

where ϕ is the standard normal density function and $I_i = \prod_{q=1}^g I_{iq}$

$$I_{iq} = \mathbf{1}_{(d_{iq}=0, z_{iqj} < 0, j=1, \dots, p_q-1)} + \sum_{k=1}^{p_q-1} \mathbf{1}_{(d_{iq}=k, z_{iqk} = \max_{1 \leq l \leq p_q-1} (z_{iql}, 0))}$$

where $\mathbf{1}_E$ is a indicator function equal to 1 when expression E is true and 0 when E is false. Each I_i is thus an indicator function evaluating to 1 if the choice vector d_i is compatible with the latent vector z_i .

To implement our MCMC algorithms, we build on the following:

- Assuming $\beta \sim N(b, C)$ as a prior distribution for β and using standard Bayesian linear model results, $\beta|\Sigma, z, d$ has a multivariate normal distribution:

$$\beta|\Sigma, z, d \sim N(\widehat{\beta}, V_\beta),$$

where $V_\beta = \left(\sum_{i=1}^N X_i^T \Sigma^{-1} X_i + C^{-1}\right)^{-1}$ and $\widehat{\beta} = V_\beta \left(\sum_{i=1}^N X_i^T \Sigma^{-1} Z_i + C^{-1} b\right)$.

- The latent variable $z_{iqj}|\beta, \Sigma, d_i, z_{iq(-j)}, z_{i(-q)}$ has a truncated normal distribution that can be represented:

$$\begin{aligned} p(z_{iqj}|\beta, \Sigma, d_i, z_{iq(-j)}, z_{i(-q)}) &\propto I_i \times p(z_{iqj}|\beta, \Sigma, z_{iq(-j)}, z_{i(-q)}) \\ &= \left(\mathbf{1}_{(d_{iq}=j)} \mathbf{1}_{(z_{iqj} \geq \max_{1 \leq l \leq p_q-1} (z_{i,q,l}, 0))} + \mathbf{1}_{d_{iq} \neq j} \mathbf{1}_{(z_{iqj} \leq \max_{1 \leq l \leq p_q-1} (z_{i,q,l}, 0))} \right) \times \phi(z_{iqj}; \mu_{iqj}, \Sigma_{iqj}) \end{aligned}$$

where μ_{iqj} and Σ_{iqj} are the conditional mean and variance of z_{iqj} given $z_{iq(-j)}, z_{i(-q)}$. With missing $d_{ij}, z_{iqj}|\beta, \Sigma, d_{iq}, z_{iq(-j)}, z_{i(-q)}$ has a univariate normal distribution with mean μ_{iqj} and variance Σ_{iqj} as before, but without truncation.

- Assuming Σ has prior density $p(\Sigma)$, we have $p(\Sigma|\beta, Z, d)$ is proportional to $p(\Sigma) \times \prod_{i=1}^n \phi(z_i; X_i\beta, \Sigma)$. It is not easy to directly draw simulations from the posterior distribution of the covariance matrix Σ with g diagonal elements equal to 1. In the next section, we elaborate in detail the steps involved in drawing $p(\Sigma|\beta, Z, d)$ using the PX-MH algorithm.

3.2 Parameter-extended Metropolis-Hastings (PX-MH) step

We first give a brief review of the PX-MH algorithm proposed by Zhang et al. (2006). To sample a correlation matrix, R , in a multivariate probit model, Zhang et al. sample a covariance matrix, W , using the decomposition $W = D^{1/2}RD^{1/2}$ where D is a diagonal matrix of artificial variance components governed by a joint prior distribution, $p(R, D)$, for the correlation matrix R and D . We present the PX-MH algorithm as follows.

Set initial value of $(R^{(0)}, D^{(0)})$ through setting $W^{(0)} = D^{(0)\frac{1}{2}}R^{(0)}D^{(0)\frac{1}{2}}$ to an initial covariance matrix.

Then, at iteration $(t + 1)$

1. Generate (R^*, D^*) by generating $W^* = D^{*\frac{1}{2}}R^*D^{*\frac{1}{2}}$ from *Wishart* $(m, W^{(t)})$.
2. Take

$$(R^{(t+1)}, D^{(t+1)}) = \begin{cases} (R^*, D^*) & \text{with probability } \alpha \\ (R^{(t)}, D^{(t)}) & \text{otherwise.} \end{cases}$$

where $\alpha = \min \left\{ \frac{p(R^*, D^*|\beta, Z, Y) q(W^{(t)}|W^*)}{p(R^{(t)}, D^{(t)}|\beta, Z, Y) q(W^*|W^{(t)})}, 1 \right\}$. Here, $p(R, D|\beta, Z, Y)$ is the joint posterior density of (R, D) and $q(\cdot|W^{(t)})$, the proposal density, is equal to the product of the Jacobian term for the transformation $(W \rightarrow R, D)$ and the Wishart density with m degrees of freedom and scale matrix equal to $W^{(t)}$.

In the MVMNP model, the covariance matrix Σ has g diagonal elements equal to 1. We decompose $\Sigma = D^0RD^0$ where R is the correlation matrix of Σ and D^0 is the diagonal standard deviation matrix with elements $(D_{p_1}^0, D_{p_1}^0, D_{p_1+p_2-1}^0, \dots, D_{p_1+\dots+p_{g-1}-g-1}^0)$ equal to 1. Then we consider a diagonal matrix D replacing those elements of D^0 equal to 1 with unknown parameters $(v_1, v_{p_1}, v_{p_1+p_2-1}, \dots, v_{p_1+\dots+p_{g-1}-g-1})$. Therefore, the matrix $W = DRD$ is a covariance matrix without restrictions on the diagonal elements. We use the above PX-MH algorithm to sample W , thereby obtaining a draw of Σ . A slight distinction between sampling Σ in the MVMNP model and sampling R in the multivariate probit model is that some of the diagonal elements of D are identified parameters in the MVMNP model, while for the multivariate probit model, all the diagonal elements of D are artificial; this distinction does not alter the character of the algorithm, however.

For the prior distribution of Σ , we use a PXW prior proposed by Zhang et al. (2006), with density given by the product of the Jacobian term for the transformation $(W \rightarrow R, D)$ and the Wishart density with m_0 degrees of freedom and scale matrix equal to Λ . The scale matrix Λ reflects the prior guess for the covariance matrix Σ with higher values of m_0 representing greater prior precision.

Including the g artificial parameters, the joint posterior density of β, R, D, Z given d is

$$p(\beta, R, D, z|d) \propto p(\beta) \times p(R, D) \times \prod_{i=1}^N [I_i \times \phi(z_i; X_i \beta, \Sigma)].$$

The conditional distributions for β and z_i given other parameters are the same as described in Section 3.1. Through this joint posterior density, we have $p(R, D|\beta, Z, d)$ is proportional to

$p(R, D) \times \prod_{i=1}^N \phi(z_i; X_i \beta, \Sigma)$. As suggested above, the prior density $p(R, D)$ can be specified by letting the joint prior distribution of (R, D) be from the $PXW(m_0, \Lambda)$ family of distributions. Therefore, one cycle of the algorithm consists of Gibbs steps to sample β and each component of the latent variable z_i , and a Metropolis-Hastings step for sampling (R, D) , with Σ generated as a byproduct of the PX-MH step.

3.3 Software implementation

We implemented our algorithm in C using the GNU scientific library (Galassi et al., 2006). For the univariate MNP model, we compared our algorithm to algorithms proposed by McCulloch and Rossi (1994) (MR) and Imai and van Dyk (2005a) (IvD). We implemented the MR algorithm using the `rmnpGibbs` function in the `bayesm` package for R (Rossi, Allenby, & McCulloch, 2006) and the IvD algorithm using the MNP R package (Imai & van Dyk, 2005b). Using a variety of data sets, including simulated examples from McCulloch and Rossi (1994) and the detergent brand choice example included in the MNP package, we obtained similar results to both MR and IvD. Our convergence performance was comparable to MR, which is shown by Imai and van Dyk (2005b) to be slower than the IvD algorithm. Because our primary interest is to provide greater modeling flexibility, based on a Bayesian framework for multivariate nominal measures that allows prior information to be incorporated in a flexible and intuitive manner, we do not view the slower convergence of our method relative to IvD as a fatal flaw, but it does suggest that it is important to be careful in assessing convergence in multivariate applications. In the following sections, we illustrate the use of our algorithm through simulated data on multivariate nominal measures, followed by an analysis of data from a cancer-control study.

4 Illustration using simulated data

To illustrate our MCMC algorithm for the MVMNP model, we use the following simulated example to investigate posterior inference for unknown parameters.

We generated a data set with sample size equal to 2,000. Each subject i was assumed to have two nominal measures for each person, y_{i1} and y_{i2} , with y_{i1} having three categorical levels and y_{i2} having four categorical levels. We use z_{i1} to denote the two-dimensional latent variable corresponding to y_{i1} and z_{i2} to denote the three-dimensional latent variable corresponding to y_{i2} . The covariance matrix for z_{i1} was a 2 by 2 matrix with the correlation equal to 0.4 and variances equal to 1 and 0.81, respectively, and the covariance matrix for z_{i2} was taken to be an AR(1) correlation matrix with first-order correlations equal to 0.5 and variance vector equal to (1, 0.81, 0.90). The correlations between elements of z_{i1} and z_{i2} were all set to 0.2. The 2 by 1 covariate matrix X_i was generated from iid uniform (-0.5, 0.5), and the regression parameter β was set to 2.0.

To perform inference using the MVMNP model, we consider two alternative prior formulations for β and Σ . First, we take the prior distribution for β to be $N(0, 100)$, which is very weakly informative, and we assume Σ has a $PXW(m_0 = 8, I)$ distribution, i.e. a prior guess that the covariance matrix is equal to the identity matrix with eight degrees of freedom. A proper prior distribution for the five by five covariance matrix Σ requires m_0 to be greater than or equal to

6, and thus the prior distribution reflects a weakly informative belief in a scenario where the levels of the nominal variable have no association with one another. We also examined a strongly informative prior scenario with a $N(0, 1)$ prior distribution for β and a $PXW(m_0 = 100, CS(0.4))$ prior distribution for Σ , where $CS(0.4)$ indicates a compound symmetry structure with equal correlation 0.4. The degrees-of-freedom parameter m_0 in this case reflects a strong prior belief that the covariance matrix has a $CS(0.4)$ structure. We label the first approach PXW_ID_Weak and the second approach PXW_CS_Strong . The implications of these scenarios for the regression parameter β , the three variance parameters (σ_{22} , σ_{44} and σ_{55}), and the correlations r_{ij} are displayed in Figure 1, showing that the PXW_CS_Strong scenario gives much tighter information for all of the parameters than the PXW_ID_Weak scenario does.

We ran the MCMC algorithm for 101,000 iterations, discarding the first 1,000 iterations as a burn-in period for each of these two prior distribution scenarios. The posterior mean and standard deviation for each parameter are presented in Table 1, and the marginal posterior densities are shown in Figure 2. We see that the correlation parameters and the covariance parameters appear to depend somewhat on the specification of prior distributions. Not surprisingly, the posterior means of the correlations under the PXW_CS_Strong scenario are pulled toward the assumed value of 0.4. Also, the posterior standard deviations under the PXW_CS_Strong scenario are uniformly smaller than those under the PXW_ID_Weak scenario. The coverage of true values appears satisfactory in both scenarios, partly because substantial posterior uncertainty remains. It stands to reason that better prior specification may give better estimated values, but, in general, inference appears to be fairly robust to the choice between these two groups of priors, presumably because the sample size of 2,000 is sufficient to dominate either prior scenario.

The convergence of the MCMC algorithm was assessed by several procedures recommended by Cowles and Carlin (1996). We calculated Gelman and Rubin's potential scale reduction factor, $\sqrt{\hat{R}}$ for five dispersed chains with the first 1,000 iterations discarded as burn-in (Gelman & Rubin, 1992). A jumping distribution degrees-of-freedom parameter of $m = 2200$ gave an acceptance rate of about 10% for the PX-MH step of the algorithm. Although this is below the value of 23% recommended by Gelman, Roberts, and Gilks (1996), we find in practice that higher values for m substantially increased autocorrelations. For the single regression parameter β , ten correlation parameters r_{ij} and three variance parameters (σ_{22} , σ_{44} and σ_{55}), the values of $\sqrt{\hat{R}}$ were all below 1.1 after 40,000 iterations and declined consistently through a further 60,000 iterations. The multivariate potential scale reduction factor for these 14 parameters was 1.04 after 40,000 iterations, improving to 1.02 at 100,000 iterations.

5 Application to a cancer-control study

We illustrate our method for MVMNP models using incomplete contingency table data from a study on adherence to clinical recommendations among women diagnosed with breast abnormalities (Mojica, Bastani, Boscardin, & Ponce, 2006). The study aimed to test the effect of a telephone counseling intervention to encourage clinical follow-up for diagnosis of breast abnormalities, with telephone counselling or usual care subject to random assignment. The data were collected on a sample of 1,671 women who presented with a breast abnormality at two Los Angeles county hospitals (Hospitals A and B). Six months after enrollment, outcome data were collected via medical chart reviews and a computer-assisted telephone interview.

Here we investigate the between-measure and within-measure associations across the levels of six nominal variables: *Final Diagnosis* (Yes/No), *Patient Language* (Spanish/Other), *Ethnicity* (Hispanic/Non-Hispanic), *Age at Referral* (3 Categories), *Type of Referring Clinic* (grouped into 3 categories) and *Work Situation* (grouped into 4 categories). Descriptive

information about these nominal variables is presented in Table 2, showing that there are substantial rates of missingness for several of the measures. Of the $n = 1671$ subjects, only 314 had complete data on all six measures.

We used the proposed MCMC algorithm for MVMNP models including missing values and obtained posterior means and standard deviations for the 10×10 correlation matrix of the nominal variables. As discussed at the end of section 2.2, we set $X_i = I_{10}$ for each subject. We chose an independent $N(0, 100)$ prior for each regression parameter and a $PXW(m_0 = 18, I)$ prior for the covariance matrix for the latent vector z_i . We ran 401,000 iterations with 1,000 burn-in iterations, yielding a roughly 15% acceptance rate for the proposed draws in the PX-MH step. Values of univariate potential scale reduction factors (PSRF) for all 10 regression, 45 correlation and 4 variance parameters were well below 1.10 at 200,000 iterations with the exception of three correlation parameters (r_{45} , r_{67} , and r_{89}) for which the PSRF's were near 1.3 and a single variance parameter ($\sigma_{10,10}$) for which the PSRF was 1.5. The multivariate scale reduction factor for these 59 parameters was 1.3. A further 200,000 iterations reduced all scale reduction factors below 1.1.

In Table 3, we present the posterior mean correlation matrix and variances for the latent vector $z_i = (z_{i1}, \dots, z_{i10})$ with corresponding posterior standard deviations. The estimated correlations are not typically large in absolute values, with the correlation between the latent variables for Spanish language and Hispanic ethnicity an unsurprising exception that emerges as highly significant despite a large proportion of missing values in the patient language variables. This finding reinforces the ability of the proposed MCMC algorithm in Section 3 to handle data with many missing values. The negative correlations in the third row with the non-reference Working Situation levels suggest a slight tendency for Hispanic subjects to be homemakers. Obtaining final diagnosis is slightly more likely for older women, Working women are less likely than unemployed women to be referred from an Emergency or Medical Walk-in clinic. The four variance components for z_{i5} , z_{i7} , z_{i9} and z_{i10} , had 95% credible intervals that include 1, suggesting that the latent variables underlying the multivariate nominal measures are roughly on the same scale.

Since the data can be represented in a multi-way table, we now compare our approach to a simple log-linear model with all main effects and two-way interactions. The log-linear modeling was performed using the SAS Catmod procedure which only makes use of complete cases. To contrast the two approaches, we used a reduced set of variables: *Final Diagnosis* with 2 levels, *Age at Referral* with 3 levels, and *Type of Referring Clinic* with 3 levels. The data for this $2 \times 3 \times 3$ table have $n = 1327$ complete cases. For the MVMNP model, we calculated posterior mean fitted probabilities for the 18 cells in the contingency table; the correlation of these 18 values with the corresponding fitted probabilities from the two-way log-linear model was 0.99. Thus, as expected, our method coincides very closely to the standard log-linear model predictions in the situation of a large sample size and a simple model. We note for this reduced set of variables that the MVMNP model has 17 parameters (10 correlation parameters, 2 variance parameters, and 5 mean parameters) and the two-way log-linear model has 13 parameters (5 main effects and 8 two-way interactions). Thus the latent variable structure in the MVMNP model gives it somewhat more flexibility in this complete case setting. More importantly, the MVMNP model is able to directly accommodate incomplete data. In the original setting of six variables, the MVMNP model uses incomplete data on $n = 1671$ subjects to fit a model with 59 parameters (45 correlation parameters, 4 variance parameters, and 10 mean parameters). In contrast, a two-way log-linear model has 50 parameters (10 main effects and 40 two-way interactions); these 50 parameters would be poorly estimated using the $n = 314$ complete cases.

6 Conclusions

In this manuscript, we proposed MVMNP models for analyzing multivariate nominal data using the PX-MH algorithm in a Bayesian framework. We illustrated our methodology through a simulated example and an application to incomplete contingency table data from a cancer control study.

The MVMNP model is a general model to handle multivariate nominal data that has several advantages over MNP models and multinomial multiperiod probit models. First, the MNP model is a special case for univariate nominal data using the MVMNP model, whereas the multiperiod multinomial probit model does not have this property. Second, the MVMNP model allows general covariance structures for the latent variables, compared with the multiperiod multinomial probit model which only allows special covariance structures, such as those with AR(1) structures. Third, the prior distribution we used for the MVMNP models allows useful prior information to be incorporated into the model in a flexible manner that is naturally embedded into the sampling algorithm.

In the special case of contingency table data, the MVMNP model provides additional flexibility over a two-way log-linear model, but is more parsimonious than a saturated model. The MVMNP model can also directly accommodate incomplete data as well as covariates.

An inherent concern in the PX-MH algorithm for the MVMNP model is that with high-dimensional repeated nominal measures, the PX-MH algorithm may have high autocorrelations among the posterior draws and therefore result in slow convergence. The identified MVMNP model necessitates a complex Metropolis-Hastings step for sampling the restricted covariance matrix instead of a Gibbs sampler step for the unidentified version. Imai and van Dyk (2005a) have given a similar discussion for the slow convergence induced by solving the identification problem for the MNP model. Statistical computing strategies that offer the prospect of more rapid convergence would be a worthy area for future research.

Acknowledgements

This work was supported by NIH grants NS 30308 and AI 28697 (WJB) and MH 60213 (all).

References

- Albert JH, Chib S. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 1993;88:669–679.
- Bhat CR. A multiple discrete-continuous extreme value model: Formulation and application to discretionary time-use decisions. *Transportation Research Part B* 2005;39:679–707.
- Bhat, CR. The multiple discrete-continuous extreme value (mdcev) model: Role of utility function parameters, identification considerations, and model extensions (Tech. Rep.). Department of Civil Engineering; The University of Texas, Austin: 2006.
- Bierlaire, M. Discrete choice models. In: Labb, KTM.; Laporte, G.; Toint, P., editors. *Operations Research and Decision Aid Methodologies in Traffic and Transportation Management*. Vol. Vol. 166. Springer Verlag; New York: 1998. Series F: Computer and Systems Sciences
- Bishop, Y.; Fienberg, S.; Holland, P. *Discrete multivariate analysis*. The MIT Press; Cambridge, Massachusetts: 1975.
- Börsch-Supan A, Hajivassiliou VA. Smooth unbiased multivariate probability simulators for maximum likelihood estimation of limited dependent variable models. *Journal of Econometrics* 1993;58:347–368.
- Boscardin, WJ.; Zhang, X. Modeling the covariance and correlation matrix of repeated measures. In: Gelman, A.; Meng, X-L., editors. *Applied Bayesian Modeling and Causal Inference from an Incomplete-Data Perspective*. John Wiley & Sons; New York: 2004.

- Carlin, BP.; Louis, TA. Bayes and empirical Bayes methods for data analysis. Vol. Second ed. Chapman & Hall/CRC; New York: 2000.
- Chen, M-H.; Shao, Q-M.; Ibrahim, JG. Monte Carlo methods in Bayesian computation. Springer Series in Statistics; New York: 2000.
- Chen Z, Kuo L. Discrete choice models based on the scale mixture of multivariate normal distributions. *The Indian Journal of Statistics* 2002;64:192–213.
- Chib, S.; Greenberg, E.; Chen, Y. MCMC methods for fitting and comparing multinomial response models (Tech. Rep. No. 9802001). Economics Working Paper Archive; Washington University at St. Louis: 1998.
- Congdon, P. Bayesian statistical modelling. Wiley; New York: 2001.
- Cowles MK, Carlin BP. Markov chain Monte Carlo convergence diagnostics: a comparative review. *Journal of American Statistical Association* 1996;91:883–904.
- Daganzo, C. Multinomial Probit. Academic Press; New York: 1980.
- Galassi, M.; Davies, J.; Theiler, J.; Gough, B.; Jungman, G.; Booth, M., et al. Gnu scientific library reference manual. Network Theory Limited; Bristol: 2006.
- Gamerman, D. Markov chain Monte Carlo. CRC Press; New York: 1998.
- Gelfand AE, Smith AFM. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 1990;85:398–409.
- Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. Bayesian data analysis. Vol. Second ed. Chapman & Hall/CRC; New York: 2003.
- Gelman A, Roberts GO, Gilks WR. Efficient jumping rules for metropolis algorithms. *Bayesian Inference* 1996;5:595–605.
- Gelman A, Rubin DB. Inference from iterative simulation using multiple sequences (with discussion). *Statistical Science* 1992;7:457–511.
- Geweke J, Keane M, Runkle D. Alternative computation approaches to inference in the multinomial probit model. *Review of Economics and Statistics*, 1994 1994a;76:609–632.
- Geweke, J.; Keane, M.; Runkle, D. Proceedings of the american statistical association. American Statistical Association; Alexandria, VA: 1994b. Recursively simulating multinomial multiperiod probit probabilities.
- Geweke J, Keane M, Runkle D. Statistical inference in the multinomial multiperiod probit model. *Journal of Econometrics* 1997;80:125–165.
- Gilks, WR.; Richardson, S.; Spiegelhalter, DJ., editors. Markov chain Monte Carlo in practice. CRC Press; New York: 1996.
- Golob T, Regan A. Trucking industry adoption of information technology: a multivariate discrete choice model. *Transportation Research - Part C: Emerging Technologies* 2002;10:205–228.
- Hausman J, Wise D. A conditional probit model for qualitative choice: Discrete decisions recognizing interdependence and heterogeneous preferences. *Econometrica* 1978;45:319–339.
- Hensher D, Greene W. The mixed logit model: The state of practice and warnings for the unwary. *Transportation* 2003;30:133–176.
- Imai K, van Dyk DA. A Bayesian analysis of the multinomial probit model using marginal data augmentation. *Journal of Econometrics* 2005a;124:311–334.
- Imai K, van Dyk DA. MNP: R package for fitting the multinomial probit model. *Journal of statistical software* 2005b;14:1–32.
- Kim J, Allenby GM, Rossi PE. Modeling consumer demand for variety. *Marketing Science* 2002;21:229–250.
- Liang K-Y, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika* 1986;73:13–22.
- Liang K-Y, Zeger SL, Qaqish B. Multivariate regression analyses for categorical data. *Journal of the Royal Statistical Society. Series B* 1992;54:3–40.
- Liu, JS. Monte Carlo strategies in scientific computing. Springer; New York: 2001.
- McCulloch RE, Polson NG, Rossi PE. A Bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics* 2000;99:173–193.

- McCulloch RE, Rossi PE. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics* 1994;64:207–240.
- McFadden D. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica* 1989;57:995–1026.
- McFadden D, Ruud P. Estimation with simulation. *Review of Economics and Statistics* 1994;76:591–608.
- McFadden D, Train K. Mixed MNL models for discrete response. *Journal of Applied Econometrics* 2000;15:447–470.
- Mojica CM, Bastani R, Boscardin WJ, Ponce NA. Lowincome women with breast abnormalities: System predictors of timely diagnostic resolution. *Cancer Control*. 2006(Conditionally accepted)
- Muthen B. Latent variable structural equation modeling with categorical data. *Journal of Econometrics* 1983;22:43–65.
- Nobile A. A hybrid Markov chain for the Bayesian analysis of the multinomial probit model. *Statistics and Computing* 1998;8:229–242.
- Nobile A. Comment: Bayesian multinomial probit models with normalization constraint. *Journal of Econometrics* 2000;99:335–345.
- Rendtel, U.; Kaltenborn, U. The stability of simulation based estimation of the multiperiod multinomial probit model with individual specific covariates (Tech. Rep.). *Diskussionsbeitrge des Fachbereichs Wirtschaftswissenschaft der FU Berlin (Volkswirtschaftliche Reihe)*; Berlin: 2004.
- Robert, CP.; Casella, G. Monte Carlo statistical methods. Vol. Second ed. Springer; New York: 2004.
- Rossi, P.; Allenby, GM.; McCulloch, R. Bayesian statistics and marketing. Wiley; New York: 2006.
- Sivakumar A, Bhat CR, Ökten G. Simulation estimation of mixed discrete choice models with the use of randomized quasi-monte carlo sequences: A comparative study. *Transportation Research Record* 2005;1921:112–122.
- Tanner MA, Wong WH. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 1987;82:528–540.
- Train, K. A comparison of hierarchical Bayes and maximum simulated likelihood for mixed logit (Tech. Rep.). Department of Economics; University of California, Berkeley: 2001.
- Train, K.; Sonnier, G. Mixed logit with bounded distributions of correlated partworths. In: Alberini, A.; Scarpa, R., editors. *Applications of Simulation Methods in Environmental Resource Economics*. Springer; New York: 2005.
- Zeger SL. The analysis of discrete longitudinal data: Commentary. *Statistics in Medicine* 1987;7:161–168.
- Zeger SL, Liang K-Y. Longitudinal data analysis for discrete and continuous outcomes. *Biometrics* 1986;42:121–130. [PubMed: 3719049]
- Zhang X, Boscardin WJ, Belin TR. Sampling correlation matrices in Bayesian models with correlated latent variables. *Journal of Computational Graphics and Statistics* 2006;15:880–896.
- Ziegler, A. Simulated classical tests in the multiperiod multinomial probit model (Tech. Rep.). Center for European Economic Research; Mannheim: 2002.

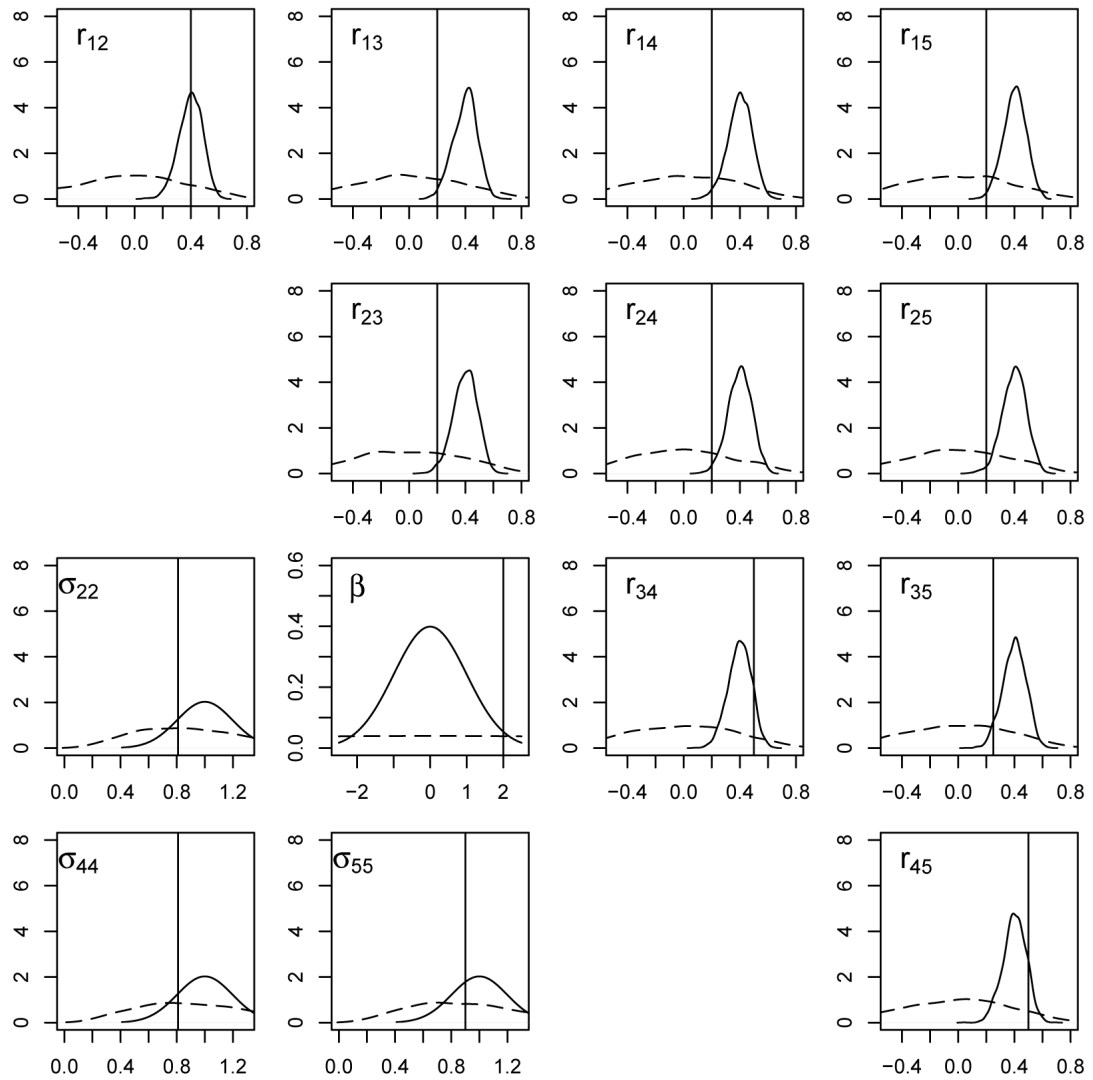


Figure 1.

Prior density plots for the regression parameter (β), variance parameters (σ_{22} , σ_{44} , σ_{55}), and correlations (r_{ij}). The dotted lines are for the PXW_ID_Weak scenario, the solid lines are for the PXW_CS_Strong scenario, and vertical lines have been drawn at the true value for each parameter.

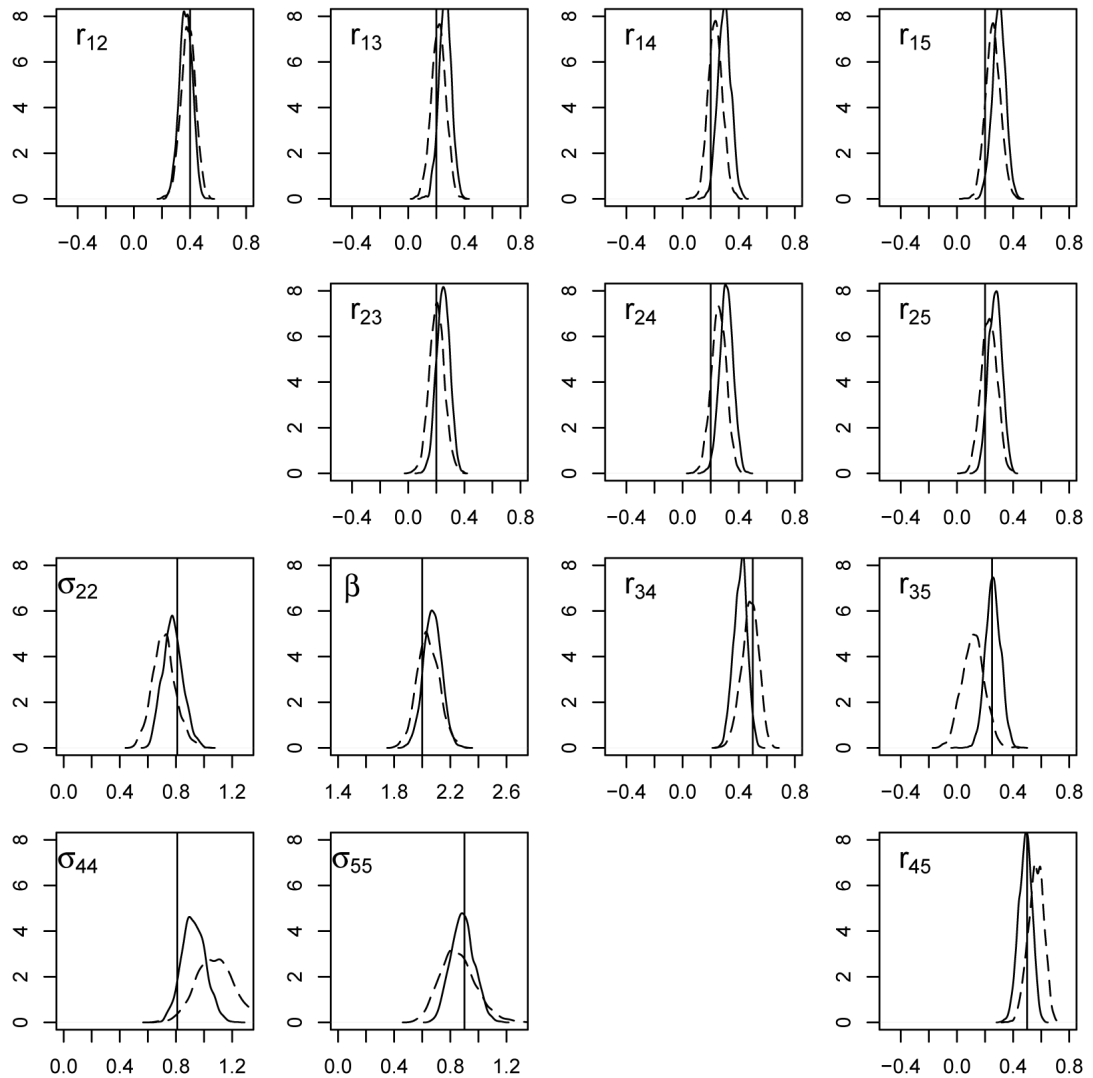


Figure 2.

Posterior density plots for the regression parameter (β), variance parameters (σ_{22} , σ_{44} , σ_{55}), and correlations (r_{ij}). Dotted lines correspond to the PXW_ID_Weak scenario, and solid lines correspond to the PXW_CS_Strong scenario. Vertical lines have been drawn at the true values for each parameter. The plots show some sensitivity to prior specification for variance parameters, somewhat less sensitivity in correlation parameters (with r_{35} a possible exception), and little sensitivity for the regression parameter.

Table 1

Posterior means (posterior standard deviations) for the regression parameter (β), estimable variance parameters (σ_{22} , σ_{44} and σ_{55}) and correlation parameters (r_{ij}) under PXW_ID_Weak and PXW_CS_Strong scenarios. This table shows that the posterior means for all parameters, except r_{35} , are similar under these two groups of priors and the posterior standard deviations under the PXW_CS_Strong prior are uniformly smaller than those under the PXW_ID_Weak prior.

Parameters	True	PXW_I_weak	PXW_CS_strong
β	2.00	2.04 (0.08)	2.07 (0.07)
σ_{22}	0.81	0.72 (0.08)	0.78 (0.08)
σ_{44}	0.90	1.08 (0.14)	0.93 (0.09)
σ_{55}	0.81	0.85 (0.13)	0.89 (0.09)
r_{12}	0.40	0.39 (0.23)	0.37 (0.21)
r_{13}	0.20	0.22 (0.23)	0.27 (0.21)
r_{14}	0.20	0.23 (0.23)	0.30 (0.21)
r_{15}	0.20	0.26 (0.23)	0.30 (0.21)
r_{23}	0.20	0.20 (0.23)	0.26 (0.22)
r_{24}	0.20	0.26 (0.23)	0.31 (0.22)
r_{25}	0.20	0.23 (0.24)	0.27 (0.22)
r_{34}	0.50	0.48 (0.25)	0.42 (0.22)
r_{35}	0.25	0.12 (0.28)	0.25 (0.24)
r_{45}	0.50	0.56 (0.23)	0.48 (0.22)

Table 2

Description of variables used in cancer-control data analysis.

Variable	Categories	Percentage	Missing %	Latent Variables
<i>Final Diagnosis</i>	Yes No	55.1 44.9	0.0	$z_{i,1}$ (reference)
<i>Patient Language</i>	Spanish Other	61.0 39.0	63.3	$z_{i,2}$ (reference)
<i>Ethnicity</i>	Hispanic Non-Hispanic	85.0 15.0	36.2	$z_{i,3}$ (reference)
<i>Age at Referral</i>	< 40 40–49 ≥ 50	26.3 31.5 42.2	2.8	(reference) $z_{i,4}$ $z_{i,5}$
<i>Type of Referring Clinic</i>	Breast and Tumor Clinic Emergency and Medical Walk-in Other	35.4 22.4 42.2	18.5	$z_{i,6}$ $z_{i,7}$ (reference)
<i>Working Situation</i>	Full-time Part-time Unemployed Homemaker	18.4 20.0 23.6 38.1	36.0	$z_{i,8}$ $z_{i,9}$ $z_{i,10}$ (reference)

Table 3
 Posterior means (posterior standard deviations) for the variances and correlations of the latent vector $z_i = (z_{i1}, \dots, z_{i,10})$.

Diagnosis	Language			Ethnicity			Age			Referring Clinic			Working Situation		
	z_{i1}	z_{i2}	z_{i3}	z_{i4}	z_{i5}	z_{i6}	z_{i7}	z_{i8}	z_{i9}	$z_{i,10}$					
z_{i1}	1.00 (0.00)	0.05 (0.25)	0.05 (0.25)	0.19 (0.23)	-0.07 (0.21)	0.00 (0.21)	-0.01 (0.21)	0.03 (0.24)	0.00 (0.24)	0.03 (0.24)					
z_{i2}	0.05 (0.25)	1.00 (0.00)	0.85 (0.20)	-0.12 (0.24)	0.06 (0.27)	0.08 (0.26)	-0.03 (0.26)	-0.17 (0.28)	0.04 (0.31)	-0.39 (0.27)					
z_{i3}	0.05 (0.25)	0.85 (0.20)	1.00 (0.00)	-0.12 (0.26)	-0.01 (0.25)	0.05 (0.25)	-0.17 (0.25)	-0.23 (0.26)	-0.10 (0.27)	-0.39 (0.24)					
z_{i4}	0.19 (0.23)	-0.12 (0.24)	-0.12 (0.26)	1.00 (0.00)	-0.25 (0.41)	-0.21 (0.25)	-0.48 (0.23)	0.17 (0.26)	0.04 (0.27)	0.00 (0.27)					
z_{i5}	-0.07 (0.21)	0.06 (0.27)	-0.01 (0.25)	-0.25 (0.41)	0.90 (0.24)	0.06 (0.23)	-0.02 (0.24)	0.00 (0.25)	-0.01 (0.26)	-0.01 (0.25)					
z_{i6}	0.00 (0.21)	0.08 (0.26)	0.05 (0.25)	-0.21 (0.25)	0.06 (0.23)	1.00 (0.00)	-0.03 (0.47)	0.10 (0.26)	-0.14 (0.26)	0.08 (0.26)					
z_{i7}	-0.01 (0.21)	-0.03 (0.26)	-0.17 (0.25)	-0.48 (0.23)	-0.02 (0.24)	-0.03 (0.47)	0.83 (0.38)	-0.23 (0.26)	-0.12 (0.26)	0.21 (0.26)					
z_{i8}	0.03 (0.24)	-0.17 (0.28)	-0.23 (0.26)	0.17 (0.26)	0.00 (0.25)	0.10 (0.26)	-0.23 (0.26)	1.00 (0.00)	0.00 (0.45)	0.07 (0.45)					
z_{i9}	0.00 (0.24)	0.04 (0.31)	-0.10 (0.27)	0.04 (0.27)	-0.01 (0.26)	-0.14 (0.26)	-0.12 (0.26)	0.00 (0.45)	0.93 (0.31)	0.02 (0.44)					
$z_{i,10}$	0.03 (0.23)	-0.31 (0.27)	-0.39 (0.24)	0.00 (0.27)	-0.01 (0.25)	0.08 (0.26)	0.21 (0.26)	0.07 (0.45)	0.02 (0.44)	1.18 (0.28)					