A Bayesian Latent Variable Mixture Model for Longitudinal Fetal Growth

James C. Slaughter[*],
Vanderbilt University School of Medicine, Department of Biostatistics, S-2323 Medical Center North, 1161 21st Ave South, Nashville, TN 37232-2158

Amy H Herring, and
The University of North Carolina, Department of Biostatistics, McGavran-Greenberg Hall, CB 7420, Chapel Hill, NC 27599-7420

John M Thorp
University of North Carolina School of Medicine, 4012 Old Clinical Building, CB 7570, Chapel Hill, NC 27599-7570

Summary

Fetal growth restriction is a leading cause of perinatal morbidity and mortality that could be reduced if high risk infants are identified early in pregnancy. We propose a Bayesian model for aggregating 18 longitudinal ultrasound measurements of fetal size and blood flow into three underlying, continuous latent factors. Our procedure is more flexible than typical latent variable methods in that we relax the normality assumptions by allowing the latent factors to follow finite mixture distributions. Using mixture distributions also permits us to cluster individuals with similar observed characteristics and identify latent classes of subjects who are more likely to be growth or blood flow restricted during pregnancy. We also use our latent variable mixture distribution model to identify a clinically-meaningful latent class of subjects with low birth weight and early gestational age. We then examine the association of latent classes of intrauterine growth restriction with latent classes of birth outcomes as well as observed maternal covariates including fetal gender and maternal race, parity, body mass index (BMI), and height. Our methods identified a latent class of subjects who have increased blood flow restriction and below average intrauterine size during pregnancy who were more likely to be growth restricted at birth than a class of individuals with typical size and blood flow.

Keywords

Bayesian methods; birth weight; correlated data; intrauterine growth restriction; pre-term birth; latent variables; small for gestational age

1. Introduction

We consider a latent variable model for identifying intrauterine growth restriction using multiple longitudinal ultrasound measurements collected during pregnancy (Figure 1: $S \rightarrow$

[*]james.c.slaughter@vanderbilt.edu.

$\eta_1 \rightarrow y$). Growth restriction during pregnancy, or intrauterine growth restriction, cannot be observed directly, but ultrasound measurements of fetal size and blood flow are useful indicators of this concept. Fetal size measurements, including the head circumference (HC), abdominal circumference (AC), biparietal diameter (BPD) and femur length (FL), have be used identify "growth restricted" individuals when one of these measurements falls below a specific gestational-age adjusted threshold (Maulik, 2006). Women with high uterine and umbilical artery blood flow resistance are at an increased risk for delivering a growth-restricted infant because the fetus may not receive enough oxygen or nutrients during pregnancy (Dugoff et al., 2005; Hugo et al., 2007).

Blood flow resistance can be measured using multiple Doppler ultrasound measurements on different arteries (uterine and umbilical), at different locations (left and right), and at multiple times during the pregnancy.

While many ultrasound measurements are available to diagnose intrauterine growth restriction, it is not clear how best to combine and use all of the measurements. The most common approach compares the measurements of the HC, AC, BPD, and FL to predicted sizes for a gestational age based on population studies to define second trimester growth restriction (e.g. Filly and Hadlock, 2000; Dugoff et al., 2005). This approach does not address situations in which different measurements disagree. Other researchers have suggested using a simple average of all of the fetal size measurements for prediction of subsequent growth restriction (Ott, 1994). Instead, we aggregate the multiple ultrasound measurements into underlying latent factors (Figure 1: $\eta_1 \rightarrow y$), which we postulate arise from a fixed number of latent classes (Figure 1: $S \rightarrow \eta_1$).

We are interested in examining the association of growth restriction from the 15th to 24th week of pregnancy with functions of birth weight and gestational age at delivery (Figure 1: $S \rightarrow T$). In general, populations with lower mean birth weight and earlier gestational age have poorer infant survival, but closer inspection reveals that the increased risk is due to a higher proportion of births in the lower tails of the birth weight and gestational age distributions (Buekens et al., 2000). For birth weight in particular, Wilcox (2001) has described the skewed distributions as arising from a "predominant" and (high-risk) "residual" component that contains the long lower tails (Figure 2). In an attempt to identify high-risk infants, outcomes such as low birth weight (< 2500 grams) and pre-term birth (< 37 weeks) have been studied extensively, but these dichotomies are prone to misclassify individuals into the predominant and residual components. Instead, we develop a latent variable model to identify a latent class of subjects who are more likely to belong to the residual component and fall in the high-risk tails of the birth weight and gestational age distributions (Figure 1: $T \rightarrow \eta_{21} \rightarrow z_1$ and z_2). Additionally, low birth weight can be due to slow growth throughout pregnancy, so we also consider the association between intrauterine growth restriction class and birth weight z-scores (Figure 1: $S \rightarrow z_3$), which are a gestational-age-standardized measure of growth restriction at birth (Oken et al., 2003).

We allow the latent factors to follow finite mixture distributions to relax the assumption of marginal normality of the growth measures and group individuals into latent classes. The longitudinal ultrasound measurements allow us to identify intrauterine growth restriction classes. Using the birth outcomes, we formalize the concept of "residual" and "predominant" distribution components to identify a group of subjects who are more likely to have low birth weight and early gestational age. We then examine the association between classes of intrauterine growth restriction and (1) having clinically-relevant early gestational age and lower birth weight, as measured by belonging to the "residual" class, and (2) growth restriction at birth, as measured by birth weight Z-scores. Finally, we incorporate observed covariates known to be related to fetal size including maternal race, parity (number of prior

live births), BMI, and fetal gender by allowing the covariates to predict latent class membership.

Latent variable models attempt to reduce dimensionality by summarizing several observed variables into a smaller number of (continuous) latent factors or by assuming that each subject can be classified into one of a number of fixed, known latent classes. Several latent class models have been suggested to accommodate longitudinally-collected categorical data by conceptualizing that they are indicative of an underlying latent ordinal state or nominal latent class (Reboussin et al., 1999; Reboussin and Anthony, 2001; Huang and Bandeen-Roche, 2004). Additionally, authors have proposed general growth mixture models (GGMM) which assume that longitudinal measurements will follow a growth trajectory given a latent class (Muthén and Shedden, 1999; Muthén et al., 2002; Elliott et al., 2005). Our approach first aggregates non-commensurate longitudinal outcomes into a fixed number of continuous latent factors while allowing the latent factors to follow mixture distributions. In this way we allow the latent factors to follow latent trajectories where the mean and variance of the latent factor is conditional on latent class. As in the GGMM, interest lies in examining the association between latent classes and observed covariates, and our model also allows for relationships among latent classes.

The remainder of this paper is organized as follows. In Section 2, we propose our general model for latent variable mixture distributions and compare our method to previous approaches. We provide a Bayesian approach to fitting such a model while focusing on our fetal development application. Section 3 contains the results of our analysis, followed by a discussion in Section 4.

2. Methods

2.1 Fetal Growth Model

To analyze fetal growth, we propose the following general latent variable mixture model with covariates (Figure 1)

$$y_{ij}|\eta_{i1} \sim N\left(\boldsymbol{W}_{ij}^{'}\Gamma_j + \eta_{i1}^{'}\Lambda_j, \tau_{y,j}^{-1}\right) \quad (1)$$

$$\eta_{i1}|S_i{=}k \sim N_{q1}\left(\mu_{1k}, \Sigma_{1k}\right) \quad (2)$$

$$S_i|\boldsymbol{x}_i \sim Multi\left(1; \pi\left(\omega_1, \boldsymbol{x}_i\right), \ldots, \pi\left(\omega_K, \boldsymbol{x}_i\right)\right) \quad (3)$$

$$\pi\left(\omega_k, \boldsymbol{x}_i\right) = \Phi\left(\omega_k^{'}\boldsymbol{x}_i\right) \quad (4)$$

$$z_{ig}|\eta_{i2} \sim N\left(\eta_{i2}^{'}\Theta_g + \sum_{k=1}^{K}\beta_{g,k}I\left(S_i{=}k\right), \tau_{z,g}^{-1}\right) \quad (5)$$

$$\eta_{i2}|T_i{=}l \sim N_{q2}\left(\mu_{2l}, \Sigma_{2l}\right) \quad (6)$$

$$Pr\left(T_i{=}0|S_i\right) \sim Bin\left(g\left(\alpha_k, S_i\right)\right) \quad (7)$$

$$g\left(\alpha_k, S_i\right) = \Phi\left(\sum_{k=1}^{K} \alpha_k I\left(S_i = k\right)\right) \quad (8)$$

where $\Phi(\cdot)$ is the normal cumulative distribution function and $I(S_i = k)$ is the indicator function that takes a value of one if $S_i = k$ and zero otherwise. We assume that each of our observed outcomes measured during pregnancy, y_{ij} ($i = 1, \ldots, n$ and $j = 1, \ldots, p_1$), follows a normal distribution with a mean that is a function of observed covariates $W_{ij}(r_j \times 1)$ and latent factors $\eta_{i1}(q_1 \times 1)$ with parameter vectors $\Gamma_j(r_j \times 1)$ and $\Lambda_j(q_1 \times 1)$, respectively. The $\tau_{y,j}$ depend on j because the y_{ij} may be measured on different scales for each j. Outcomes measured at birth, z_{ig} ($g = 1, \ldots, p_2$), such as birth weight and gestational age, also follow a normal distribution with a mean that is a function of latent immaturity (η_{i2}) with parameters $\theta_g(q_2 \times 1)$ and latent class S_i with covariates $\beta_g = [\beta_{g,1}, \ldots, \beta_{g,K}]'$. We allow $\eta_{i1}(q_1 \times 1)$ and $\eta_{i2}(q_2 \times 1)$ to follow finite mixture distributions. It is often convenient to express mixture models using a missing data formulation in which each η_{i1} and η_{i2} is presumed to arise from a specific, but unknown, underlying component (Dempster et al., 1977). Specifically, for η_{i1} and η_{i2}, respectively, we introduce latent class allocation variables $S_i \in \{1, \ldots, K\}$ where $\Pr(S_i = k) = \pi_{s,k}$ for early growth restriction and $T_i \in \{1, \ldots, L\}$ where $\Pr(T_i = l) = \pi_{t,l}$ for latent immaturity. This specification is useful for computational purposes and allows us to naturally group subjects with similar latent variable characteristics. We can then jointly examine possible associations between measured covariates x_i and S_i using parameters ω as well as associations between T_i and S_i using parameters a by following Bayesian techniques for fitting probit regression models (Albert and Chib, 1993).

Latent variable and mixture models are over-identified and require parameter constraints for identifiability and interpretation. One common approach sets the covariance matrices Σ_{1k} and Σ_{2l} in (2) and (6) to be equal to the identity matrix so that all of the elements Λ_j and Θ_g in (1) and (5) can be identified an interpreted as factor loadings. Instead, we estimate an equivalent model in which Σ_{1k} has diagonal elements ($\tau_{1,1,k}, \ldots, \tau_{1,q1,k}$) and Σ_{2l} has diagonal elements ($\tau_{2,1,l}, \ldots, \tau_{2,q2,l}$) with all covariance terms fixed to zero. This specification requires fixing q_1 factor loadings in Λ (where Λ is a stacked matrix of the Λ_j) and q_2 factor loadings in Θ (a stacked matrix of the Θ_j) to one so that each of the q_1 latent factors in η_1 and q_2 latent factors in η_2 will have a scale that is commensurate with a specific ultrasound or birth outcome. Such a specification aids in specifying appropriate prior distributions (2.6) as well as easing interpretation of the latent factors. Elliott et al. (2005) pursues an alternative approach in which Λ is assumed to follow a polynomial function of known time points so that Σ_{1k} can be estimated. Their approach may be more appropriate when the longitudinal observations are measured at more time points than in our application.

2.2 Mixture Distribution with Covariates

The sensitivity of parametric latent variable models to distributional assumptions limits their general use, especially as the assumptions are often difficult to evaluate. Our data are further complicated by the unusual distribution of gestational age at delivery, which has a long left tail for early ages and a near truncation on the right due to medically-induced labor for longer gestations (Figure 2). Because symmetry and normality are particularly unlikely to hold in our example, we propose an alternative method in which the latent factors follow mixture distributions. Proposed approaches that are robust to outliers (Lee and Xia, 2006) would be appropriate if we had symmetrically heavy-tailed distributions. Early births are particularly indicative of future problems, so that cutoffs at 37 (pre-term) and 32 (very pre-term) weeks of gestation have been established in the reproductive health literature, with gestational age at delivery treated as a binary outcome. Our mixture model approach will

allow us to identify latent classes of subjects who are more likely to deliver early, without limiting the analysis to these pre-defined cut points. We then propose a model for examining the association between a latent growth restriction class, based on measures during pregnancy, and latent immaturity.

To relax the marginal normality assumption, we allow the latent factors to follow finite mixture distributions given by $f(\eta_{i1}) \sim \Sigma_{k=1}^{K} \pi_{1k} N_3 \left([\mu_{11k}, \mu_{12k}, \mu_{13k}]', Dg \left(\tau_{11k}^{-1}, \tau_{12k}^{1}, \tau_{13k}^{-1} \right) \right)$ and $f(\eta_{i2}) \sim \Sigma_{l=1}^{L} \pi_{21l} N \left(\mu_{21l}, \tau_{21l}^{-1} \right)$ where $D_g \left(\tau_{11k}^{-1}, \tau_{12k}^{-1}, \tau_{13k}^{-1} \right)$ is a diagonal covariance matrix with elements τ_{1mk}. In our fetal growth example, we use a two-component mixture for early growth restriction ($K = 2$) and a three-component mixture for latent immaturity ($L = 3$).

Latent class models are subject to additional identifiability complications due to the fact that the likelihood is symmetric across the possible permutations of class membership (Stephens, 2000). Consequently, assignment to a particular class k during one iteration of the Gibbs sampler may not have the same meaning in terms of model structure as assignment to class k at another iteration of the Gibbs sampler. For the latent immaturity variable η_2, we impose the identifiability constraint that $\mu_{211} < \mu_{212} < \mu_{213}$ to deal with the "label switching" problem and to ensure that subjects with the lower birth weights and earlier gestational ages will be assigned to the first latent class $T_i = 1$.

We then propose a probit model for the probability of belonging to the immature group as a stochastic function of early growth restriction class (S_i)

$$ \Pr \quad (T_i = 1) = \Phi \left(\sum_{k=1}^{K} I(S_i = k) \alpha_k \right), \quad (9) $$

where Φ is the normal cumulative distribution function. We also use a probit model to examine the association between observed covariates and early growth restriction class, $\Pr(S_i = 2) = \Phi(x'_i \omega_k)$, where x_i ($r \times 1$) includes covariates of interest with parameters ω ($r \times 1$). In the fetal development example, we consider the covariates fetal gender and maternal race, parity, gender, body mass index (BMI), and height. To fit these probit regression models, we use the data augmentation algorithm of Albert and Chib (1993).

Our Bayesian approach for fitting finite mixture distributions is based on the work of Diebolt and Robert (1994) and Richardson and Green (1997), but is complicated by needing to model latent rather than observed variables. To model the association of latent group membership with covariates, we utilize data augmentation. In this procedure, auxiliary data augment the observed data so we can use both the auxiliary and observed data to calculate the posterior distribution of the parameters of interest (van Dyk and Meng, 2001). Alternatively, Fokoue (2005) has proposed an EM algorithm for a normal latent variable mixture model with covariates, but does not model associations between the latent classes.

2.3 Model selection

A complication in mixture distribution models with a finite number of components is the method used to select the number of mixture components. A more general model could treat the number of mixture components as parameters to be estimated (Richardson and Green, 1997) or use a Dirichlet process in which the number of components is countably infinite (Dunson, 2006), but neither of these approaches would be able to readily incorporate covariates that predict class membership. When possible, we prefer an approach for selecting the number of mixture components that is guided by the application. For example, in the reproductive epidemiology literature birth weight and gestational age have been

described as arising from a predominant and residual component that is indicative of early gestational age and low birth weight (Wilcox, 2001; Buekens et al., 2000). We use a three component mixture for latent immaturity class (T_i) in order to identify the residual distribution, and thus identify individuals who are at increased risk for mortality and other forms of morbidity. In our data, a mixture with only two components did not identify the residual group, and a model with more than three components only improves the fit of the predominant distribution and introduces a needlessly complex interpretation.

When the applied problem is not helpful in selecting the number of mixture components, some statistical tools are available. Bayesian approaches for comparing complex hierarchical models in which the number of parameters is not clearly defined include using the deviance information criterion (DIC; Spiegelhalter et al., 2002), but, for mixture models in particular, the DIC is thought to favor overly-complex models (Richardson, in discussion of Spiegelhalter et al., 2002; Celeux et al., 2006). Alternatively, we used posterior prediction intervals to determine if too few classes were chosen, leading to a model that under fits the data (Lynch and Western, 2004).

2.4 Longitudinal Ultrasound Measurements

In our fetal development example analysis, we observe eighteen total ultrasound measurements, which we summarize with three underlying latent factors, $\boldsymbol{\eta}_{i1} = [\eta_{i11}, \eta_{i12}, \eta_{i13}]'$. Four fetal size measurements are obtained at two time points. Specifically, the abdominal circumference is measured at an approximate 15-week ultrasound ($y_{i,1}$) and 24-week ultrasound ($y_{i,5}$). Head circumference ($y_{i,2}$ and $y_{i,6}$), femur length ($y_{i,3}$ and $y_{i,7}$), and biparietal diameter ($y_{i,4}$ and $y_{i,8}$) are also measured at these two time points. We model the correlation between the fetal size outcomes using the latent early fetal growth variable (η_{i11}). Additionally, blood flow resistance is measured by the pulsatility index (PI) and systolic-diastolic ratio (S/D) in the left and right uterine arteries at two time points ($y_{i,9}, \ldots, y_{i,16}$), and the S/D and resistance index (RI) measure resistance in the umbilical artery ($y_{i,17}, y_{i,18}$) around week 24. The highest correlations (all ρ 0.90) were observed among the S/D and PI (or RI for umbilical artery) within a given artery, location, and time, and the lowest correlation between the uterine and umbilical artery measurements (Supplementary Web Table 1). We model the correlation between the uterine artery measurements using a second latent factor (η_{i12}), and consider the umbilical artery S/D and RI to be error-prone realizations of a third latent factor (η_{i13}). We refer to these three latent factors as, collectively, measuring latent early growth restriction.

We formally express the relationship of the longitudinal ultrasound measurements with latent factors and time since the last menstrual period (LMP, denoted W) using the measurement model

$$y_{ij} = \gamma_{0j} + \lambda_j \eta_{i11} + \gamma_{1j} W_i + \gamma_{2j} W_i^2 + \epsilon_{ij}, \quad j = 1, \ldots, 8 \quad (10)$$

$$y_{ij} = \gamma_{0j} + \lambda_j \eta_{i12} + \epsilon_{ij}, \quad j = 9, \ldots, 16 \quad (11)$$

$$y_{ij} = \gamma_{0j} + \lambda_j \eta_{i13} + \epsilon_{ij}, \quad j = 17, 18 \quad (12)$$

where $\epsilon_{ij} \sim N\left(0, \tau_j^{-1}\right), j = 1, \ldots, 18$ and are independent of η_{i11}, η_{i12}, and η_{i13}. The ultrasound size measurements ($j = 1, \ldots, 8$) increase with time, so we allow them to be functions of the reported days from LMP to the ultrasound. Plots of the blood restriction measurements ($j = 9, \ldots, 18$) versus time since LMP showed no change in restriction over

the range of times observed in our study, so we do not include any $\mathbf{\Gamma}$ parameters or time covariates in (11) or (12). Additionally, in (10), we restrict $\lambda_j = \lambda_{j+4}$, $\gamma_{0,j} = \gamma_{0,j+4}$, $\gamma_{1,j} = \gamma_{1,j+4}$, $\gamma_{2,j} = \gamma_{2,j+4}$, and $\tau_j = \tau_{j+4}$, $j = 1, \ldots, 4$, which assumes separate growth curves for the longitudinally-collected HC, FL, AC, and BPD size measurements. Our measurement model assumes that $\mathrm{cov}(y_{ij}, y_{ij'}|S_i = k)$ is given by either $\lambda_j \lambda_{j'} \tau_{11k}^{-1}$ for $j = 1, \ldots, 8$, or $\lambda_j \lambda_{j'} \tau_{12k}^{-1}$ for $j = 9, \ldots, 16$, or $\lambda_j \lambda_{j'} \tau_{13k}^{-1}$ for $j = 17, 18$ with all other covariances being zero. For identifiability, we fix $\gamma_{0,4} = \gamma_{0,16} = \gamma_{0,18} = 0$ and $\lambda_4 = \lambda_{16} = \lambda_{18} = 1$ so that η_{i11}, η_{i12}, and η_{i13} will have location and scale that is commensurate with the biparietal diameter ($y_{i,4}$), S/D in the right uterine artery at week 24 ($y_{i,16}$), and S/D in the umbilical artery ($y_{i,18}$), respectively.

2.5 Cross-sectional birth outcomes

For the outcomes measured at birth, we propose the following form for the measurement model (5) for birth weight (z_{i1}), gestational age at delivery (z_{i2}), and birth weight Z-score (z_{i3}). The correlation among birth weight and gestational age is captured using one latent immaturity variable, $\eta_{i2} = \eta_{i21}$,

$$z_{i1} = \theta_{01} + \theta_{11}\eta_{i21} + \delta_{i1}, \quad \delta_{i1} \sim N\left(0, \tau_{z_1}^{-1}\right) \quad (13)$$

$$z_{i2} = \theta_{02} + \theta_{12}\eta_{i21} + \delta_{i2}, \quad \delta_{i2} \sim N\left(0, \tau_{z_2}^{-1}\right) \quad (14)$$

$$z_{i3} = \sum_{k=1}^{K} \beta_k I(S_i = k) + \delta_{i3}, \quad \delta_{i3} \sim N\left(0, \tau_{z_3}^{-1}\right) \quad (15)$$

with independent error terms δ_{ig}, $g = 1, 2, 3$. We fix $\theta_{02} = 0$ and $\theta_{12} = 1$ in (14) so that η_{i21} has a location and scale that is commensurate with gestational age at delivery. Birth weight Z-scores are calculated by comparing the observed birth weight and gestational age to the expected weight for age from from approximately 6.7 million US births in 1999 and 2000 (Oken et al., 2003). For any given gestational age, the birth weight Z-scores follow a standard normal distribution relative to the reference population (but not necessarily in our PIN group) so that individuals with negative values are believed to have some degree of growth restriction at birth. By definition, birth weight Z-scores are independent of gestational age, so we do not allow z_{i3} to be a function of η_{i21}. The latent factor η_{i21} is thus related to the timing of delivery so that smaller values will be indicative of earlier gestational ages and lower birth weight due to earlier age.

2.6 Prior distributions

To complete a Bayesian specification of the model, prior distributions must be specified for each parameter. In general, we use proper but appropriately vague priors for all parameters to obtain complete conditionals that are of known form. We use conditionally conjugate priors $p\left(\Lambda_j|\tau_{y,j}\right) \sim N\left(\mu_{0,\Lambda_j}, \tau_{y,j}^{-1}\Sigma_{0,\Lambda_j}\right)$, $p\left(\Gamma_j|\tau_{y,j}\right) \sim N\left(\mu_{0,\Gamma_j}, \tau_{y,j}^{-1}\Sigma_{0,\Gamma_j}\right)$, $p\left(\tau_{y,j}\right) \sim \Gamma\left(\frac{c_{y,j}}{2}, \frac{d_{y,j}}{2}\right)$, $p\left(\Theta_g|\tau_{z,g}\right) \sim N\left(\mu_{0,\Theta_g}, \tau_{z,g}^{-1}\Sigma_{0,\Theta_g}\right)$, $p\left(\beta_k|\tau_{z,g}\right) \sim N\left(\mu_{0,\beta_k}, \tau_{z,g}^{-1}\Sigma_{0,\beta_k}\right)$, and $p\left(\tau_{z,g}\right) \sim \Gamma\left(\frac{c_{y,g}}{2}, \frac{d_{y,g}}{2}\right)$. We then choose $\mu_{0,\Lambda_j} = \mu_{0,\Gamma_j} = \mu_{0,\Theta_g} = \mu_{0,\beta_k} = \mathbf{0}$, $\Sigma_{0,\Lambda_j} = \Sigma_{0,\Gamma_j} = \Sigma_{0,\Theta_j} = \Sigma_{0,\beta_k} = 100^2\mathbf{I}$, and $c_{y,j} = d_{y,j} = c_{z,j} = d_{z,j} = .01$ for every j where $\mathbf{0}$ is a conforming vector of zeros and $\mathbf{I}$ is a conforming identity matrix. We also assume $p(\boldsymbol{\omega}) \sim N(\boldsymbol{\mu}_{0,\omega}, \Sigma_{0,\omega}$ and $p(\mathbf{a}) \sim N(\boldsymbol{\mu}_{0,a}, \Sigma_{0,a}$ with $\boldsymbol{\mu}_{0,\omega} = \boldsymbol{\mu}_{0,a} = \mathbf{0}$ and $\Sigma_{0,\omega} = \Sigma_{0,a} = 100^2\mathbf{I}$.

Constraining latent immaturity (η_2) to have the same location and scale as gestational age at delivery simplifies specifying appropriate prior distributions for the mixture component means. For $l = 1, 2, 3$ we assume $p(\mu_{21l}) \sim N(\nu_{21l}, R^2) I(\mu_{21,l-1} < \mu_{21l} < \mu_{21,l+1})$ where $\mu_{210} = -\infty$ and $\mu_{214} = \infty$. Based on previous research by Wilcox (2001) and Gage (2002), we choose $\nu_{211} = 245$, $\nu_{212} = \nu_{213} = 280$, and $R = 10$ so that, *a priori*, we expect that the residual distribution will have a mean of 245 ± 10 (days) and the predominant distribution a mean of 280 ± 10. For the mixture component means of early growth restriction (η_1), we follow the suggestions of Richardson and Green (1997) where, in the absence of previous studies, prior parameters are specified using the observed means and ranges of the three

ultrasound variables y_4, y_{16}, and y_{18}. We use $p(\mu_{1mk}) \sim N\left(\nu_{1mk}, R_m^2\right) (m = 1, 2, 3 \text{ and } k = 1, 2)$ with observed ranges $R_1 = 10$, $R_2 = 2.17$, and $R_3 = 0.32$. The ν_{1mk} ($\nu_{1m1} = \nu_{1m2}$) are the observed means of y_4 (after adjusting for time since LMP), y_{16}, and y_{18}.

We use a hierarchical structure for specifying the prior distribution of each τ_{1mk} and τ_{21l}. Specifically, we allow $p(\tau_{1mk}|b_{0,m}) \sim \Gamma(a_{0,m}, b_{0,m})$ and $p(\tau_{21l}|b_0) \sim \Gamma(a_0, b_0)$ with $b_{0,m} \sim \Gamma(g_{0,m}, h_{0,m})$ and $b_0 \sim \Gamma(g_0, h_0)$. We choose $a_{0,m} = a_0 = 2$, $g_{0,m} = g_0 = 0.2$, $h_{0,m} = 10 * R_m^{-2}$, and $h_0 = 10 * R^{-2}$ where $\Gamma(a, b)$ is the gamma distribution with mean $a \div b$ and variance $a \div b^2$. By choosing $a_{0,m} > 1 > g_{0,m}$ (and $a_0 > 1 > g_0$) we express the general belief that, for each k (and l), the τ_{1mk} (and τ_{21l}) are similar, but we have no information on their absolute size. Finally, we assume that $\boldsymbol{\pi}_s = [\pi_{s1}, \pi_{s2}]'$ and $\boldsymbol{\pi}_t = [\pi_{t1}, \pi_{t2}, \pi_{t3}]'$ follow independent, symmetric Dirichlet distributions, $p(\boldsymbol{\pi}_s) \sim D(d_1, \ldots, d_1)$ and $p(\boldsymbol{\pi}_t) \sim D(d_2, \ldots, d_2)$ and choose $d_1 = d_2 = 1$ to be appropriately vague.

3. Results

To perform the analysis, we used Gibbs sampling (Geman and Geman, 1984), conducted in Matlab, with the complete conditionals given in the Supplementary Web Appendix. We ran five chains from disparate starting values and monitored convergence using the CODA package for R (R Development Core Team, 2004). After removing an initial burn in of 15,000 iterations, all parameters were judged to have converged by the Gelman-Rubin and Geweke diagnostic measures (Geweke, 1991; Gelman and Rubin, 1992). For example, all Gelman-Rubin statistics ($\hat{R}$) were found to be less than 1.01, where $\hat{R} = 1$ at convergence and values less than 1.2 are generally considered sufficient for convergence (Gelman et al., 2004). We used the remaining 35,000 iterations for inference, and summarize our results using posterior means and 95% credible intervals (CrI) for parameters or functions of parameters (e.g. odds ratios) that are of interest.

We applied our methods to 532 subjects in the the Pregnancy, Infection, and Nutrition (PIN) cohort study of prenatal influences on pregnancy outcomes (Savitz et al., 1999). We included all singleton, live born infants who had complete ultrasound, birth weight, gestational age, and covariate information in this analysis. Characteristics of the study subjects are presented in Table 1. As expected, birth weight and gestational age had skewed distributions toward early birth and low weight, respectively.

For early growth restriction, our latent variable mixture model allowed two groups of subjects based on their multiple ultrasound measurements of fetal size and blood flow resistance. In an alternative three class mode, too few subjects belonged to the third group for any meaningful interpretation. For ease of exposition, we refer to these groups as the "normal" and "restricted" groups. On average, a majority of subjects belong to the normal group (posterior mean = 67%, 95% CrI = [59%,74%]). Figure 3 displays a comparison of these two groups for the four measures of fetal size, the PI at various times and locations, and the RI obtained in the umbilical artery at week 24 by Doppler ultrasound. Controlling

for time since LMP, the early growth restricted group had, on average, smaller fetal size measurements and greater resistance to blood flow than the normal group. The restricted group also had significantly larger S/D ratios, indicating greater blood flow resistance, at all times and locations (results not shown). Furthermore, we found evidence of growth restriction at birth in the early growth restricted group. The average birth weight Z-score was −0.34 (95% CrI: [−0.51, −0.18]) in the early growth restricted group and 0.11 (95% CrI: [0.00, 0.22]) in the normal group.

We also examined the association of fetal gender and maternal black race, parity, height, and BMI with group membership. A one inch increase in maternal height was associated with a 0.82 (95% CrI: [0.57,1.00]) fold decrease in the odds of belonging to the early growth restricted group. Black race (posterior odds = 1.33, 95% CrI: [0.77,2.14]), being nulliparous (posterior odds = 1.12, 95% CrI: [0.74,1.65]), and female infant gender (posterior odds = 1.41, 95% CrI: [0.91,2.12]) were not significantly associated with belonging to the early growth restricted group. We also found no linear association between group membership and BMI, but also examined BMI by previously established categories ranging from underweight to obese (WHO Expert Committee, 1995). Children of underweight (BMI < 18.5 kg/m^2) and obese women (BMI   30 kg/m^2) had a relatively high 0.36 (95% CrI: [0.23,0.49]) and 0.38 (95% CrI: [0.27,0.50]), respectively, posterior probability of belonging to the early growth restricted group. Children of women with a BMI in the normal range (BMI $\in$ [18.5, 25) kg/m^2) or moderately obese women (BMI $\in$ [25, 30) kg/m^2) had lower posterior probabilities of 0.31 (95% CrI: [0.23,0.39]) and 0.31 (95% CrI: [0.18,0.45]), respectively.

Particular interest lies in identifying fetuses that have a tendency to be born the most early and with the least weight, which was our motivation in considering latent immaturity classes (***T***). To help understand the type of subjects who belong to this immaturity class, Figure 4 displays the posterior probability of belonging to the residual distribution given birth weight and gestational age. *A priori*, each subject had an equal probability of belonging to the residual distribution, but the posterior distribution of class membership is strongly related to gestational age and birth weight. All subjects born before 224 days (very pre-term) and approximately 30% of subjects born before 259 days (pre-term) are assigned to the immaturity latent class. This corresponds to a 3.7% (95% CrI: [1.2%, 10.4%]) marginal probability of belonging to the residual component. For birth weight, about 90% of subjects with very low birth weight (< 1500g) and 40% of subjects with low birth weight (< 2500 g) belong to the latent class indicative of the residual component.

Finally, we examined the association between our two latent class variables. Using a probit regression model, individuals in the restricted latent class during the second trimester had 3.45 fold greater odds (95% CrI: [0.86,58.9]) of belonging to the residual component of the distribution at birth. The posterior parameter estimates for a_1 and a_2 as defined in (9) are given in Supplementary Web Figure 5 along with the posterior estimates of belonging to the residual distribution conditional on growth restriction during pregnancy. The large posterior odds ratio indicates that belonging to the restricted class is potentially an important predictor of future immaturity, but the credible interval is too wide to make a definitive statement in this dataset.

In our analysis, our primary inference is about latent variables and latent classes, and we make a number of distributional assumptions, so it is important to evaluate the fit of our model. Our latent variable mixture distribution approach fits the observed birth weight and gestational age distributions functions well. Figure 2 contains a histograms of the observed birth weight and gestational age with a line indicating our model fit. Supplementary Web Figure 1 depicts the observed and estimated cumulative distribution functions (CDFs) for

gestational age at delivery while Supplementary Web Figure 2 contains these CDFs for birth weight. Tail areas are magnified in the figures and indicate a good model fit in this area. We also calculated the posterior predictive distributions for observed data and then compared these distributions to our observed data graphically (Lynch and Western, 2004). Supplementary Figure 3 shows the posterior distributions and observed data for the S/D ratio obtained in the right uterine artery at week 15 for 90 subjects. Our model appears to fit the observed data well; posterior predictive plots for other ultrasound measurements, birth outcomes, and subjects also do not indicate a lack of fit.

4. Discussion

We develop a Bayesian approach for analyzing multiple correlated pregnancy outcomes measured during pregnancy by ultrasound and collected routinely at birth using latent variable mixture models. We found evidence in favor of the existence of a latent class of subjects who were more likely to have smaller fetal size measurements and restricted blood flow during the second trimester. These subjects who appeared growth restricted during pregnancy were (1) more likely to have unusually low birth weight for their gestational age (growth restriction at birth), and (2) may be at increased risk for belonging to the residual distribution characterized by earlier gestational age and lower birth weight due to early delivery. Finally, we did not find that height, BMI, black race, gender, or parity were associated with growth restriction during th second trimester of pregnancy.

Our mixture distribution approach is particularly appropriate for formalizing the concept of predominant and residual components of the birth weight and gestational age distributions developed by Wilcox (2001). Using the *T* latent class, we estimated that approximately 3.7% of births lie in the residual distribution, which is in the 2% to 5% range previously estimated for birth weight (Wilcox, 2001). A population with a higher proportion of births in the residual component will be at an increased risk for infant mortality and morbidity, but two populations may differ in their predominant distributions without a corresponding change in outcome (Shan and Ohlsson, 2002; Harder et al., 2007). Common approaches assess mean changes in birth weight or gestational age using linear regression models (e.g. Curry et al., 1998; Glinianaia et al., 2004), but such models assume that shifts in the mean will correspond to an increase in the percentage of the population at particularly high risk. From our perspective, changes in the mean could be due to a shift in the predominant component of the distribution, or a more concerning increase in the proportion of babies born in the residual component, or a combination of the two (Buekens et al., 2000). Our model is able to identify individuals who, based on several longitudinal second trimester ultrasound measurements, are at increased the risk of falling in the residual component at birth.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

References

Albert JH, Chib S. Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association. 1993; 88:669–679.

Buekens P, Notzon F, Kotelchuck M, Wilcox A. Why do Mexican Americans give birth to few low-birth-weight infants? American Journal of Epidemiology. 2000; 152:347–351. [PubMed: 10968379]

Celeux G, Forbes F, Robert C, Titterington D. Deviance information criteria for missing data models. Bayesian Analysis. 2006; 1:651–674.

Curry M, Perrin N, Wall E. Effects of abuse on maternal complications and birth weight in adult and adolescent women. Obstetrics and Gynecology. 1998; 92:530–534. [PubMed: 9764624]

Dempster AP, Laird NM, Rubin DB. Maximum likelihood from incomplete data via the EM algorithm. Journal of the Royal Statistical Society Series B. 1977; 39:1–38.

Diebolt J, Robert CP. Estimation of finite mixture distributions through Bayesian sampling. Journal of the Royal Statistical Society Series B. 1994; 56:363–375.

Dugoff L, Lynch A, Cioffi-Ragan D, Hobbins J, Schultz L, Malone F, D'alton M. First trimester uterine artery doppler abnormalities predict subsequent intrauterine growth restriction. American Journal of Obstetrics and Gynecology. 2005; 193:1208–1212. [PubMed: 16157139]

Dunson DB. Bayesian dynamic modeling of latent trait distributions. Biostatistics. 2006; 7:551–568. [PubMed: 16488893]

Elliott MR, Gallo JJ, Ten Have TR, Bogner HR, Katz IR. Using a Bayesian latet growth curve model to identify trajectories of positive affect and negative events following myocardial infarction. Biostatistics. 2005; 6:119–143. [PubMed: 15618532]

Enstrom B, Gelfand AE, Swamy GK, Miranda ML. Joint modeling of birthweight and gestational age. Joint Statistical Meetings Presentation. 2006

Filly, RA.; Hadlock, FP. Ultrasonography in Obstetrics and Gynecology: Sonographic determination of menstrual age. WB Saunders Co; 2000. p. 146-170.chapter 6

Fokoue E. Mixtures of factor analyzers: An extension with covariates. Journal of Multivariate Analysis. 2005; 95:370–384.

Gage TB. Modeling birthweight and gestational age distributions: Additive vs. multiplicative processes. American Journal of Human Biology. 2002; 14:728–734. [PubMed: 12400033]

Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. Bayesian data analysis. 2 edition. Chapman and Hall; 2004.

Gelman A, Rubin DB. Inference from iterative simulation using multiple sequences. Statistical Science. 1992; 7:457–511.

Geman S, Geman D. Stochastic relaxation, Gibbs distributions and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence. 1984; 6:721–741. [PubMed: 22499653]

Geweke, J. Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In: Bernado, JM.; Berger, JO.; Dawid, AP.; Smith, AFM., editors. Bayesian Statistics. Clarendon Press; 1991. p. 4

Glinianaia S, Rankin J, Bell R, Pless-Mulloli T, Howel D. Particulate air pollution and fetal health a systematic review of the epidemiologic evidence. Epidemiology. 2004; 15:36–45. [PubMed: 14712145]

Harder T, Rodekamp E, Schellong K, Dudenhausen JW, Plagemann A. Birth weight and subsequent risk of type 2 diabetes: A meta-analysis. American Journal of Epidemiology. 2007

Huang G, Bandeen-Roche K. Building an identifiable latent class model with covariate effects on underlying and measured variables. Psychometrika. 2004; 69:5–32.

Hugo E, Odendaal H, Grove D. Evaluation of the use of umbilical artery doppler flow studies and outcome of pregnancies at a secondary hospital. Journal of Maternal-Fetal and Neonatal Medicine. 2007; 20:233–239. [PubMed: 17437225]

Lee S, Xia Y. Maximum likelihood methods in treating outliers and symmetrically heavy-tailed distributions for nonlinear structural equation models with missing data. Psychometrika. 2006; 71:565–585.

Lynch SM, Western B. Bayesian posterior predictive checks for complex models. Sociological Methods and Research. 2004; 32:301–335.

Maulik D. Fetal growth compromise: Definitions, standards, and classification. Clinical Obstetrics and Gynecology. 2006; 49:214–218. [PubMed: 16721101]

Muthén B, Brown C, Masyn K, Jo B, Khoo S, Yang C, Wang C, Kellam S, Carlin J, Liao J. General growth mixture modeling for randomized preventive interventions. Biostatistics. 2002; 3:459–475. [PubMed: 12933592]

Muthén B, Shedden K. Finite mixture modeling with mixture outcomes using the EM algorithm. Biometrics. 1999; 55:463–469. [PubMed: 11318201]

Oken E, Kleinman KP, Rich-Edwards J, Gillman MW. A nearly continuous measure of birth weight for gestational age using a United States national reference. BMC Pediatrics. 2003

Ott W. Accurate gestational dating - revisited. American Journal of Perinatology. 1994; 11:404–408. [PubMed: 7857430]

R Development Core Team. R: A language and environment for statistical computing. 2004.

R Foundation for Statistical Computing. Vienna; Austria: ISBN 3-900051-00-3

Reboussin B, Anthony J. Latent class marginal regression models for modelling youthful drug involvement and its suspected influences. Statistics in Medicine. 2001; 20:623–639. [PubMed: 11223904]

Reboussin B, Liang K, Reboussin D. Estimating equations for a latent transition model with multiple discrete indicators. Biometrics. 1999; 55:839–845. [PubMed: 11315015]

Richardson S, Green P. On Bayesian analysis of mixtures with an unknown number of components. Journal of the Royal Statistical Society Series B. 1997; 59:731–758.

Savitz D, Dole N, Williams J, Thorp J, McDonald T, Carter A, Eucker B. Determinants of participation in an epidemiological study of preterm delivery. Paediatric and Perinatal Epidemiology. 1999; 13:114–125. [PubMed: 9987790]

Shan P, Ohlsson A. Literature Review of low birth weight, including small for gestational age and preterm birth. Toronto Public Health. 2002

Spiegelhalter D, Best N, Carlin B, Van der Linde A. Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society Series B. 2002; 64:583–616.

Stephens M. Dealing with label switching in mixture models. Journal of the Royal Statistical Society Series B. 2000; 62:795–809.

van Dyk D, Meng X. The art of data augmentation. Journal of Computational and Graphical Statistics. 2001; 10:1–50.

WHO Expert Committee. Physical status: The use and interpretation of anthropometry. World Health Organization; 1995. Technical Report 854

Wilcox AJ. On the importance - and the unimportance - of birthweight. International Journal of Epidemiology. 2001; 30:1233–1241. [PubMed: 11821313]
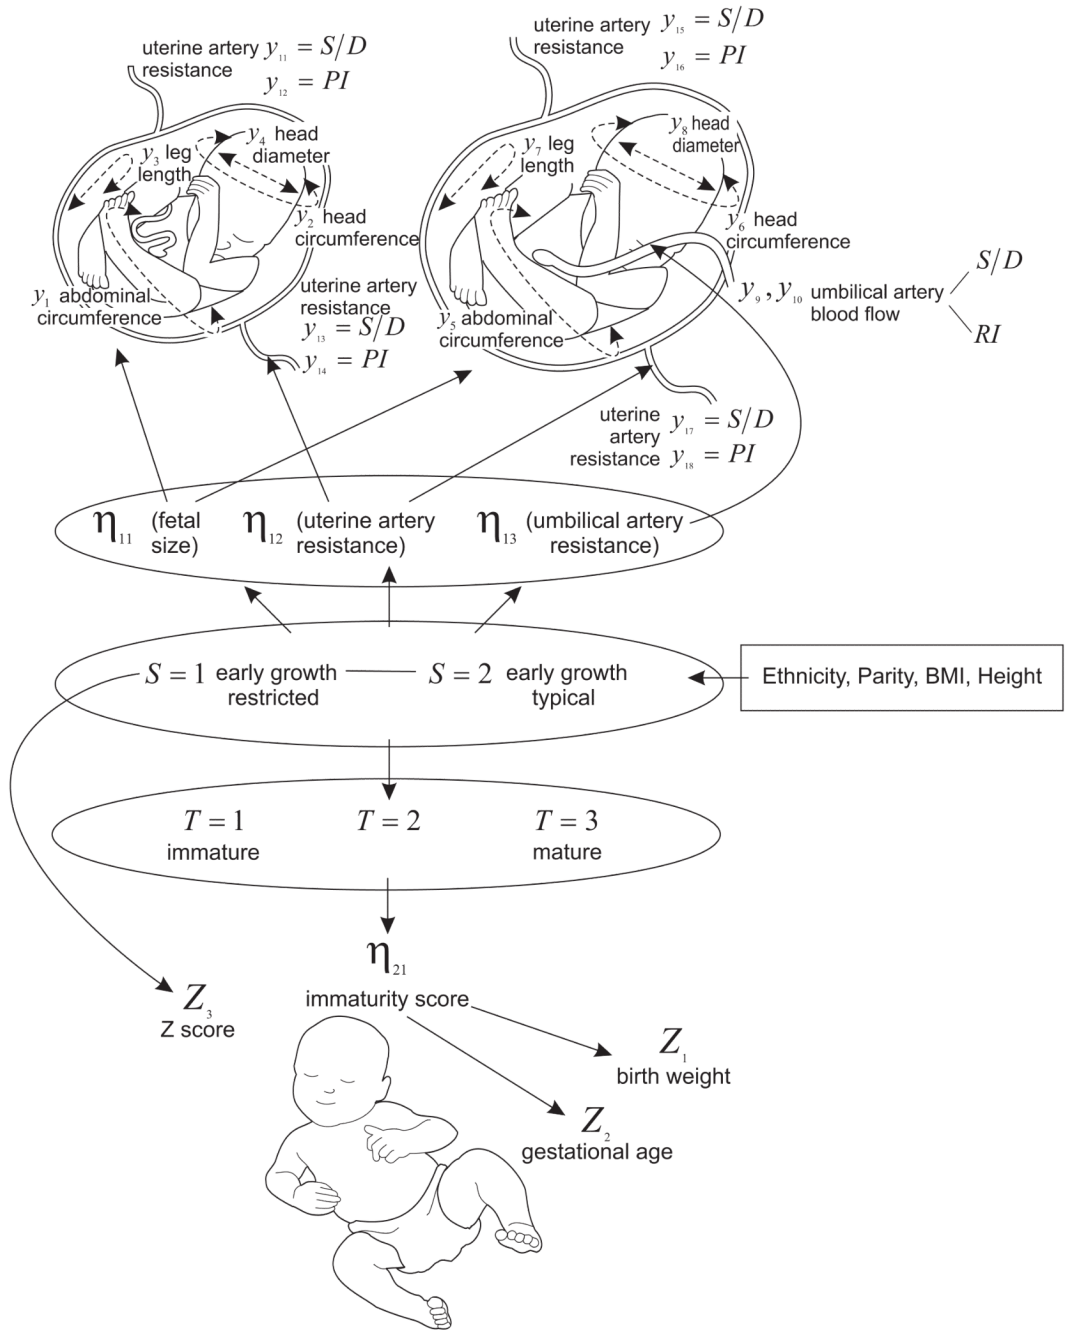
Figure 1.
Path diagram illustrating the dependencies in the proposed fetal growth model. Circles represent latent variables, squares indicate observed variables, and arrows show association.
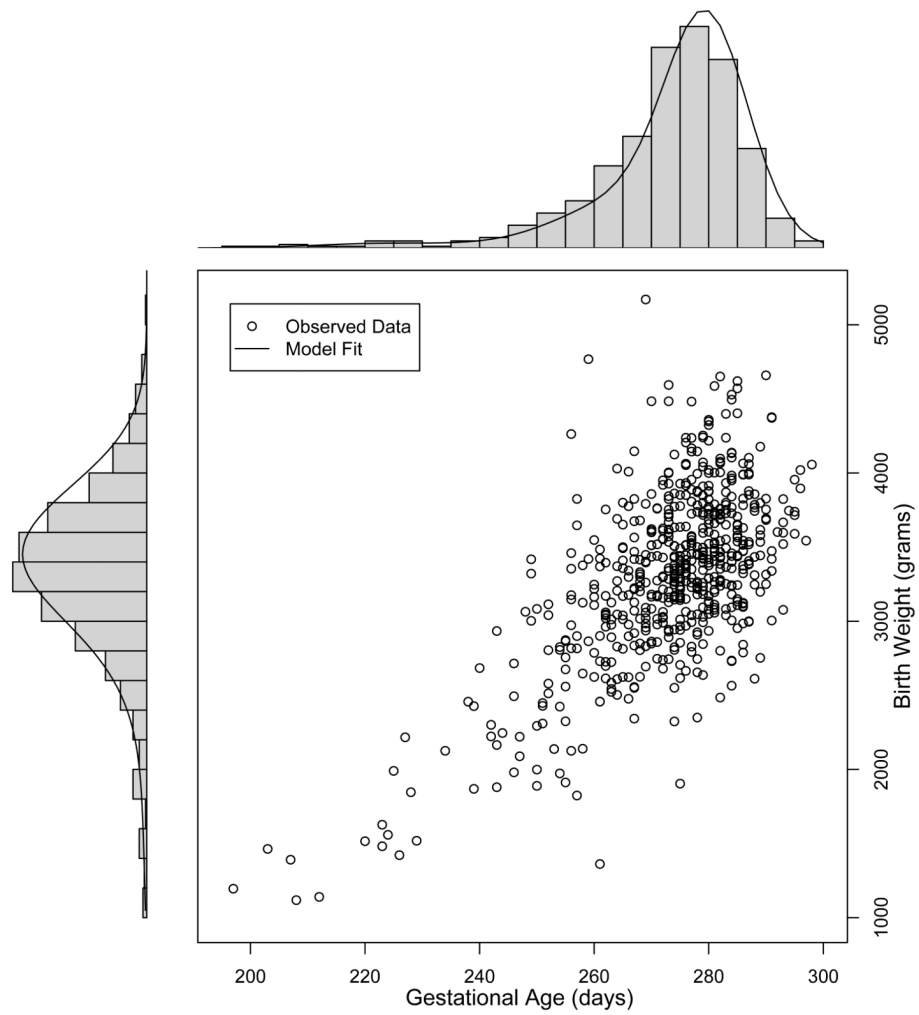
Figure 2.
Joint and marginal distributions of birth weight and gestational age at delivery. The fit of our model is given by the solid line on the histograms.
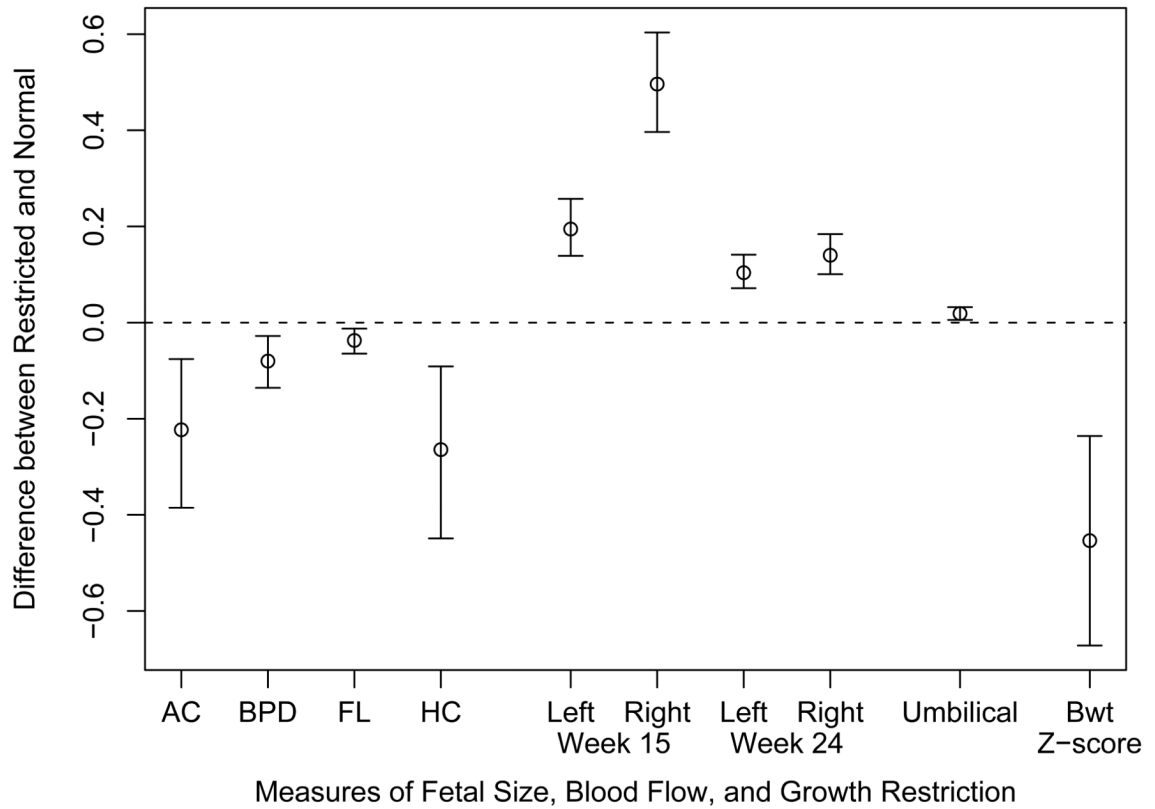
Figure 3.
Differences in ultrasound measurements of fetal size, blood restriction, and birth weight Z-scores for the restricted and normal latent classes. Fetal size measurements include the abdominal circumference (AC), biparietal diameter (BPD), femur length (FL), and head circumference (HC). Difference are estimated using parameter functions $\lambda_1(\mu_{112} - \mu_{111})$, $\lambda_2(\mu_{112} - \mu_{111})$, $\lambda_3(\mu_{112} - \mu_{111})$, and $(\mu_{112} - \mu_{111})$, respectively. The pulsatility index in the left and right uterine arteries at the week 15 and week 24 ultrasound are estimated by $\lambda_{10}(\mu_{122} - \mu_{121})$, $\lambda_{12}(\mu_{122} - \mu_{121})$, $\lambda_{14}(\mu_{122} - \mu_{121})$, and $(\mu_{122} - \mu_{121})$. The resistance index in the umbilical artery is given by $(\mu_{132} - \mu_{131})$.
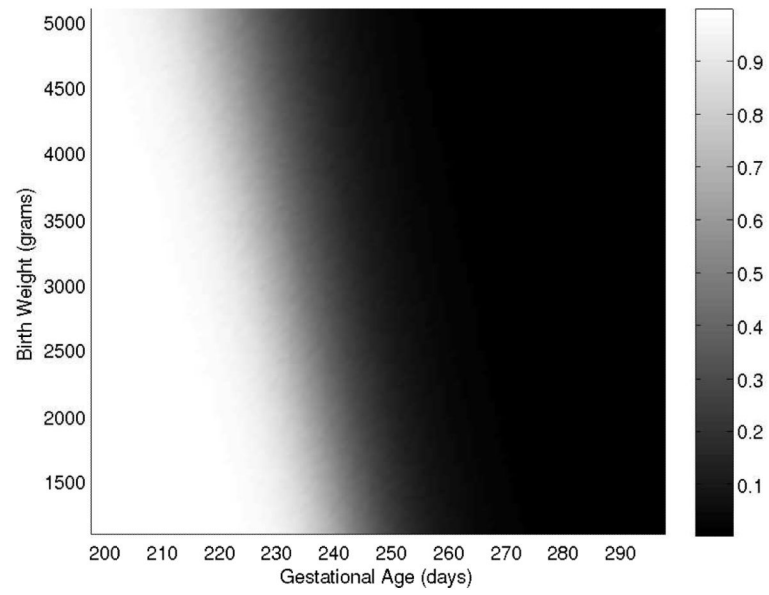
Figure 4.
Posterior probability of being assigned to the residual distribution latent class ($T = 1$) by gestational age and birth weight.

Table 1

Percentiles (continuous) or percentages(dichotomous) of variables summarizing the 532 PIN subjects studied

Characteristic	Percent	$_{10}\mathbf{50}_{90}$
Birth weight (grams)		$_{2624}\mathbf{3370}_{3991}$
Gestational age (days)		$_{257}\mathbf{276}_{287}$
Maternal height (in)		$_{62}\mathbf{65}_{68}$
Maternal BMI (kg/m^2)		$_{19.5}\mathbf{23.5}_{34.4}$
Small for gestational age	8.7%	
Pre-term birth (< 37 weeks)	10.7%	
Very pre-term birth (< 32 weeks)	1.3%	
Low birth weight (< 2500 g)	7.1%	
Very low birth weight (< 1500 g)	1.1%	
Maternal Black Race	17.5%	
Parity 1	55.1%	
Female infant gender	47.8%	