Published in final edited form as:

Biometrics. 2010 June ; 66(2): 444–454. doi:10.1111/j.1541-0420.2009.01309.x.

Semiparametric Bayesian Analysis of Nutritional Epidemiology Data in the Presence of Measurement Error

Samiran Sinha,

Department of Statistics, Texas A&M University, College Station, Texas 77843-3143,
sinha@stat.tamu.edu

Bani K. Mallick,

Department of Statistics, Texas A&M University, College Station, Texas 77843-3143,
bmallick@stat.tamu.edu

Victor Kipnis, and

Biometry Research Group, Division of Cancer Prevention and Control, National Cancer Institute,
Bethesda, Maryland, kipnisv@mail.nih.gov

Raymond J. Carroll

Department of Statistics, Texas A&M University, College Station, Texas 77843-3143,
carroll@stat.tamu.edu

Summary

We propose a semiparametric Bayesian method for handling measurement error in nutritional epidemiological data. Our goal is to estimate nonparametrically the form of association between a disease and exposure variable while the true values of the exposure are never observed. Motivated by nutritional epidemiological data we consider the setting where a surrogate covariate is recorded in the primary data, and a calibration data set contains information on the surrogate variable and repeated measurements of an unbiased instrumental variable of the true exposure. We develop a flexible Bayesian method where not only is the relationship between the disease and exposure variable treated semiparametrically, but also the relationship between the surrogate and the true exposure is modeled semiparametrically. The two nonparametric functions are modeled simultaneously via B-splines. In addition, we model the distribution of the exposure variable as a Dirichlet process mixture of normal distributions, thus making its modeling essentially nonparametric and placing this work into the context of functional measurement error modeling. We apply our method to the NIH-AARP Diet and Health Study and examine its performance in a simulation study.

Keywords

B-splines; Dirichlet process prior; Gibbs sampling; Measurement error; Metropolis-Hastings algorithm; Partly linear model

1 Introduction

We consider the logistic regression of disease status Y on covariates $(X, \mathbf{Z})$, where X is unobservable and can only be measured with error and the goal is to understand its effect on Y . In doing so we assume that the logit of the disease probability is a linear function of $\mathbf{Z}$ and possibly a nonlinear function of X but its exact form is unknown, thus yielding a

partially linear logistic model. Although we work with a logistic model, the method is applicable to any distributional model of partially linear form.

The commonly used classical measurement error model assumes that instead of observing X , for every subject one observes, possibly after data transformation, a surrogate W which is unbiased for X involving purely random error with constant variance (Carroll et al., 2006). However, in large nutritional epidemiological studies, dietary intake for all individuals is usually measured by a food frequency questionnaire (FFQ), which we denote here by Q . It has become well appreciated in the literature that FFQs have substantial measurement error, both random and systematic, therefore violating the classical measurement error assumptions. Our motivating example is the NIH-AARP Diet and Health Study (<http://dietandhealth.cancer.gov/>), the details of which can be found in Schatzkin et al. (2001). Important to note that any case-control study for finding diet-cancer association is subject to differential recall bias between cases and controls. Secondly a homogeneous population with a narrow range of fat intake usually fails to find any association between fat intake and breast cancer. To circumvent these issues, in the NIH-AARP Diet and Health Study a large and diverse population was targeted where diet was assessed prior to diagnosis. In this study, the initial cohort of 617,119 men and women who responded to a FFQ in 1995–1996 has been followed for the evaluation of possible diet-cancer associations. The latest database till the year 2003 contained information on 27 types of cancer, and breast cancer is one of them, which is considered as the response variable in this paper. To adjust for FFQ measurement error in estimated relationships, the NIH-AARP study includes a so-called calibration sub-study with approximately 2000 men and women who, in addition to the FFQ, were administered two non-consecutive 24-hour dietary recalls (24hr) denoted by W . Following the study design, we assume W follows the classical measurement error model. Thus the study design consists of the binary response Y (occurrence or non-occurrence of invasive breast cancer during follow up), covariates without error Z , true unobserved main exposure X , observed exposure Q that measures X with both bias and random error. In addition, a calibration sub-study includes both Q and W along with Z .

To handle the substantial random and systematic measurement error in the FFQ in a semi-parametric fashion, we assume that conditional on (X, Z) , the surrogate variable Q has a partially linear model, where the effect of Z is linear and the effect of X on Q is still unknown, but assumed to be a smooth and monotone function. In fact for our data example, we found that a linear regression between Q and the average of W is not sufficient to explain the nature of association between these two variables, which is an indication of a possible nonlinear association between Q and X . Although a simple logistic regression of the occurrence of invasive breast cancer (Y) on the percentage of non-alcohol energy from total fat measured via FFQ (Q) revealed a significant association, we want to measure how the risk changes with the true value of the percentage of non-alcohol energy from total fat (X) taking into account that X is unobserved and Q contains substantial measurement error. Therefore, we will model the effect of X on the logit of the success probability of Y via splines, and the effect of X on Q via monotone splines. Since we will develop a likelihood based inference for our models, we need to specify the distribution of the latent variable X . Since the histogram of the average of W obtained from the calibration data and that of Q from the cohort study (see Web Figure 1) do not show strong evidence for the normal distribution assumption for X , we model the distribution of X nonparametrically via an infinite mixture of normal distributions. In fact the pattern of food intake is likely to vary as

Supplementary Materials

Web Appendix, and Figures referenced in Sections 1, 3, and 4 are available under the Paper Information link at the Biometrics website <http://www.biometrics.tibs.org>.

the cohort members have a diverse background and are from 6 different US cities and two metropolitan areas with large minority populations.

Before describing the novelty of our disease model and the calibration model and our approach to handling them, we first point out that the regression calibration approach is not applicable in our context. In parametric linear logistic model for the disease probability i.e.,

$$\text{pr}(Y=1|X, \mathbf{Z})=H\{X\xi+\mathbf{Z}^T\boldsymbol{\zeta}_{\text{risk}}\}, \quad (1)$$

where $H(u) = 1/\{1 + \exp(-u)\}$ is the logistic distribution function, it has become common to apply regression calibration (RC) adjustment for measurement error, where one regresses W on Q in the calibration sub-study, then replaces X in (1) by the predictions from this regression. However, this method is well-known to be undesirable in semiparametric models such as the partially linear logistic model, and indeed in our simulations it performs quite poorly. The reason is that RC assumes that in the induced observed data model $\text{pr}(Y=1|W, \mathbf{Z})$ the effect of W is conferred only through $E(X/W, \mathbf{Z})$. However, if the effect of X in the logit of $\text{pr}(Y=1|X, \mathbf{Z})$ is nonlinear, the actual observed data model may be quite far from the assumed induced observed data model. An example of this is known in ordinary nonparametric regression, where if $E(Y/X) = \sin(2X)$, $X, U \sim \text{Normal}(0, 1)$, and $W = X + U$, then the approximation $E(Y/W) \approx \sin\{2E(X/W)\}$ is systematically biased and out of phase with the true regression function.

Carroll and Hall (1988), Fan and Truong (1993) and Delaigle and Hall (2008) considered deconvolution method to deal with classical measurement error in X for a nonparametric regression of Y on X . However, such methods did not consider systematically biased surrogate such as FFQ. Also the methods were not designed to handle the partially linear logistic model.

There are some related Bayesian methodologies, but none of them handle the generality of the problem we confront. Berry et al. (2002) used smoothing splines and regression splines in the classical measurement error problem to a linear model set up, but not to the important case of binary data. Carroll et al. (2004) used Bayesian spline-based regression when an instrument is available for all study participants. In addition, both papers assumed that the unknown X is normally distributed. Mallick and Gelfand (1996) considered covariate measurement error in the generalized linear model with unknown link function, where the distribution of X was modeled via a multivariate normal distribution. Müller and Roeder (1997) considered the multivariate normal mixture of Dirichlet process prior for handling covariate measurement error in case-control studies. Bayesian nonparametric regression approaches without measurement error in the covariate for binary data have been considered by Wood and Kohn (1998), Wood et al. (2002) and Holmes and Mallick (2003), among others. In summary, all the above mentioned papers considered either (a) nonparametric regression without any measurement error or (b) measurement error in covariates while the regression model is parametrically specified. None of these papers allowed for the partially linear model for the response variable Y , nor did they even begin to address partially linear calibration model to handle systematic bias in FFQ. Johnson et al. (2007) considered a problem similar to ours, although the data structures are slightly different, since in their example the FFQ is replicated. The important differences are that in place of the semiparametric risk model for Y given $(X, \mathbf{Z})$, they used the parametric model (1), and in place of the partially linear calibration model for Q given $(X, \mathbf{Z})$, they used a linear model. The important similarity is that they, like us, used a mixture of Dirichlet process (DP) model for the distribution of the latent variable.

In summary the three novel features of our approach are the following. First, we consider a semiparametric logistic model with a nonparametric component subject to measurement error. Second, we allow for the fact that in actual epidemiological practice, the vast majority of the data only have a systematically biased measure of the true risk covariate, and we handle this feature via a semiparametric model with a monotone nonparametric component, the monotonicity being natural in the scientific context. Although the idea of systematic and random bias in Q has appeared in many papers in nutritional epidemiology (Kipnis et al., 2001, 2003), our consideration of a semiparametric model for the systematic component of the bias is new. Third, we model the distribution of the unobserved covariate nonparametrically via a Dirichlet process mixture of normal distributions. While the use of such models in general is not new (Johnson et al., 2007), using the idea in a semiparametric context has not previously been investigated. To the best of our knowledge then, this is the first work in the semiparametric measurement error field where two smooth nonparametric functions are estimated simultaneously, one in the exposure-response association, and the other in the association between the surrogate variable and the true exposure, while at the same time treating the distribution of the true exposure variable essentially nonparametrically. Moreover, the estimation of two nonparametric functions, where one is dependent on the other, is not an easy task, especially for binary regression when the covariate distribution is unknown. The simulation study and data analysis show the importance of such flexible models.

An outline of the paper is as follows. The model and assumptions are described in Sections 2, while the method of estimation is described in Section 3. Section 4 contains the data analysis of the NIH-AARP Study. A simulation study is described in Section 5. Section 6 contains concluding remarks.

2 Model and Assumptions

2.1 Outline

In this section, we give the model structure for the observed data (Y, Q, Z, W) given the latent long term diet X (Section 2.2), the model structure for X (Section 2.3), and the models for the semiparametric functions that are in the model (Section 2.4).

2.2 Model Structure for Observed Data

Let Y be the binary response variable, X the covariate of interest, Q the surrogate variable for X , and Z a vector of error-free covariates. The primary data consist of (Y_i, Q_i, Z_i) for $i = 1, \dots, n$. In order to obtain information about the relationship of Q_i and X_i we assume that there is an external calibration data where we observe Q, Z , and repeated measurements of some unbiased surrogate variable W . Therefore, the calibration data are (Q_j, Z_j, W_{ij}) for $j = 1, \dots, J$, and $i = n + 1, \dots, N$, where $N = n + m$. The basic risk model is

$$\text{pr}(Y=1|X, Z) = H\{\theta_{\text{risk}}(X) + Z^T \zeta_{\text{risk}}\}, \quad (2)$$

where $\theta_{\text{risk}}(\cdot)$ is an unknown function, and Q is related to (X, Z) via

$$Q = \theta_{\text{cal}}(X) + Z^T \zeta_{\text{cal}} + U_Q, \quad (3)$$

where $\theta_{\text{cal}}(\cdot)$ is an unknown smooth monotone function. We further assume that the within-person random measurement error U_{Q_i} is nondifferential, and we make the standard assumption that U_{Q_i} follows a normal distribution with mean zero and variance σ_Q^2 . Let Δ_i be

an indicator variable corresponding to subject i , which takes on value 1 if the i^{th} subject belongs to the calibration study and 0 otherwise. Therefore, when $\Delta_i = 1$ we observe $(Q_i, Z_i, W_{ij}, j = 1, \dots, J)$, and when $\Delta_i = 0$ we observe (Y_i, Q_i, Z_i) for the i^{th} subject. For identifiability, we require that the distribution of Q given (X, Z) is the same in both primary and calibration data, e.g. the latter is a randomly chosen subset of the primary data.

For the unbiased surrogate variable, there are $j = 1, \dots, J$ replicates and we assume $W_{ij} = X_i + U_{Wij}$. Furthermore, we assume that conditional on X , U_{Wij} are independent and identically distributed $\text{Normal}(0, \sigma_w^2)$, and conditional on (X, Z) , U_{Qi} is assumed to be independent of U_{Wij} . Note that the assumptions of normality and independence for the errors U_{Qi} and U_{Wij} can be relaxed to any bivariate parametric model.

2.3 Distribution of the Unobserved Covariate X

A likelihood based approach requires a distribution for the unobserved covariate X . Since the misspecification of this distribution may result in biased estimates of the quantities of interest, we model the distribution in a flexible semiparametric fashion. We assume that conditional on Z_i , μ_i , and σ_i^2 , $X_i \sim \text{Normal}(\mu_i + Z_i^T \zeta_x, \sigma_i^2)$, and assume that a priori the parameters (μ_i, σ_i^2) come from a distribution G , that means, $\varphi_i^T = (\mu_i, \sigma_i^2) \sim G$, and a priori $G \sim DP(\alpha_0 G_0)$, where G_0 is the base probability measure defined on (χ, β) . Here $\chi = (-\infty, \infty) \times (0, \infty)$ and β is the σ -algebra generated by χ . Under G_0 , $[\mu | \sigma^2, \tau] \sim \text{Normal}(0, \tau \sigma^2)$, $\sigma^2 \sim IG(a_\sigma, b_\sigma)$, and the hyperparameter $\tau \sim IG(a_\tau, b_\tau)$, where $IG(a_\tau, b_\tau)$ denotes the inverse-gamma distribution with mean $1/\{b_\tau(a_\tau - 1)\}$. Since for any set A , $G(A)$ has mean $G_0(A)$ and variance $G_0(A)\{1 - G_0(A)\}/(\alpha_0 + 1)$, the parameter α_0 plays an important role in determining the concentration of the random process G around the base probability measure G_0 . The stick-breaking representation of the Dirichlet process indicates that G is almost surely a discrete probability measure with infinite many mass points, which in turn implies that the distribution of X is an infinite mixture of normal distributions, and the parameters of the component distribution come from the base probability measure G_0 . In addition, the

mixing probabilities π_k 's are obtained via $\pi_k = \gamma_k \prod_{j=1}^{k-1} (1 - \gamma_j)$, where $\gamma_k = \text{Beta}(1, \alpha_0)$, and $\sum_{k=1}^{\infty} \pi_k = 1$.

2.4 Semiparametric Models for $\theta_{\text{risk}}(\bullet)$ and $\theta_{\text{cal}}(\bullet)$

We wish to estimate $\theta_{\text{risk}}(X)$ in the interval $[a, b]$ and for this purpose we use B-splines. Let $B_{\text{risk}}^T(X) = \{B_{1\text{risk}}(X), \dots, B_{p\text{risk}}(X)\}$ be a set of cubic spline basis functions for k_{risk} fixed knot points. With $p = k_{\text{risk}} + 4$, our model is

$$\theta_{\text{risk}}(X) = \sum_{j=1}^p B_{j\text{risk}}(X) \beta_j.$$

Also based on k_{cal} fixed knot points, with $q = k_{\text{cal}} + 4$, we express $\theta_{\text{cal}}(\cdot)$ as

$$\theta_{\text{cal}}(X) = \sum_{j=1}^q B_{j\text{cal}}(X) \alpha_j,$$

where $B_{\text{cal}}^T(X) = \{B_{1\text{cal}}(X), \dots, B_{q\text{cal}}(X)\}$ is also a set of cubic spline basis. The use of cubic splines implies that the nonparametric functions have continuous second-derivatives. The chosen knot points for the two nonparametric functions could be the same, and consequently the spline basis functions would be the same. In order to estimate $\theta_{\text{risk}}(\bullet)$ properly, $\theta_{\text{cal}}(\bullet)$

should be a monotonic nondecreasing function of its argument, which is a natural assumption regarding the conditional mean of Q given X . Because the B-spline basis functions are nonnegative, under the monotonicity constraint for $\theta_{\text{cal}}(X)$, the first derivative $\theta_{\text{cal}}^{(1)}(X)$ is non-negative if the coefficients α_j 's are nondecreasing, i.e., $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_q$. This property of B-splines has been previously used by Leitenstorfer and Tutz (2007) among many others. Also we chose a moderately large number of knot points which will make the inference insensitive to the location of knots (Ruppert, 2002). Next we describe the method of estimation within the Bayesian paradigm.

3 Method of Estimation

Let $\beta^T = (\beta_1, \dots, \beta_p)$, $\alpha^T = (\alpha_1, \dots, \alpha_q)$, $\phi = (\phi_1, \dots, \phi_{n+m})$, and define $\theta = (\alpha^T, \beta^T, \zeta_{\text{risk}}^T, \zeta_{\text{cal}}^T, \zeta_x^T, \sigma_Q^2, \sigma_w^2)$. The observed data likelihood function is

$$L_o = \int \prod_{i=1}^{n+m} \left\{ f^{1-\Delta_i}(Y_i, Q_i | Z_i, \theta, \phi) f^{\Delta_i}(Q_i, W_{ij}, j=1, \dots, J | Z_i, \theta, \phi) \right\} dG(\phi), \text{ where}$$

$$\begin{aligned} f(Y_i, Q_i | Z_i, \theta, \phi) &= \int f(Y_i, Q_i, X_i | Z_i, \theta, \phi) dX_i \\ &= \int \text{pr}^{Y_i}(Y=1 | X_i, Z_i, \theta) \text{pr}^{1-Y_i}(Y=0 | X_i, Z_i, \theta) f(Q_i | X_i, Z_i, \theta) f(X_i | Z_i, \theta, \phi) dX_i \\ &= \int \frac{\exp \{Y_i[\theta_{\text{risk}}(X_i) + Z_i^T \zeta_{\text{risk}}]\}}{1 + \exp \{\theta_{\text{risk}}(X_i) + Z_i^T \zeta_{\text{risk}}\}} \frac{1}{(\sigma_Q^2)^{1/2}} \exp \left[-\frac{1}{2\sigma_Q^2} \{Q_i - \theta_{\text{cal}}(X_i) - Z_i^T \zeta_{\text{cal}}\}^2 \right] \\ &\quad \times \frac{1}{(\sigma_i^2)^{1/2}} \exp \left\{ -\frac{1}{2\sigma_i^2} (X_i - \mu_i - Z_i^T \zeta_x)^2 \right\} dX_i, \end{aligned}$$

$$\begin{aligned} f(Q_i, W_{i1}, \dots, W_{ij} | Z_i, \theta, \phi) &= \int f(Q_i, W_{i1}, \dots, W_{ij}, X_i | Z_i, \theta, \phi) dX_i \\ &= \int f(Q_i | X_i, Z_i, \theta) f(W_{i1}, \dots, W_{ij} | X_i, Z_i, \theta) f(X_i | Z_i, \theta, \phi) dX_i \\ &= \int \frac{1}{(\sigma_Q^2)^{1/2}} \exp \left[-\frac{1}{2\sigma_Q^2} \{Q_i - \theta_{\text{cal}}(X_i) - Z_i^T \zeta_{\text{cal}}\}^2 \right] \times \frac{1}{(\sigma_w^2)^{J/2}} \exp \left\{ -\frac{1}{2\sigma_w^2} \sum_{j=1}^J (W_{ij} - X_i)^2 \right\} \\ &\quad \times \frac{1}{(\sigma_i^2)^{1/2}} \exp \left\{ -\frac{1}{2\sigma_i^2} (X_i - \mu_i - Z_i^T \zeta_x)^2 \right\} dX_i. \end{aligned}$$

Before describing prior specification and posterior inference, we would like to mention that one could adopt a complete nonparametric method by modeling the joint distribution of X , Q and Z for $Y = 0$ and $Y = 1$ using dependent Dirichlet process prior. However, that method would involve multidimensional Dirichlet process integrals which is quite challenging in terms of computation. Therefore, to reduce the dimensionality and computational complexity of the problem we model the distribution of X conditional on Z where Z is always observed without any error.

3.1 Specification of Priors

In order to avoid oversmoothing due to large number of knot points in the spline models, following Ruppert et al. (2003) we use the following prior distribution for α and β ,

$$\beta, \pi(\alpha | \delta_\alpha) \propto \exp\{-(\delta_\alpha/2) \sum_{j=3}^q (\Delta^2 \alpha_j)^2 - (\delta_\alpha/2V) \alpha^T \alpha\}, \text{ and } \pi(\beta | \delta_\beta) \propto \exp\{-(\delta_\beta/2) \sum_{j=3}^p (\Delta^2 \beta_j)^2 - (\delta_\beta/2V) \beta^T \beta\}$$

, where δ_α and δ_β are the two penalty parameters corresponding to α and β . Note that the priors are proper for any choice of $V > 0$, and we set $V = 10^8$. Also, Δ represents the first

order difference operator, i.e., $\Delta\alpha_j = \alpha_j - \alpha_{j-1}$, and $\Delta^2\alpha_j = \alpha_j - 2\alpha_{j-1} + \alpha_{j-2}$. We assume $\delta_\alpha \sim \text{Gamma}(a_\alpha, b_\alpha)$ and $\delta_\beta \sim \text{Gamma}(a_\beta, b_\beta)$. For σ_Q^2, σ_W^2 , we use $IG(a_Q, b_Q)$ and $IG(a_W, b_W)$ priors respectively. For each component of $\zeta_{\text{cal}}, \zeta_{\text{risk}}, \zeta_x$ we use a $\text{Normal}(0, \sigma_\zeta^2)$ prior.

As mentioned earlier, for $(\mu_i, \sigma_i^2) \sim G$ and a priori $G \sim DP(\alpha_0 G_0)$. The Dirichlet process prior involves two parameters G_0 and α_0 . Note that the tuning parameter of the Dirichlet process prior α_0 has a dual role. On the one hand it determines how concentrated G is around G_0 , and on the other hand it determines the number of clusters, i.e., the number of distinct mass points of G . Hence the choice of α_0 is an important one. Following Escobar and West (1995) one may put a prior on α_0 . However, we will estimate α_0 empirically following an existing characterization of its maximum likelihood estimator (McAuliffe et al, 2006). In summary, for our proposed method one needs to specify the following quantities: $V, a_\alpha, b_\alpha, a_\beta, b_\beta, a_Q, b_Q, a_W, b_W, \sigma_\zeta^2$ and G_0 .

3.2 Posterior Inference

Our primary goal is to estimate $\zeta_{\text{risk}}, \zeta_{\text{cal}}, \beta, \alpha, \sigma_Q^2$, and σ_W^2 . Posterior means are used as the parameter estimates. As we can see the posterior means cannot be calculated analytically, we propose an efficient posterior sampling algorithm for simulating random numbers from the posterior distributions, and the details are given in the Web Appendix. At each MCMC iteration we estimate α_0 empirically by solving the following equation:

$$\sum_{r=1}^N \frac{\alpha_0}{\alpha_0 + r - 1} = E(k|X_i, i=1, \dots, N, \alpha_0), \quad (4)$$

where the right hand side of (4) is obtained by taking average of k , the number of distinct ϕ_i 's, of the last 30 Gibbs sampling iterations, i.e., we take

$$E(k|X_i, i=1, \dots, N, \alpha_0) \approx \sum_{l=1}^{30} k_l / 30.$$

Note, as $\phi_1, \dots, \phi_N$ are independent and identically distributed observations from G , and G has a $DP(\alpha_0 G_0)$ prior, the posterior distribution of G given $\phi_1, \dots, \phi_N$ is again a

$DP(\alpha_0^* G_0^*)$, where $\alpha^* = \alpha_0 + N$ and $G_0^* = (\alpha_0 + N)^{-1}(\alpha_0 G_0 + \sum_{i=1}^N \delta_{\phi_i})$. Thus for a future realization ϕ_{N+1} , $\text{pr}(\phi_{N+1} \in A | \phi_1, \dots, \phi_N) = E\{G(A) | G \sim DP(\alpha_0^* G_0^*)\} = G_0^*(A)$. Therefore, the predictive distribution of $\phi_{N+1} | \phi_1, \dots, \phi_N$ is G_0^* .

4 Application to the NIH-AARP Diet and Health Data

4.1 Background

In this illustration of our methodology, we consider the association between percentage of non-alcohol energy from total fat and the risk of invasive breast cancer (Y) in the NIH-AARP cohort. Here are some details about this study.

During 1995–1996 3.5 million baseline questionnaires were mailed to current members of the AARP (formally the American Association of Retired Persons), aged 50–71 years, and 567,169 members completed them satisfactorily (Schatzkin et al, 2001). Of these, the investigators excluded an additional withdrawal ($n=1$), duplicate records ($n=179$), subjects who moved out of the eight states included in the study before returning the baseline questionnaire ($n=321$), or were found to have died before study entry ($n=261$). After these

exclusions, the NIH-AARP baseline cohort included 566,407 persons that we considered in our analysis. Thiebaut et al. (2007) analyzed this cohort with non-alcohol energy from total fat as well as fatty acid as the main exposure variables, in addition to several potential confounders. In our illustration, we use non-alcohol energy from total fat as the main exposure and body mass index (BMI) measured in the metric unit kg/m^2 as the error-free covariate Z . We considered only postmenopausal women who did not report a personal history of any cancer, did not develop breast cancer in situ during the follow-up and whose BMI values were not missing. Of the remaining 144,366 subjects we excluded those women who reported extreme values (i.e., more than two interquartile range above the 75th percentile or below the 25th percentile on the logarithmic scale) for total fat, energy, or percentage of non-alcohol energy from total fat, thus leaving us with 142,364 subjects, of whom 2,724 women developed invasive breast cancer during the follow-up period. The average follow-up time for these subjects was 4.52 years. Based on the daily alcohol consumption in grams, total fat intake in grams, and total energy intake in Kcals, the percentage of non-alcohol energy from total fat was calculated by using the following formula

$$\text{Percentage of non-alcohol energy from fat} = \frac{\text{energy from total fat}}{\text{non-alcohol energy}} \times 100,$$

where energy from total fat was obtained by multiplying total daily consumption of fat by 9 Kcals, and non-alcohol energy was calculated by subtracting energy from alcohol from total energy. Energy from alcohol was computed by multiplying alcohol consumption measured in grams by 7 Kcals. The percentage of non-alcohol energy from total fat measured from the FFQ was treated as the surrogate measurement Q of X .

The NIH-AARP calibration sub-study was designed to calibrate the FFQ used in the main study with two non-consecutive 24-hour recall telephone interviews as the unbiased surrogate measure. While there is concern that 24-hour recalls may not be actually unbiased for true intake, these are the data we have and their use can be justified as a so-called alloyed gold standard. Out of the 1953 subjects of both genders we considered only the 919 women who completed two 24-hour recall interviews and whose BMI values were not missing. Note that these women were all postmenopausal and cancer free at the start of the study and they were a random sample from the main study population. The percentage of non-alcohol energy from total fat obtained from the two 24-hour recalls is considered as the unbiased surrogate, W . We defined the response variable Y as one if a person develops invasive breast cancer during the follow-up period and zero otherwise. Before the analysis we transformed BMI by dividing by 10, and X was the logarithm of the percentage of non-alcohol energy from total fat. In the preliminary analysis we fit model (1) without any adjustment for measurement error using the primary data. The result showed that X was significantly positively associated with the risk of invasive breast cancer with an odds ratio 1.176 with a 95% confidence interval (1.016, 1.363) and p -value= 0.030. Also, we found that BMI was positively associated with the disease with an odds ratio 1.069 with a 95% confidence interval (1.008, 1.135) and p -value= 0.026.

4.2 Semiparametric Results

First, we used our semiparametric Bayesian (SPB) approach. We experimented with 15, 11, 9, and 6 knot points, and the results were fairly insensitive to the number of knot points. Therefore here we discuss the results for 6 knot points: (2.9, 3.0, 3.2, 3.4, 3.6, 3.80), and we used cubic B-splines, so there were $p = 10$ spline basis functions. We used the same knot points for $\theta_{\text{risk}}(\cdot)$ and $\theta_{\text{cal}}(\cdot)$. For δ_α and δ_β we used Gamma(0.001, 1000) priors. We used

$IG(25, 0.25)$ prior for σ_Q^2 and σ_w^2 , and $Normal(0, 10^2)$ prior for ζ_{risk} , ζ_{cal} and ζ_x . For the distribution G_0 of the Dirichlet process prior, the priors chosen were $\tau \sim IG(2.1, 2.0)$ and $\sigma^2 \sim IG(2.1, 0.91)$. The prior parameters were chosen in such a way that the priors cover a wide range of values for the parameter of interest. We ran the Gibbs sampling for 20,000 iterations and we discarded the first 5000 iterations as burn-in samples. The choice of the initial parameters is an important step for proper and timely convergence of the MCMC method, and in the next section we discuss this issue in detail. The solid lines of Figure 1 show the estimated $\theta_{risk}(\cdot)$ along with a 95% credible interval. Overall, the method shows that the risk of having breast cancer increases with the logarithm of the percentage of non-alcohol energy from total fat. Figure 2 shows the estimated $\theta_{cal}(\cdot)$ function along with 95% credible interval. This figure also shows the scatter plot of $(Q - \hat{\zeta}_{cal}Z, \bar{W})$ from the calibration data. Both figures clearly indicate that there is some sort of nonlinearity in both $\theta_{risk}(\cdot)$ and $\theta_{cal}(\cdot)$. Table 1 presents the posterior mean, posterior standard deviation, and 95% credible intervals for ζ_{risk} , ζ_{cal} , ζ_x , σ_Q^2 and σ_w^2 .

Using the Dirichlet process prior we found that the average number of mixing components is 2.11 with a posterior standard deviation 0.33, indicating that the true distribution of X given Z is not a normal distribution but a mixture of normals which is also evident from the histogram of a MCMC sample of X values (Figure 3). Figure 4 shows the histogram of α_0 , the tuning parameter of the Dirichlet process prior. Web Figure 2 shows a plot of the number of clusters of the Dirichlet process prior in the MCMC samples. Following a referee's comment we also have added a summary of cluster sizes using boxplot based on the MCMC sample (Web Figure 3) when the number of clusters are 2, 3, or 4. Web Figure 4 shows an estimated $G_0^* = E(G|data)$ based on the MCMC sample of ϕ_{N+1} .

The null hypothesis of interest is that there is no risk of fat on breast cancer: under this null hypothesis θ_{risk} should be constant and thus it is equivalent to test $\beta_1 = \dots = \beta_p$ as

$\sum_{j=1}^p B_{jrisk}(X) = 1$. Suppose we consider M MCMC samples taking every 300th observation after the burn-in period. Then the test statistic for testing the null hypothesis is $F = [(M - p + 1) / \{(M - 1)(p - 1)\}] T^2$, where

$T^2 = M \bar{D}^T S_D^{-1} \bar{D}$, $\bar{D}^T = (\bar{D}_1, \dots, \bar{D}_{p-1})$, $\bar{D}_k = \frac{1}{M} \sum_{j=1}^M D_{kj}$, $k=1, \dots, (p-1)$. Here $D_{kj} = \beta_{kj} - \beta_{pj}$, β_{kj} is the j^{th} value of β_k out of M observations,

$S_D = \sum_{j=1}^M (D_j - \bar{D})(D_j - \bar{D})^T / (M - 1)$ with $D_i^T = (D_{1j}, \dots, D_{p-1j})$. Under the null hypothesis, for sufficiently large M , the statistic follows an $F_{(p-1), (M-p+1)}$ distribution. The observed value of the test statistic was 14.78 with a p-value 2.01×10^{-10} . Thus, we rejected the null hypothesis at 1% level of significance, and concluded that θ_{risk} was not a constant function. Before conducting this test we made sure that there is no significant correlation between D_{kj} and $D_{k,j+1}$ for any k . Note that if X changes from 3.0 to 3.1, according to our semiparametric Bayesian method, $\theta_{risk}(3.0) = -4.226$ and $\theta_{risk}(2.9) = -4.223$, and thus the risk increases by 0.35%. In contrast, $\theta_{risk}(3.4) = -4.113$, and $\theta_{risk}(3.3) = -4.162$ thus the risk increases by 5.03% for changing X from 3.3 to 3.4, which signifies that the rate of change of the relative risk is not constant in our method. In the SPB method, the estimated odds ratio for BMI is 1.057 with a 95% credible interval (1.021, 1.087) for an increase of 10 kg/m^2 BMI. In addition, following a referee's comment we fitted model (1) by replacing X by $\hat{X} = \hat{\eta}_0 + \hat{\eta}_1 Q + \hat{\eta}_3 Z$, where $\hat{\eta}_0$, $\hat{\eta}_1$, $\hat{\eta}_3$ are the MLE of the simple linear regression of the average of W on Q and Z based on the sub-study. The estimated odds ratio is 1.36 for one unit change in X , with a p-value = 0.03, or more specifically the risk increases by 3.04% for 0.1 unit increase in X for all values of X .

4.3 Parametric Calibration Model

One of the unique features of our approach is the semiparametric model for the calibration function given in (3), where $\theta_{\text{cal}}(\cdot)$ is an unknown smooth monotone function. In addition, we model the distribution of X in a semiparametric manner. Here we compare the results to those that assume this calibration function is linear in the unmeasured X , and that the unmeasured X is normally distributed given Z . With a slight abuse of terminology we will refer this method as the Parametric Bayes (PRB) approach. In the PRB method we modeled the surrogate variable as $Q_i = \psi_0 + \psi_1 X_i + Z_i \zeta_{\text{cal}} + U_{Q_i}$, $U_{Q_i} \sim \text{Normal}(0, \sigma_Q^2)$ and assumed that $X_i \sim \text{Normal}(\mu_0 + Z_i \zeta_x, \sigma_x^2)$. All other components and assumptions remained the same as the SPB approach. The parameters were estimated by MCMC. For this method we used the same prior that we had used in the SPB method. For ψ_0 , ψ_1 , and μ_0 we used $\text{Normal}(0, 10^2)$ priors, and used $IG(2.1, 0.91)$ prior for σ_x^2 . The posterior mean, standard deviation and 95% credible intervals for the parameters are given in Table 1. It is obvious that the estimate of σ_Q^2 is significantly larger in the PRB approach compared to the SPB method, which is an indicator that possibly $E(Q|X, Z)$ is better explained by the partly linear function $\theta_{\text{cal}}(X) + Z\zeta_{\text{cal}}$ than a linear function of X and Z . The dashed lines in Figure 1 show the estimated θ_{risk} for the Parametric Bayes approach along with a 95% credible interval. For this method, using the above F-test we again concluded that θ_{risk} is not a constant function in the interval of interest.

4.4 Model Checking

It is possible to compare the SPB and PRB models, i.e., to compare whether either $\theta_{\text{cal}}(\cdot)$ is nonlinear or the distribution of X is non-normal. In order to compare these two models we calculated the marginal likelihood under each model. Suppose $\Theta^T = (\theta^T, \varphi)$ is the set of all parameters, then the marginal likelihood under the Semiparametric Bayes model is

$m_{\text{SPB}} = \left[\prod_{i=1}^n \int f^{1-\Delta_i}(Y_i, Q_i | Z_i; \Theta) f^{\Delta_i}(Q_i, W_{ij} | Z_i; \Theta) \pi(\Theta) d\Theta \right]$, where the expressions for $f(Y_i, Q_i | Z_i; \Theta)$ and $f(Q_i, W_{ij} | Z_i; \Theta)$ are given in Section 3. Following Newton and Raftery (1994) we estimate the marginal likelihood via

$$\hat{m}_{\text{SPB}} = \left\{ \frac{1}{M} \sum_{m=1}^M \frac{1}{\prod_{i=1}^n f^{1-\Delta_i}(Y_i, Q_i | Z_i; \Theta^{(m)}) f^{\Delta_i}(Q_i, W_{ij} | Z_i; \Theta^{(m)})} \right\}^{-1},$$

where $\Theta^{(1)}, \dots, \Theta^{(M)}$ are M draws from the posterior distribution of Θ . Note that $f(Y_i, Q_i | Z_i; \Theta^{(m)}) = \int f(Y_i, Q_i | X_i, Z_i; \Theta^{(m)}) f(X_i | Z_i; \Theta^{(m)}) dX_i$ and $f(Q_i, W_{ij} | Z_i; \Theta^{(m)}) = \int f(Q_i, W_{ij} | X_i, Z_i; \Theta^{(m)}) f(X_i | Z_i; \Theta^{(m)}) dX_i$; these integrals were calculated by numerical quadrature. In a similar manner we estimated m_{PRB} via $\hat{m}_{\text{PRB}}$. Finally, we obtain $\hat{m}_{\text{SPB}} / \hat{m}_{\text{PRB}} = 111.21$, which suggest that the SPB fits the data far better than does the PRB model.

5 Simulation

To test our methodology, we performed a simulation study that has some of the complexities of the NIH-AARP study, in terms of design and nonlinearity. First we created a cohort of size $n = 10,000$ with Y, X, Q, W_1 and W_2 . We considered two different cases. In case I we simulated X from the $\text{Normal}(\gamma_0 = 3.5, \sigma_X = 0.5)$ distribution and in case II X was simulated from $(1/2)\text{Gamma}(65, 0.0455) + (1/2)\text{Normal}(3.8, 0.1^2)$. The second case clearly indicates that the distribution of X is not normal. For each case, after simulating X , we generated the unbiased variable W as $W_{ij} = X_i + U_{W_{ij}}$, where $U_{W_{ij}} \sim \text{Normal}(0, \sigma_W^2)$, $j = 1, 2$. We took two

different values for σ_w , 0.2 and 0.5, where the value 0.5 mimics the large measurement error situation common in nutritional epidemiology. The surrogate variable Q was generated as $Q_i = \theta_{\text{cal}}(X_i) + U_{Q_i}$, $U_{Q_i} \sim \text{Normal}(0, \sigma_Q^2)$ and we used two different values for σ_Q , 0.1 and 0.2. We considered a situation where $\theta_{\text{cal}}(X) = 1.25 + 0.6X$, a linear function of X and another situation where $\theta_{\text{cal}}(X) = 2 + 3 \exp\{4(X - 3.5)\} / [1 + \exp\{4(X - 3.5)\}]$. Conditional on X , the binary disease variable Y was generated with success probability $H\{\theta_{\text{risk}}(X)\}$, where $\theta_{\text{risk}}(X) = -2.8 + 1.5 \exp\{10(X - 3.5)\} / [1 + \exp\{10(X - 3.5)\}]$. On average the above probability resulted in 10% – 12% subjects with $Y = 1$. To create a calibration data from the same cohort we randomly drew $m = 1000$ subjects and obtained the variables Q , and the unbiased surrogate variables W_1 and W_2 . Note that for the primary data we only considered (Y_i, Q_i) , $i = 1, \dots, 10,000$ whereas for the calibration data we considered (Q_j, W_{j1}, W_{j2}) , $j = 1, \dots, 1000$.

Under each scenario we generated $R = 200$ data sets, and each simulated data set was analyzed by the SPB approach and by the PRB approach. In both the SPB and PRB approaches, the nonparametric effect of X on the disease risk θ_{risk} was modeled via B-splines with 9 knot points, and the knots were placed at every 10th percentile of the distribution of the average values of W starting with the 10th percentile. One should note the two main differences between the SPB and PRB approach. In the PRB approach X was always modeled as a normal random variable whereas in the SPB method X was modeled as a Dirichlet process mixture of normal random variable. In addition, in the PRB approach, $\theta_{\text{cal}}(X)$ was assumed to be linear in X , whereas in the SPB method we do not assume any specific form for θ_{cal} , rather we simply impose the restriction that $\theta_{\text{cal}}(X)$ is a smooth and a monotone nondecreasing function of X . One can of course construct many other methods to be compared with the SPB method. We have chosen to use the PRB approach to investigate the extent to which less flexible models in terms of the distribution of the latent covariate X and the calibration model affect results.

For both the SPB and PRB methods we used a Gamma(0.001, 1000) distribution for δ_β . In the analysis, we used $IG(25, 0.25)$ prior for σ_Q^2 and σ_w^2 . The PRB approach involves ψ_0 and ψ_1 , and we put a Normal(0, 10²) prior on each of them. Also, in the PRB approach we used Normal(0, 10²) prior for μ_0 and $IG(2.1, 0.91)$ prior for σ_x^2 . In the SPB method we need to estimate $\theta_{\text{cal}}(X)$ using B-splines, which involves α_j , $j = 1, \dots, q$, and we used a Gamma(0.001, 1000) prior for δ_α . The previously mentioned knot points are also used for estimating $\theta_{\text{cal}}(X)$. In the Dirichlet process prior we used $IG(2.1, 2)$ and $IG(2.1, 0.91)$ priors for τ and σ^2 respectively.

For the initial values of X and α , we first regress $(W_1 + W_2)/2$ on Q in the calibration study.

Let $\hat{\eta}_0$, $\hat{\eta}_1$ and $\hat{\sigma}_\eta^2$ be the estimated intercept, slope parameter, and the residual variance of the regression obtained using the calibration study. For the true exposure variable we set the initial value of $X_i = \hat{\eta}_0 + \hat{\eta}_1 Q_i$, for $i = 1, \dots, n$ and set $X_i = (W_{i1} + W_{i2})/2$ for $i = (n + 1), \dots, (n + m)$. The initial value of σ was obtained by fitting penalized nonparametric regression through $\{Q_i, (W_{i1} + W_{i2})/2\}$ using the calibration data. We set the initial penalty parameters $\delta_\alpha = 0.005$ and $\delta_\beta = 0.005$.

The performance of the methods were assessed via the following two statistics. The major concern of the paper is how well we estimate the risk function $\theta_{\text{risk}}(X)$. We considered a set of grid points from 2.8 to 4.2 with 0.01 increment. Letting g_j be the j^{th} grid point, for the s^{th}

data set, we estimated $\theta_{\text{risk}}(g_j)$ by $\hat{\theta}_{s, \text{risk}}(g_j)$ and let $\hat{\theta}_{\text{risk}}(g_j) = \sum_{s=1}^{200} \hat{\theta}_{s, \text{risk}}(g_j) / 200$. We computed integrated square bias as $\sum_{j=1}^{141} \{\hat{\theta}_{\text{risk}}(g_j) - \theta_{\text{risk}}(g_j)\}^2 / 141$ and integrated mean

square error as $\sum_{j=1}^{141} \sum_{s=1}^{200} \{\widehat{\theta}_{s,\text{risk}}(g_j) - \theta_{\text{risk}}(g_j)\}^2 / (141 \times 200)$. In addition to the above methods, each data set was analyzed by the *naive* method where we estimated $\theta_{\text{risk}}(X)$ using the data on (Y_i, Q_i) and assuming $X_i = Q_i$. For this method we modeled $\theta_{\text{risk}}(X)$ via B-splines and estimated β and δ_β using the logistic likelihood in a Bayesian framework.

The results of the simulation study are presented in Table 2. From the results one can quickly make the following important observations. For both the SPB and PRB methods, bias increases with σ_Q and σ_W . The value of integrated squared bias due to PRB method is significantly higher than the SPB approach when the underlying model assumptions of the PRB approach are violated. When $E(Q/X)$ is a linear function of X , and X follows a normal distribution, the performance of the SPB and PRB methods are equivalent for moderate values of σ_W . Note that the integrated squared bias of the naive method is significantly larger than that for the SPB method. Also we found that except the situation when $\theta_{\text{cal}}(X)$ is a nonlinear function of its argument and X is generated from a mixture distribution, the bias of the PRB approach is smaller than that for the naive method. Generally the variance of the SPB method is higher than the PRB approach, and the variance of the naive approach is smaller among all three approaches.

Along with the proposed method each data set was analyzed by the regression calibration technique assuming X was a linear function of Q . As expected from Carroll, et al. (2006), the method behaved very badly in comparison to the SPB and PRB methods, and we do not present the results here.

6 Discussion

In this paper, we considered the logistic regression when an error-prone covariate X is modeled nonparametrically, while exactly measured covariates Z are modeled parametrically, the result being a partially linear model. The logistic framework could be readily generalized to any distributional model.

There are two additional important and novel components to our approach. First, we recognize that in current practice, the surrogate Q for the unobserved X is not unbiased for X , and hence we have a situation where the classical measurement error model does not hold. We model the relationship of Q to (X, Z) also in a partially linear fashion, while forcing the nonparametric function $\theta_{\text{cal}}(X)$ for X to be monotone.

Finally, in a small subset of individuals, a so-called calibration study, we observed replicated unbiased versions W of X . We then model the distribution of the unobserved X nonparametrically through a Dirichlet process. As stated in the introduction, this appears to be the first paper in semiparametric measurement error models where the two functions, one of them monotone, are estimated jointly nonparametrically, while the distribution of X is modeled nonparametrically.

The simulation study shows that our method generally vastly outperforms the regression calibration method. In terms of integrated mean square error our SPB method shows much better performance than the less flexible PRB approach. The simulation study also indicates that the prices of no model assumption (naive method) and wrong model assumption (PRB method) could be significant. Instead of a single covariate with nonlinear effect, our method could also be used for the generalized additive model with multiple covariates.

One of the challenging parts of the method is the computation which requires appropriate choice of the initial values of the parameters and proposal distributions of the Metropolis-Hastings algorithm for the Gibbs sampling which are discussed in the appendix. The

computation were done using Fortran and R, and the computer code is available at <http://www.stat.tamu.edu/~sinha/research.html>.

Acknowledgments

The research of Mallick and Carroll was supported by grants from the National Cancer Institute (CA57030, CA104620) and in part by Award Number KUS-CI-016-04, made by King Abdullah University of Science and Technology (KAUST).

References

- Berry SM, Carroll RJ, Ruppert D. Bayesian smoothing and regression splines for measurement error problems. *Journal of the American Statistical Association*. 2002; 97:160–169.
- Carroll RJ, Ruppert D, Tosteson TD, Crainiceanu C, Karagas MR. Nonparametric regression and instrumental variables. *Journal of the American Statistical Association*. 2004; 99:736–750.
- Carroll, RJ.; Ruppert, D.; Stepanski, LA.; Crainiceanu, C. *Measurement Error in Nonlinear Models: A Modern Perspective*. Second Edition. New York: Chapman and Hall; 2006.
- Carroll RJ, Hall P. Optimal rates of convergence for deconvolving a density. *Journal of the American Statistical Association*. 1988; 83:1184–1186.
- Delaigle A, Hall P. Using SIMEX for smoothing parameter choices in errors in variables problems. *Journal of the American Statistical Association*. 2008; 103:280–287.
- Escobar MD, West M. Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*. 1995; 90:577–588.
- Fan J, Truong YK. Nonparametric regression with errors in variables. *Annals of Statistics*. 1993; 21:1900–1925.
- Gilks WR, Wild P. Adaptive rejection sampling for Gibbs sampling. *Applied Statistics*. 1992; 41:337–348.
- Holmes CC, Mallick BK. Generalized nonlinear modeling with multivariate free-knot regression splines. *Journal of the American Statistical Association*. 2003; 98:352–368.
- Johnson BA, Herring AH, Ibrahim JG, Siega-Riz AM. Structured measurement error in nutritional epidemiology: applications in the pregnancy, infection, and nutrition (PIN) study. *Journal of the American Statistical Association*. 2007; 102:856–866. [PubMed: 18584067]
- Kipnis V, Midthune D, Freedman LS, Bingham S, Schatzkin A, Subar A, Carroll RJ. Empirical evidence of correlated biases in dietary assessment instruments and its implications. *American Journal of Epidemiology*. 2001; 153:394–403. [PubMed: 11207158]
- Kipnis V, Subar AF, Midthune D, Freedman LS, Ballard-Barbash R, Troiano R, Bingham S, Schoeller DA, Schatzkin A, Carroll RJ. The structure of dietary measurement error: results of the OPEN biomarker study. *American Journal of Epidemiology*. 2003; 158:14–21. [PubMed: 12835281]
- Leitenstorfer F, Tutz G. Generalized monotonic regression based on B-splines with an application to air pollution data. *Biostatistics*. 2007; 8:654–673. [PubMed: 17062722]
- Mallick BK, Gelfand AE. Semiparametric errors-in-variables models: A Bayesian approach. *Journal of Statistical Planning and Inference*. 1996; 52:307–321.
- McAuliffe JD, Blei DM, Jordan MI. Nonparametric empirical Bayes for the Dirichlet process mixture model. *Statistics and Computing*. 2006; 16:5–14.
- Müller P, Roeder K. A Bayesian semiparametric model for case-control studies with errors in variables. *Biometrika*. 1997; 84:523–537.
- Newton MA, Raftery AE. Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society, Series B*. 1994; 56:3–26.
- Ruppert D. Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics*. 2002; 11:735–757.
- Ruppert, D.; Wand, MP.; Carroll, RJ. *Semiparametric Regression*. Cambridge University Press; 2003.
- Schatzkin A, Subar AF, Thompson FE, Harlan LC, Tangrea J, Hollenbeck AR, Hurwitz PE, Coyle LSNMDS, Freedman LS, Brown CC, Midthune D, Kipnis V. Design and serendipity in

- establishing a large cohort with wide dietary intake distributions. *American Journal of Epidemiology*. 2001; 154:1119–1125. [PubMed: 11744517]
- Thiébaud ACM, Kipnis V, Chang S-C, Subar AF, Thompson FE, Rosenberg PS, Hollenbeck AR, Leitzmann M, Schatzkin A. Dietary fat and postmenopausal invasive breast cancer in the National Institutes of Health – AARP diet and health study cohort. *Journal of the National Cancer Institute*. 2007; 99:451–462. [PubMed: 17374835]
- Wood S, Kohn R. A Bayesian approach to robust binary nonparametric regression. *Journal of the American Statistical Association*. 1998; 93:203–213.
- Wood S, Kohn R, Shivley T, Jiang W. Model selection in spline nonparametric regression. *Journal of the Royal Statistical Society, Series B*. 2002; 64:119–139.

Figure 1.

Plot of the estimated $\theta_{risk}(X)$ along with a pointwise 95% credible interval using the semiparametric Bayesian approach (SPB, solid lines) and the parametric Bayesian approach (PRB, dashed lines) approach.

Figure 2.

This figure shows the estimated $\theta_{cal}(X)$ using the semiparametric Bayesian (SPB) method, along with a 95% pointwise credible interval. This figure also contains the scatter plot of $(Q - \hat{\zeta}_{cal}Z, W)$ of the calibration data, and a line with zero intercept and slope one.

Figure 3. Histogram of X obtained from the MCMC sample using our semiparametric Bayesian (SPB) method along with the density plot of a normal distribution with the same mean and variance as of X .

Figure 4.

An empirical distribution of α_0 of the Dirichlet process prior which is determined according to Equation (4) in each MCMC samples of the semiparametric Bayesian (SPB) method.

Table 1

Results of the NIH-AARP data analysis by the semiparametric Bayesian (SPB) and parametric Bayesian (PRB) methods.

Method	Parameter	Posterior mean	Posterior standard deviation	95% credible interval
SPB	ζ_{risk}	0.0516	0.0173	(0.0205, 0.0830)
	ζ_{cal}	0.0008	0.0039	(-0.0046, 0.0096)
	σ_Q^2	0.0260	0.0025	(0.0229, 0.0295)
	ζ_x	0.0697	0.0038	(0.0598, 0.07677)
	σ_W^2	0.1104	0.0145	(0.0889, 0.1368)
PRB	ζ_{risk}	0.0397	0.0234	(-0.0134, 0.0754)
	ψ_0	1.3960	0.1365	(1.0577, 1.6327)
	ψ_1	0.5915	0.0463	(0.5137, 0.7036)
	ζ_{cal}	0.0268	0.0052	(0.0134, 0.0327)
	σ_Q^2	0.0496	0.0019	(0.0458, 0.0533)
	γ_0	3.1550	0.0562	(3.0576, 3.2253)
	ζ_x	0.0839	0.0199	(0.0563, 0.1187)
	σ_X^2	0.0528	0.0065	(0.0370, 0.0637)
	σ_W^2	0.0916	0.0043	(0.0836, 0.1065)

Table 2

Results of the simulation study based on 200 simulations.

Case	σ_Q	Method	ISB	IMSE	ISB	IMSE	ISB	IMSE
<hr/>								
I			$\theta_{\text{cal}} = 1.25 + 0.6X_i$					
			$\sigma_W = 0.2$			$\sigma_W = 0.5$		
			Naive					
0.1	SPB		0.0062	0.0149	0.0080	0.0155	0.0905	0.1004
		PRB	0.0058	0.0140	0.0067	0.0147		
0.2	SPB		0.0065	0.0152	0.0082	0.0164	0.0915	0.1011
		PRB	0.0060	0.0145	0.0069	0.0150		
<hr/>								
I			$\theta_{\text{cal}}(X) = 2 + 3 \exp\{4(X_i - 3.5)\}/[1 + \exp\{4(X_i - 3.5)\}]$					
			$\sigma_W = 0.2$			$\sigma_W = 0.5$		
			Naive					
0.1	SPB		0.0087	0.0161	0.0089	0.0164	0.0462	0.0551
		PRB	0.0114	0.0194	0.0144	0.0236		
0.2	SPB		0.0088	0.0169	0.0092	0.0189	0.0470	0.0575
		PRB	0.0118	0.0198	0.0150	0.0249		
<hr/>								
II			$\theta_{\text{cal}}(X) = 1.25 + 0.6X_i$					
			$\sigma_W = 0.2$			$\sigma_W = 0.5$		
			Naive					
0.1	SPB		0.0149	0.0354	0.0221	0.0429	0.1550	0.1635
		PRB	0.0283	0.0423	0.0560	0.0801		
0.2	SPB		0.0168	0.0368	0.0245	0.0450	0.2004	0.2208
		PRB	0.0289	0.0439	0.0805	0.1030		
<hr/>								
II			$\theta_{\text{cal}}(X) = 2 + 3 \exp\{4(X_i - 3.5)\}/[1 + \exp\{4(X_i - 3.5)\}]$					
			$\sigma_W = 0.2$			$\sigma_W = 0.5$		
			Naive					
0.1	SPB		0.0184	0.0394	0.0335	0.0552	0.0610	0.0668
		PRB	0.0606	0.0749	0.0684	0.0831		
0.2	SPB		0.0232	0.0450	0.0347	0.0560	0.0650	0.0710
		PRB	0.0799	0.0965	0.0874	0.1054		

Here ISB and IMSE represent integrated square bias and the integrated mean square error. Here $n = 10,000$ and $m = 1,000$. In case I X is generated from $\text{Normal}(3.5, 0.5^2)$, and in case II X is generated from $(1/2)\text{Normal}(3.8, 0.1^2) + (1/2)\text{Gamma}(65, 0.0455)$. Note that $U_{\mathcal{Q}} = \text{Normal}(0, \sigma_{\mathcal{Q}}^2)$, and $W = X + U_W$, where $U_W = \text{Normal}(0, \sigma_W^2)$. SPB and PRB stand for the proposed semiparametric Bayesian approach and the parametric Bayesian approach to the analysis of the simulated data, respectively.