



RESEARCH ARTICLE

Improving estimation and prediction in linear regression incorporating external information from an established reduced model

Wenting Cheng¹  | Jeremy M. G. Taylor¹ | Pantel S. Vokonas^{2,3} | Sung Kyun Park⁴  | Bhramar Mukherjee¹

¹Department of Biostatistics, University of Michigan, Ann Arbor, Michigan, USA

²VA Normative Aging Study, Veterans Affairs Boston Health Care System, Boston, Massachusetts, USA

³Department of Medicine, Boston University School of Medicine, Boston, Massachusetts, USA

⁴Department of Epidemiology, University of Michigan, Ann Arbor, Michigan, USA

Correspondence

Wenting Cheng, Department of Biostatistics, University of Michigan, Ann Arbor, Michigan, USA.

Email: chengwt@umich.edu

We consider a situation where there is rich historical data available for the coefficients and their standard errors in a linear regression model describing the association between a continuous outcome variable \mathbf{Y} and a set of predicting factors \mathbf{X} , from a large study. We would like to use this summary information for improving inference in an expanded model of interest, \mathbf{Y} given \mathbf{X} , \mathbf{B} . The additional variable \mathbf{B} is a new biomarker, measured on a small number of subjects in a new dataset. We formulate the problem in an inferential framework where the historical information is translated in terms of nonlinear constraints on the parameter space and propose both frequentist and Bayes solutions to this problem. We show that a Bayesian transformation approach proposed by Gunn and Dunson is a simple and effective computational method to conduct approximate Bayesian inference for this constrained parameter problem. The simulation results comparing these methods indicate that historical information on $E(\mathbf{Y}|\mathbf{X})$ can improve the efficiency of estimation and enhance the predictive power in the regression model of interest $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$. We illustrate our methodology by enhancing a published prediction model for bone lead levels in terms of blood lead and other covariates, with a new biomarker defined through a genetic risk score.

KEYWORDS

Bayesian methods, constrained estimation, prediction models

1 | INTRODUCTION

In clinical biomedicine, there are many well-known models describing the association between a measure of disease and patient characteristics, treating the measure of disease as the outcome and patient characteristics as predicting variables. Examples include Framingham risk score,¹ Prostate Cancer Prevention Trial calculator,² and Gail model.³ They can make predictions for future patients, based on their individual characteristics. These models could then be used in the settings of early detection and screening, or help decisions on treatment after diagnosis, or monitor for progression after treatment. While these models are well established, it is possible that including some additional candidate biomarkers and constructing an expanded model will improve the prediction ability.

The challenge of estimating the expanded model is that the additional biomarkers are measured only on a small number of subjects in a new dataset; thus, inference in the expanded model tends to give relatively poor coefficient estimates with large standard errors and low prediction accuracy. It is natural to consider incorporating information that is available from an established model into the expanded model to improve the estimates of the parameters and the prediction ability of the newly developed model. Such external information is often available; however, it may not come in a direct or convenient form. We consider a situation where the outcome is a continuous marker of disease risk and the established regression model is described in an article, in which the estimated regression coefficients and their standard errors are presented in tables. The expanded model, however, includes one additional biomarker as a predicting variable. How to incorporate this coefficient information in a principled way is a nontrivial statistical problem.

The use of external information is a popular strategy for improving efficiency in statistical inference. Often, the information can be expressed as constraints on the regression coefficients and one can conduct constrained maximum likelihood inference. The problem of inference for regression coefficients from linear regression subject to a set of constraints has been considered from the Bayesian perspective, either by discarding draws violating the constraints⁴ or translating the constraints as informative priors. Geweke uses noninformative priors and an indicator function representing the inequality constraints. The posterior distributions are then computed using importance sampling. Although this idea is easy to implement, it could be extremely slow computationally, especially when the truncation region has a small probability. Dunson and Neelon⁵ as well as Gunn and Dunson⁶ propose a simple approach to handle constraints by generating sample draws from the unconstrained posterior distribution and mapping these draws to the constrained space. Their interest is primarily in order-restricted inference, and they choose the constrained draw that minimizes the Mahalanobis distance between the unconstrained draws and the ordered draws, across different choices of ordered draws.

There is literature emerging on new frequentist proposals to incorporate external information. Chen et al⁷ propose a linear regression shrinkage method for predictions in a small dataset calibrated by a larger but biased dataset. Imbens and Lancaster⁸ investigate how aggregate data (eg, the population average of the response) could be used to improve maximum likelihood estimates in a regression model. They show that the gains from incorporating such information could be substantial. Qin⁹ proposes that the aggregate data can be incorporated into the empirical likelihood and the combination of empirical and parametric likelihood could provide valid inference for the regression coefficients. Qin et al¹⁰ consider auxiliary information (eg, disease prevalence at different levels of risk factors) as constraints on the regression coefficients and the joint covariate distribution. They use empirical likelihood and general estimating equations for estimation. Chatterjee et al¹¹ use summary-level information from external data sources of large sample size to calibrate the current regression model.

To introduce notation, let \mathbf{Y} denote the outcome, which is assumed to be continuous, and we have a set of standard risk covariates \mathbf{X} (there is no assumption regarding the distribution of \mathbf{X}) and a new continuous covariate \mathbf{B} measured on a small dataset. The model of primary interest is a regression model that describes the joint effect of \mathbf{X} , \mathbf{B} on \mathbf{Y} :

$$E(\mathbf{Y}|\mathbf{X}, \mathbf{B}) = \mathbf{X}\boldsymbol{\gamma}_X + \mathbf{B}\boldsymbol{\gamma}_B. \quad (1)$$

We could also estimate $E(\mathbf{B}|\mathbf{X})$ in this small dataset from a model of the form:

$$E(\mathbf{B}|\mathbf{X}) = \mathbf{X}\boldsymbol{\theta}. \quad (2)$$

A large, well-characterized previous study describes the association between \mathbf{X} and \mathbf{Y} through a regression model:

$$E(\mathbf{Y}|\mathbf{X}) = \mathbf{X}\boldsymbol{\beta}. \quad (3)$$

The knowledge we obtain from the previous study is summary-level information on model (3): estimated regression coefficients and their standard errors. We assume we know the empirical variances (standard error squared) of regression coefficients, but not their covariances, since the estimated covariance matrices are rarely reported in publications. We use $\hat{\boldsymbol{\beta}}$ and $\hat{\mathbf{S}}$ to denote the reported coefficient estimates and their standard errors.

We formulate the problem in an inferential framework where the external information from (3) is translated in terms of nonlinear constraints on the regression parameters and propose both frequentist and Bayes solutions to this problem. The goal is to improve the estimation of linear regression coefficients $\boldsymbol{\gamma}$ and prediction power of model (1) incorporating external coefficients information from (3), when the sample size is small. When the current dataset is large, the potential gain by incorporating external information from an established reduced model may be limited.

The following is the structure of the remainder of this article: In Section 2, we discuss how to transform the available external information into constraints on the regression coefficients. We show 2 unconstrained solutions based on current

data, ignoring historical information and propose 4 constrained solutions that use the historical information: constrained maximum likelihood, partial regression, informative full Bayes, and Bayesian transformation approach. We present a simulation study in Section 3. Section 4 is an application of the approaches to enhance a prediction model for bone lead levels based on data from the Normative Aging Study published in 2009 with new genetic marker information. We discuss the findings and possibilities for future work in Section 5.

2 | STATISTICAL APPROACHES

2.1 | Relationship equations

Assume that \mathbf{X} has $p + 1$ dimensions (including an intercept). $\mathbf{X}_0 = 1$ by notational convention. From models (1), (2), and (3), we find that the regression of \mathbf{Y} on \mathbf{X}, \mathbf{B} is a linear function of $E(\mathbf{B}|\mathbf{X})$:

$$E(\mathbf{Y}|\mathbf{X}) = E(E(\mathbf{Y}|\mathbf{X}, \mathbf{B})|\mathbf{X}) = E(\mathbf{X}\boldsymbol{\gamma}_X + \mathbf{B}\boldsymbol{\gamma}_B|\mathbf{X}) = \mathbf{X}\boldsymbol{\gamma}_X + \boldsymbol{\gamma}_B E(\mathbf{B}|\mathbf{X}). \tag{4}$$

We are going to estimate model (1), $E(\mathbf{Y}|\mathbf{X}, \mathbf{B}) = \gamma_0 + \gamma_1\mathbf{X}_1 + \dots + \gamma_p\mathbf{X}_p + \boldsymbol{\gamma}_{p+1}\mathbf{B}$ and model (2), $E(\mathbf{B}|\mathbf{X}) = \theta_0 + \theta_1\mathbf{X}_1 + \dots + \theta_p\mathbf{X}_p$ from the current small dataset. Using Equation 4 and the historical information in model (3), we have the following equation:

$$\bar{\beta}_0 + \bar{\beta}_1\mathbf{X}_1 + \dots + \bar{\beta}_p\mathbf{X}_p = \gamma_0 + \gamma_1\mathbf{X}_1 + \dots + \gamma_p\mathbf{X}_p + \boldsymbol{\gamma}_{p+1}(\theta_0 + \theta_1\mathbf{X}_1 + \dots + \theta_p\mathbf{X}_p), \tag{5}$$

which implies the relationship between parameters in models (1) to (3) as

$$\beta_j = \gamma_j + \boldsymbol{\gamma}_{p+1}\theta_j, j = 0, \dots, p. \tag{6}$$

We summarize the necessary assumptions required for constructing the relationship equations (6) in Table 1. Essentially, when constructing the relationship equations connecting the parameters $\boldsymbol{\beta}, \boldsymbol{\gamma}$ and $\boldsymbol{\theta}$, from the models for $E(\mathbf{Y}|\mathbf{X}), E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$, and $E(\mathbf{B}|\mathbf{X})$ respectively, we assume that \mathbf{Y} and \mathbf{B} are continuous variables. We do assume that although \mathbf{B} is not available in the external historical data, if it were available, $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ is linear in (\mathbf{X}, \mathbf{B}) , $E(\mathbf{Y}|\mathbf{X})$ is linear in \mathbf{X} , and $E(\mathbf{B}|\mathbf{X})$ is linear in \mathbf{X} . We further assume that the models for $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ and $E(\mathbf{Y}|\mathbf{X})$ are correctly specified in both the internal and the external study and they are the same in both internal and external populations. Moreover, by using the summary information on the regression coefficients from the large dataset, we implicitly assume that the available results provide consistent estimates of $\boldsymbol{\beta}$ and standard errors of $\boldsymbol{\beta}$. No additional distributional assumptions are needed to establish the constraints.

TABLE 1 A summary of necessary theoretical assumptions required in constructing the relationship equations

Distribution	Internal dataset	External dataset
$E(\mathbf{Y} \mathbf{X}, \mathbf{B})$	Linear in (\mathbf{X}, \mathbf{B}) ; Correctly specified	Linear in (\mathbf{X}, \mathbf{B}) ; Correctly specified
$E(\mathbf{Y} \mathbf{X}, \mathbf{B})$ is the same in the 2 datasets		
$E(\mathbf{Y} \mathbf{X})$	Linear in \mathbf{X} ; Correctly specified	Linear in \mathbf{X} ; Correctly specified
$E(\mathbf{Y} \mathbf{X})$ is the same in the 2 datasets		
$E(\mathbf{B} \mathbf{X})$	Linear in \mathbf{X} ; Correctly specified	Linear in \mathbf{X} Correctly specified
$E(\mathbf{B} \mathbf{X})$ is the same in the 2 datasets		
$(\mathbf{Y}, \mathbf{X}, \mathbf{B})$	No assumptions about the joint distribution of $(\mathbf{Y}, \mathbf{X}, \mathbf{B})$ in either dataset other than the assumption on $E(\mathbf{Y} \mathbf{X}, \mathbf{B})$	
(\mathbf{X}, \mathbf{B})	No assumptions about the joint distribution of (\mathbf{X}, \mathbf{B}) in either dataset other than the assumption on $E(\mathbf{B} \mathbf{X})$	
(\mathbf{X})	Distribution of (\mathbf{X}) does not have to be the same in the two datasets	

2.2 | Unconstrained solutions

2.2.1 | Direct regression

Without constraints, γ can be estimated by ordinary least squares directly:

$$\min_{\gamma} \sum_{i=1}^n \left(Y_i - \sum_{j=0}^p \gamma_j X_{ij} - \gamma_{p+1} B_i \right)^2. \quad (7)$$

This estimates parameters in model (1). To estimate parameters in model (2), we obtain least squares estimates of θ by considering:

$$\min_{\theta} \sum_{i=1}^n \left(B_i - \sum_{j=0}^p \theta_j X_{ij} \right)^2. \quad (8)$$

2.2.2 | Standard Bayes

Analogous to direct regression without constraints, we perform standard parametric Bayesian linear regression with noninformative conjugate priors.¹² We use standard Bayes linear regression procedures to fit model (1) and model (2) separately.

For model (1), the likelihood function is derived from $\mathbf{Y}|\mathbf{X}, \mathbf{B}, \boldsymbol{\gamma}, \sigma_1^2 \sim N_n(\sum_{j=0}^p \gamma_j \mathbf{X}_j + \gamma_{p+1} \mathbf{B}, \sigma_1^2 \mathbf{I}_n)$. We follow the standard prescription of choosing independent normal-inverse-gamma conjugate priors for the regression coefficients $\boldsymbol{\gamma}$ and residual variance σ_1^2 as we are trying to avoid the use of any historical prior information in this analysis¹³: $\pi(\boldsymbol{\gamma}, \sigma_1^2) = \pi(\boldsymbol{\gamma}) \cdot \pi(\sigma_1^2) = N_{(p+2)}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) \cdot \text{inverse-gamma}(\frac{\nu_0}{2}, \frac{\nu_0 \sigma_0^2}{2}) = N(0, 100^2 \mathbf{I}_{(p+2)(p+2)}) \cdot \text{IG}(0.01, 0.01)$. We choose the prior variance to be large for the purpose of choosing a proper prior on the regression coefficients that is flat. The joint posterior distribution $p(\boldsymbol{\gamma}, \sigma_1^2 | \mathbf{Y}, \mathbf{X}, \mathbf{B})$ is then proportional to $p(\mathbf{Y}|\mathbf{X}, \mathbf{B}, \boldsymbol{\gamma}, \sigma_1^2) \cdot \pi(\boldsymbol{\gamma}) \cdot \pi(\sigma_1^2)$. The full conditional distribution of $\boldsymbol{\gamma}$ (ie, $p(\boldsymbol{\gamma} | \mathbf{Y}, \mathbf{X}, \mathbf{B}, \sigma_1^2)$) is then $N\left(\left(\boldsymbol{\Sigma}_0^{-1} + \frac{(\mathbf{X}, \mathbf{B})^T (\mathbf{X}, \mathbf{B})}{\sigma_1^2}\right)^{-1} \left(\boldsymbol{\Sigma}_0^{-1} \boldsymbol{\mu}_0 + \frac{(\mathbf{X}, \mathbf{B})^T \mathbf{Y}}{\sigma_1^2}\right), \left(\boldsymbol{\Sigma}_0^{-1} + \frac{(\mathbf{X}, \mathbf{B})^T (\mathbf{X}, \mathbf{B})}{\sigma_1^2}\right)^{-1}\right)$ and the conditional distribution of σ_1^2 is inverse-gamma $\left(\frac{n+\nu_0}{2}, \frac{1}{2}[(\mathbf{Y} - (\mathbf{X}, \mathbf{B})\boldsymbol{\gamma})^T (\mathbf{Y} - (\mathbf{X}, \mathbf{B})\boldsymbol{\gamma}) + \nu_0 \sigma_0^2]\right)$. For model (2), the prior specifications and inferences are very similar to that of model (1).

Using Markov chain sampling techniques like Gibbs sampling, standard Bayes can be implemented in a fast and easy algorithm to obtain posterior draws of $\boldsymbol{\gamma}$ and θ . Direct regression and standard Bayes do not use external information and serve as references for quantifying the amount of efficiency we gain by using external information.

2.3 | Constrained solutions

2.3.1 | Constrained maximum likelihood

The constrained maximum likelihood (constrained ML) method uses optimization of the likelihood under the constraints in (6). As we have information on both the point estimate and the standard error of $\boldsymbol{\beta}$, we will require estimates of the parameters such that the new $\boldsymbol{\beta}$ to be within d standard errors of the old point estimate. Our constrained maximum likelihood estimation optimizes the joint log-likelihood, namely,

$$\log(L) = \sum_{i=1}^n \left[-\frac{1}{2} \log(2\pi\sigma_1^2) - \frac{1}{2\sigma_1^2} \left(Y_i - \sum_{j=0}^p \gamma_j X_{ij} - \gamma_{p+1} B_i \right)^2 \right] + \sum_{i=1}^n \left[-\frac{1}{2} \log(2\pi\sigma_2^2) - \frac{1}{2\sigma_2^2} \left(B_i - \sum_{j=0}^p \theta_j X_{ij} \right)^2 \right] \quad (9)$$

subject to the set of nonlinear constraints: $\gamma_j + \gamma_{p+1}\theta_j \in [\bar{\beta}_j - d\bar{S}_j, \bar{\beta}_j + d\bar{S}_j], j = 0, \dots, p$.

This method is equivalent to minimizing the weighted sum of squared errors of $\mathbf{Y}|\mathbf{X}, \mathbf{B}$ and the squared errors of $\mathbf{B}|\mathbf{X}$, namely,

$$\min_{\boldsymbol{\gamma}, \boldsymbol{\theta}} \left\{ \frac{1}{\sigma_1^2} \sum_{i=1}^n \left(Y_i - \sum_{j=0}^p \gamma_j X_{ij} - \gamma_{p+1} B_i \right)^2 + \frac{1}{\sigma_2^2} \sum_{i=1}^n \left(B_i - \sum_{j=0}^p \theta_j X_{ij} \right)^2 \right\} \quad (10)$$

s.t. $\gamma_j + \gamma_{p+1}\theta_j \in [\bar{\beta}_j - d\bar{S}_j, \bar{\beta}_j + d\bar{S}_j], j = 0, \dots, p$.

In this optimization problem, instead of treating σ_1^2 and σ_2^2 as unknown parameters, we use $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$ as plug-in estimates, which are the OLS residual variances from $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ and $E(\mathbf{B}|\mathbf{X})$.

The width of the constrained interval is controlled by d , which is a scale parameter representing the strength of external information. From simulations, we find that fixing d as $d = 1$ may be reasonable. Since the constraints in (10) are of the form $\gamma_j + \gamma_{p+1}\theta_j \in [\bar{\beta}_j - d\bar{S}_j, \bar{\beta}_j + d\bar{S}_j]$, $j = 0, \dots, p$, fixing $d = 1$ allows one standard error departure from the observed point estimate in the historical data. Further discussions about the choice of d are in Section 5 and in the Supporting Information. This is an optimization problem with nonlinear inequality constraints. To solve it, we use function `solnp` in R package `Rsolnp`, a function that efficiently solves general nonlinear optimization problems using Lagrange multipliers. The starting point is the OLS estimates of γ and θ , namely, $\hat{\gamma}$ and $\hat{\theta}$. For computational convenience, we further specify wide lower and upper bounds for each of the parameters: $\gamma_j \in [\hat{\gamma}_j - 5\hat{SE}(\gamma_j), \hat{\gamma}_j + 5\hat{SE}(\gamma_j)]$, $j = 0, \dots, p + 1$, $\theta_j \in [\hat{\theta}_j - 5\hat{SE}(\theta_j), \hat{\theta}_j + 5\hat{SE}(\theta_j)]$, $j = 0, \dots, p$.

The standard error of the estimates in the constrained ML solution is hard to derive. Usually, the distribution of the constrained maximum likelihood estimate may be derived by expressing the constrained estimate as functions of both the unconstrained estimate and the data and then applying Taylor expansion. However, the fact that the constraints are in the form that is known as “box” constraints in the optimization literature in (10) and these box constraints involve nonlinear functions of the regression parameters makes it impossible to implement this procedure for our solution. Instead, we use the bootstrap to estimate the standard error. The bootstrap procedure is described in Appendix S1.

In simulations, we find that the constrained maximum likelihood estimate can show substantial bias for small sample sizes. As an alternative, we consider the bootstrap bias-corrected estimate $\hat{\gamma}_{bc}$, given by $\hat{\gamma}_{bc} = 2\hat{\gamma} - \bar{\gamma}$, where $\hat{\gamma}$ is the original estimate, $\bar{\gamma}$ is the mean of the bootstrap estimates.¹⁴ We use this bootstrap bias-correction procedure to modify the constrained maximum likelihood solution in the simulation studies. For the real data analysis, we also provide a bias-corrected 95% confidence interval: $(F_{\bar{\gamma}}^{-1}(\Phi(2\mathbf{b} + Z_{0.025})), F_{\bar{\gamma}}^{-1}(\Phi(2\mathbf{b} + Z_{0.975})))$ where \mathbf{b} is estimated from the bootstrap distribution by $\Phi^{-1}(\Pr(\bar{\gamma} \leq \hat{\gamma}))$, and Φ is the cumulative distribution function of the normal distribution.^{15,16} When the sample size is small, we recommend bootstrap bias-corrected constrained ML instead of the constrained ML as an alternative method for reducing the bias. However, this bootstrap bias-corrected constrained ML is not necessary for large sample sizes. Additional simulation studies in the Supporting Information demonstrate the performance of this bias-corrected bootstrap procedure in simulation settings of various sample sizes.

2.3.2 | Partial regression

Partial regression is an indirect method to estimate the amount by which a dependent variable increases when one of the predicting variables is increased by one unit with all other predicting variables held constant.¹⁷

Our adaptation of the partial regression method is an attempt to look at the relationship between the response and the new explanatory variable while preserving the effect from the old set of explanatory variables. The following 3 simple steps describe the proposed partial regression method:

1. Remove the effect of variables \mathbf{X} on the response \mathbf{Y} by computing \mathbf{r}_1 : $\mathbf{r}_1 = \mathbf{Y} - \sum_{j=0}^p \bar{\beta}_j \mathbf{X}_j$. That is, remove the effect of \mathbf{X} on \mathbf{Y} in the small dataset using historical information, as reflected through plugging in the estimate $\bar{\beta}$.
2. Remove from \mathbf{B} the effect of correlation due to variables \mathbf{X} : estimate coefficients θ from $E(\mathbf{B}|\mathbf{X})$, calculate $\mathbf{r}_2 = \mathbf{B} - \sum_{j=0}^p \hat{\theta}_j \mathbf{X}_j$. That is, regress \mathbf{B} against \mathbf{X} from our small dataset.
3. Regress \mathbf{Y} -residuals against \mathbf{B} -residuals: estimate coefficients α_0, α_1 , in $E(r_1|r_2) = \alpha_0 + \alpha_1 r_2$ by OLS. The estimated coefficients for $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ are then given by

$$\begin{cases} \hat{\gamma}_0 = \hat{\alpha}_0 - \hat{\alpha}_1 \hat{\theta}_0 + \bar{\beta}_0 \\ \hat{\gamma}_j = -\hat{\alpha}_1 \hat{\theta}_j + \bar{\beta}_j, j = 1, \dots, p \\ \hat{\gamma}_{p+1} = \hat{\alpha}_1. \end{cases} \quad (11)$$

The standard error of the partial regression estimate is not easy to derive. The partial regression estimate depends on the estimated θ from \mathbf{B} regressed on \mathbf{X} and the estimated β from historical information. Because of this dependence, the marginal distribution of the partial regression estimate is not in closed form. We use a variance approximation linearization technique to estimate the standard error of the partial regression estimate. This procedure is described in Appendix S2.

2.3.3 | Informative full Bayes

We suggest a Bayesian approach with informative priors, based on a Markov chain Monte Carlo (MCMC) technique using a Metropolis-Hastings sampling algorithm. The first step is to write down the joint likelihood function, $L(\mathbf{Y}, \mathbf{B}|\mathbf{X}) = L(\mathbf{Y}|\mathbf{X}, \mathbf{B})L(\mathbf{B}|\mathbf{X})$ with prior $\pi(\boldsymbol{\gamma}, \boldsymbol{\theta}, \sigma_1^2, \sigma_2^2)$. This joint likelihood function is a valid likelihood and therefore could be used for Bayesian inference.

$$p(\boldsymbol{\gamma}, \boldsymbol{\theta}, \sigma_1^2, \sigma_2^2 | data) \propto \left\{ \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma_1^2}} e^{-\frac{1}{2\sigma_1^2} \left(Y_i - \sum_{j=0}^p \gamma_j X_{ij} - \gamma_{p+1} B_i \right)^2} \cdot \frac{1}{\sqrt{2\pi\sigma_2^2}} e^{-\frac{1}{2\sigma_2^2} \left(B_i - \sum_{j=0}^p \frac{\beta_j - \gamma_j}{\gamma_{p+1}} X_{ij} \right)^2} \right\} \cdot \pi(\boldsymbol{\gamma}, \boldsymbol{\theta}, \sigma_1^2, \sigma_2^2) \quad (12)$$

We can reparameterize (12) in terms of variables $\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_1^2, \sigma_2^2$, by a Jacobian transformation using the constraints in (6) as the underlying transformation. The Jacobian matrix is denoted by \mathbf{J} . We further assume independent priors for $\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_1^2$ and σ_2^2 . Since we have no information for parameters $\boldsymbol{\gamma}, \sigma_1^2, \sigma_2^2$, we assume noninformative priors $N(0, 100^2 I_{(p+2) \times (p+2)})$, $IG(0.01, 0.01)$, and $IG(0.01, 0.01)$; for parameter $\boldsymbol{\beta}$, we use the constraints directly as priors:

$$\beta_j = \gamma_j + \gamma_{p+1} \theta_j \sim N(\bar{\beta}_j, \bar{S}_j^2), j = 0, \dots, p. \quad (13)$$

Then, we can rewrite the joint posterior distribution of $\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_1^2, \sigma_2^2$ as

$$p(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_1^2, \sigma_2^2 | data) \propto \left\{ \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma_1^2}} e^{-\frac{1}{2\sigma_1^2} \left(Y_i - \sum_{j=0}^p \gamma_j X_{ij} - \gamma_{p+1} B_i \right)^2} \cdot \frac{1}{\sqrt{2\pi\sigma_2^2}} e^{-\frac{1}{2\sigma_2^2} \left(B_i - \sum_{j=0}^p \frac{\beta_j - \gamma_j}{\gamma_{p+1}} X_{ij} \right)^2} \right\} \cdot \pi(\boldsymbol{\beta}) \cdot \pi(\boldsymbol{\gamma}) \cdot \pi(\sigma_1^2) \cdot \pi(\sigma_2^2) \cdot |\mathbf{J}| \quad (14)$$

After some algebraic calculations, we find that the conditional distribution of β_0, \dots, β_p are normal, each with distribution function $N(\mu_{\beta_j, n}, \sigma_{\beta_j, n}^2), j = 0, \dots, p$; the conditional distribution of $\gamma_0, \dots, \gamma_p$ are normal, each with distribution function $N(\mu_{\gamma_j, n}, \sigma_{\gamma_j, n}^2), j = 0, \dots, p$; the conditional distribution of σ_1^2 and σ_2^2 are $IG\left(\frac{v_{1,n}}{2}, \frac{v_{1,n}\sigma_{1,n}^2}{2}\right)$ and $IG\left(\frac{v_{2,n}}{2}, \frac{v_{2,n}\sigma_{2,n}^2}{2}\right)$, respectively. The full conditional distribution of γ_{p+1} does not have a closed form. We use a Metropolis-Hastings sampling algorithm to obtain samples from the full conditional of γ_{p+1} . The complete form of the full conditional distributions are presented in Appendix S3.

A common drawback of approaches based on vanilla MCMC technique with a constrained parameter space is that the convergence rate is too slow. Roberts¹⁸ and Gilks and Roberts¹⁹ suggest that the rate of convergence depends on the posterior correlation between the sample draws of the parameters. We find that in our problem, due to the nonlinear relationship between the parameters, we obtain highly correlated posterior draws in the Markov chain and thus the effective draws are only a small portion of the total draws in the chain. Also, the chain does not move rapidly through the entire support of the posterior distribution and has poor mixing properties. As a consequence, although this informative full Bayes approach provides an exact posterior distribution for all parameters, it is not computationally efficient in our problem.

2.3.4 | Bayesian transformation approach

We would like to find an approximate Bayes approach that is computationally inexpensive under constraints when compared to the informative full Bayes described in the previous section. The motivation for our approach stems from the transformation approach incorporating monotone or unimodal constraints proposed in Gunn and Dunson,⁶ which we described in Section 1. Gunn and Dunson⁶ show that under monotone transformations/order restrictions, the posterior mode with transformed draws can be shown to be consistent estimator for the true posterior mean of the proper constrained Bayes solution. Beyond such monotone constraints, the intuition is that it is desirable that a Bayesian approach could produce posterior draws which are compatible with the constraints and thus tends to result in values of draws with high density in the constrained space. These constrained posterior draws are based on minimal movement from the unconstrained draws produced by standard Bayes. Inspired by their idea, we first use the unconstrained Bayes method implemented with Gibbs sampling to characterize the posterior distribution and then map the draws from this posterior distribution into the constrained space. We modify their approach in 2 ways:

1. The constraints: The constraints in our problem are a set of inequality constraints on regression coefficients. The inequality constraints are obtained from historical information and are posed directly on the regression coefficients. These box constraints can be relaxed or strengthened depending on to what extent one wishes to use the historical information, through the choice of the window size “d.”

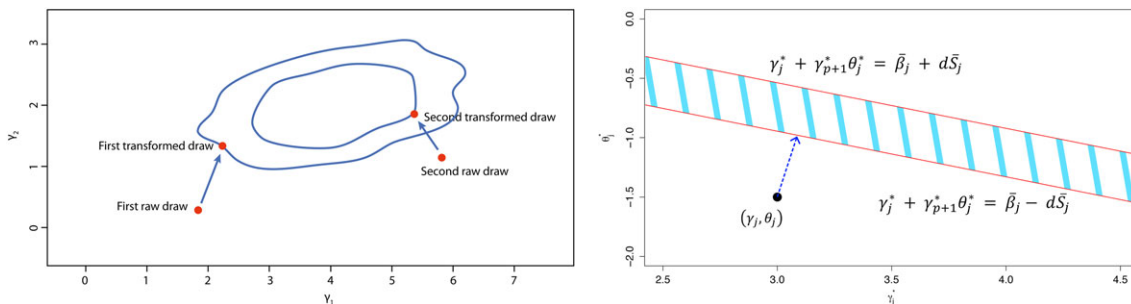
2. The distance measure: The dissimilarity between the unconstrained draws and the constrained draws is measured by normalized Euclidean distance rather than Mahalanobis distance.⁶ Mahalanobis distance is a dissimilarity measure between 2 random vectors \vec{v}_1, \vec{v}_2 with covariance matrix S as $d_{MD}(\vec{v}_1, \vec{v}_2) = \sqrt{(\vec{v}_1 - \vec{v}_2)^T S^{-1} (\vec{v}_1 - \vec{v}_2)}$. If the above covariance matrix is diagonal, the Mahalanobis distance reduces to normalized Euclidean distance: $d_{NED}(\vec{v}_1, \vec{v}_2) = \sqrt{\sum_{i=1}^N \frac{(v_{1i} - v_{2i})^2}{S_i}}$. This normalized Euclidean distance is preferred to the Mahalanobis distance for our problem because it contains only separable functions: The distance measure in each direction is detached from the distance measures in all other directions. This is especially useful for improving the computational efficiency. Also, this distance measure has a natural appeal because it only requires the knowledge of variances instead of the full variance-covariance matrix of the established model, as it is often the case that in literature the established model is presented in a table where standard errors/confidence intervals of the reported estimates are provided.

In general, our transformation approach is defined as follows: Assume a vector of parameters $\gamma_{(1 \times p)}$ are the coefficients in a regression model and subject to some inequality constraints $C : \{C_1, C_2, \dots, C_m\}$. If γ are the coefficient estimates that can be easily obtained (ie, computationally efficient) from standard Bayesian regression ignoring the constraints and $\Omega \subset \mathbb{R}^p$ is a subset of \mathbb{R}^p defined by constraints C on the elements of γ , $\Omega = \{\gamma : \gamma \text{ satisfy } C\}$. Then,

$$\forall \gamma, \gamma^* := \operatorname{argmin}(d_{NED}^2(\gamma, \gamma^*)) \quad s.t. \gamma^* \in \Omega. \tag{15}$$

Figure 1A illustrates how a draw of γ is transformed to a new γ^* . Next, we are going to apply this Bayesian transformation approach to our problem of interest. Suppose the draws from standard Bayesian linear regression on Y against X, B are $\gamma_0, \dots, \gamma_p, \gamma_{p+1}$ and the draws from standard Bayesian linear regression on B against X are $\theta_0, \dots, \theta_p$. We call these draws “raw draws,” because these draws are for the unconstrained problem. The corresponding ordinary least squares estimates are $\hat{\gamma}, \hat{\theta}$ and estimated covariance matrices are $\hat{\Sigma}_\gamma, \hat{\Sigma}_\theta$. We extract the estimated variances of the coefficients from $\hat{\Sigma}_\gamma, \hat{\Sigma}_\theta$, denote them by $s_{\gamma_0}^2, \dots, s_{\gamma_p}^2, s_{\gamma_{p+1}}^2, s_{\theta_0}^2, \dots, s_{\theta_p}^2$. Then γ^*, θ^* are obtained from the unconstrained draws $\gamma_0, \dots, \gamma_p, \gamma_{p+1}, \theta_0, \dots, \theta_p$ by solving the following optimization problem:

$$\begin{aligned} \min_{\gamma_0^*, \dots, \gamma_p^*, \gamma_{p+1}^*, \theta_0^*, \dots, \theta_p^*} & [d_{NED}^2(\gamma, \gamma^*) + d_{NED}^2(\theta, \theta^*)] = \sum_{j=0}^{p+1} \frac{(\gamma_j - \gamma_j^*)^2}{s_{\gamma_j}^2} + \sum_{k=0}^p \frac{(\theta_k - \theta_k^*)^2}{s_{\theta_k}^2} \\ \text{s.t.} & \gamma_0^* + \gamma_{p+1}^* \theta_0^* \in [\bar{\beta}_0 - d\bar{S}_0, \bar{\beta}_0 + d\bar{S}_0] \\ & \dots \\ & \gamma_p^* + \gamma_{p+1}^* \theta_p^* \in [\bar{\beta}_p - d\bar{S}_p, \bar{\beta}_p + d\bar{S}_p] \\ & \gamma_j^* \in [\hat{\gamma}_j - 5s_{\gamma_j}, \hat{\gamma}_j + 5s_{\gamma_j}], j = 0, \dots, p + 1 \\ & \theta_k^* \in [\hat{\theta}_k - 5s_{\theta_k}, \hat{\theta}_k + 5s_{\theta_k}], k = 0, \dots, p, \end{aligned} \tag{16}$$



(A) An illustration of how two constrained draws are obtained from the raw draws in the Bayesian transformation approach

(B) Two-dimensional optimization problem

FIGURE 1 Procedures of Bayesian transformation approach [Colour figure can be viewed at wileyonlinelibrary.com]

where the last 2 constraints are trivial bounds for each parameter for improving computational efficiency. The scale parameter d controls the degree of trust in the historical information as before. d is drawn from a half normal distribution to reflect the reality that there is uncertainty in the point estimates. Through simulations, we find that the choice of $|N(0, 1)|$ is reasonable.

The intuition behind these transformed draws generated by (16) is that it will produce values γ^*, θ^* subject to the box constraints that are closest to the unconstrained values γ, θ in normalized Euclidean distance. The squared normalized Euclidean distance measure (and its square root version) in the objective function in (16) minimizes the weighted Euclidean distance between (γ^*, θ^*) and (γ, θ) . The normalization is to ensure that the distance is relatively small for a particular coefficient if its OLS estimate is more precise (ie, the estimated variance is relatively small) while the distance is relatively large for those coefficients that have more uncertainty by OLS.

We obtain posterior draws based on the following steps:

- Obtain raw draws: We first obtain draws $\gamma_0, \dots, \gamma_p, \gamma_{p+1}, \theta_0, \dots, \theta_p$ from standard Bayes. The draws from Standard Bayes are collected after a burn-in period and a thinning procedure.
- Transformation of a single draw: For the first draw $\gamma_0^{(1)}, \dots, \gamma_p^{(1)}, \gamma_{p+1}^{(1)}, \theta_0^{(1)}, \dots, \theta_p^{(1)}$, generate $d^{(1)}$ from half normal distribution: $d \sim |N(0, 1)|$. Perform the transformation procedure described in (16) and obtain the first posterior draw $\gamma_0^{*(1)}, \dots, \gamma_p^{*(1)}, \gamma_{p+1}^{*(1)}, \theta_0^{*(1)}, \dots, \theta_p^{*(1)}$.
- Iterate: Repeat the above step for each draw until all draws are transformed. This means that the transformation procedure is performed 1000 times to obtain 1000 new draws.

Bayesian transformation approach is an ad hoc Bayes-type method that can be easily implemented by using draws from an unconstrained model and making them compatible with constraints derived from the external coefficient information by minimizing the normalized Euclidean distance of the raw draws from the constrained space. It is an approximate and somewhat ad hoc Bayes approach due to the fact that it does not give exact posterior inference as informative full Bayes does. Although this Bayes method does not conduct exact posterior inference, it is a pragmatic choice that was seen to both improve the estimation efficiency of the regression parameters and the predictive power in the regression model of primary interest, compared with standard Bayes without considering constraints. Getting raw draws from a standard Bayes approach is computationally efficient. The ideal Gibbs procedure with constraints is highly inefficient due to rejection of many draws outside the constrained space. The transformation approach tries to take advantage of making draws from the unconstrained model and mapping them to satisfy the constraints. The proposed Bayesian transformation approach is similar in spirit to constrained MLE in that both of them involve a mapping of the parameters from the unconstrained space to the constrained space. The difference is that the constrained MLE maps a single unconstrained MLE to the constrained space while the Bayesian transformation approach requires a set of unconstrained draws to be converted to the constrained draws, thus providing a natural measure of uncertainty.

The computational efficiency of this approach will be discussed next. The objective function is a convex function and the constraints are a set of nonlinear inequality box constraints. They involve in total $2p + 3$ parameters and $p + 1$ constraints. We could simply rely on R function `solnp` in R package `Rsolnp` to solve this minimization problem. However, since this minimization procedure needs to be applied multiple times to obtain many posterior draws, simplifying the computation is desired. Moreover, with a growing number of predicting variables, this optimization will be of higher dimension and probably harder to solve.

To find an efficient algorithm for this minimization problem, we need to simplify this multiple-parameter minimization problem. We first notice that for these constraints, the nonlinearity is due to the fact that γ_{p+1}^* appears in every constraint. All other parameters appear in pairs $(\gamma_j^*, \theta_j^*), j = 0, \dots, p$ and are only involved in one term within the sum in (16). Thus, for each pair (γ_j^*, θ_j^*) , the minimization function reduces to

$$\begin{aligned} \min_{\gamma_j^*, \theta_j^*} & \frac{(\gamma_j - \gamma_j^*)^2}{s_{\gamma_j}^2} + \frac{(\theta_j - \theta_j^*)^2}{s_{\theta_j}^2} \\ \text{s.t.} & \gamma_j^* + \gamma_{p+1}^* \theta_j^* \in [\bar{\beta}_j - d\bar{S}_j, \bar{\beta}_j + d\bar{S}_j]. \end{aligned} \quad (17)$$

In a two-dimensional space, if we fix γ_{p+1}^* , this optimization problem is trying to find coordinates (γ_j^*, θ_j^*) between 2 parallel lines $\gamma_j^* + \gamma_{p+1}^* \theta_j^* = \bar{\beta}_j - d\bar{S}_j$ and $\gamma_j^* + \gamma_{p+1}^* \theta_j^* = \bar{\beta}_j + d\bar{S}_j$ such that it is closest to a point (γ_j, θ_j) . We can further translate this problem into a problem that drops a perpendicular from the point with coordinates (γ_j, θ_j) to the line with equation $\gamma_j + \gamma_{p+1}^* \theta_j = \bar{\beta}_j - d\bar{S}_j$ and another perpendicular to the line with equation $\gamma_j + \gamma_{p+1}^* \theta_j = \bar{\beta}_j + d\bar{S}_j$ and the foot of each of these 2 perpendiculars can be easily found. Figure 1B illustrates this two-dimensional optimization problem.

As a result, by fixing γ_{p+1}^* , the minimization in (16) can be divided into $p+1$ two-dimensional minimization problems and is solved analytically by reexpressing the solution γ_j^*, θ_j^* as functions of γ_{p+1}^* . After that, the entire minimization problem is reduced to a simple one-dimensional optimization problem in γ_{p+1}^* , which can be easily solved using a one-dimensional optimization method. This iterative conditional optimization procedure is very fast. With $p = 5$ and 500 datasets, on a Mac laptop with 1.6-GHz processor, the computational time of the direct regression is about 2 seconds, 241 seconds for constrained ML, 4 seconds for the partial regression, 235 seconds for standard Bayes, 13 865 seconds for informative full Bayes, and 775 seconds for Bayesian transformation approach.

3 | SIMULATION STUDY

We present two simulation scenarios. In the first simulation scenario, the estimates of β and their standard errors are provided from an analysis of a large dataset of size 2000. The current dataset from which to estimate models $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ and $E(\mathbf{B}|\mathbf{X})$ is of size 15. Five hundred datasets are generated. The first simulation study generates data from a true model of the form $Y_i = \mu(\mathbf{X}_i, B_i) + \varepsilon_i$, where $\varepsilon_i \sim N(0, 6.5^2)$, $i = 1, \dots, 15$. $\mu(\mathbf{X}, \mathbf{B}) = 4 + 3\mathbf{X}_1 + 3\mathbf{X}_2 + 2\mathbf{B}$. $\mathbf{X}_1, \mathbf{X}_2 \stackrel{\text{i.i.d.}}{\sim} N(0, 1^2)$ and \mathbf{B} is simulated as $\mathbf{B} = 0.8\mathbf{X}_1 + 0.8\mathbf{X}_2 + N(0, 1.5^2)$. A linear regression based on the large dataset gives estimates for model $E(\mathbf{Y}|\mathbf{X}) = \beta_0 + \beta_1\mathbf{X}_1 + \beta_2\mathbf{X}_2$. The estimates and standard errors from this fit are $\bar{\beta}_0 = 4, \bar{S}_0 = 0.16, \bar{\beta}_1 = 4.6, \bar{S}_1 = 0.16, \bar{\beta}_2 = 4.6, \bar{S}_2 = 0.17$.

In the second simulation scenario, the dataset from which to estimate models $E(\mathbf{Y}|\mathbf{X}, \mathbf{B})$ and $E(\mathbf{B}|\mathbf{X})$ is of size 20. Five hundred datasets are generated. We generate data from a model with a larger number of covariates, $Y_i = \mu(\mathbf{X}_i, B_i) + \varepsilon_i$, where $\varepsilon_i \sim N(0, 6^2)$, $i = 1, \dots, 20$. $\mu(\mathbf{X}, \mathbf{B}) = 4 + 3\mathbf{X}_1 + 3\mathbf{X}_2 + 2\mathbf{X}_3 + 2\mathbf{X}_4 + 2\mathbf{B}$. $\mathbf{X}_1, \mathbf{X}_2 \stackrel{\text{i.i.d.}}{\sim} N(0, 1^2)$ and $\mathbf{X}_3, \mathbf{X}_4 \stackrel{\text{i.i.d.}}{\sim} N(0, 1.5^2)$ and \mathbf{B} is simulated as $\mathbf{B} = 0.8\mathbf{X}_1 + 0.8\mathbf{X}_2 + N(0, 1.5^2)$. A linear regression based on a large dataset of 2000 subjects gives estimates for model $E(\mathbf{Y}|\mathbf{X}) = \beta_0 + \beta_1\mathbf{X}_1 + \beta_2\mathbf{X}_2 + \beta_3\mathbf{X}_3 + \beta_4\mathbf{X}_4$. The estimates and SE's are $\bar{\beta}_0 = 4, \bar{S}_0 = 0.15, \bar{\beta}_1 = 4.6, \bar{S}_1 = 0.15, \bar{\beta}_2 = 4.6, \bar{S}_2 = 0.15, \bar{\beta}_3 = 2.1, \bar{S}_3 = 0.10, \bar{\beta}_4 = 1.9, \bar{S}_4 = 0.10$.

For comparing coefficient estimation, we report 4 quantities: the average of estimated coefficient, relative efficiency of estimated coefficient, mean squared error, and the average of the estimated standard error of the coefficient across 500 replicates. The average of estimated coefficient is defined as $\bar{\gamma}_j = \frac{1}{500} \sum_{m=1}^{500} \hat{\gamma}_{m,j}$, $j = 1, \dots, p+1$; the relative efficiency of estimated coefficient is defined as $V(\hat{\gamma}_{j,direct})/V(\hat{\gamma}_{j,method})$, where $V(\hat{\gamma}_j)$ is the Monte Carlo variance $\frac{1}{500} \sum_{m=1}^{500} (\hat{\gamma}_{m,j} - \bar{\gamma}_j)^2$ and the MSE of an estimated coefficient is defined as $\frac{1}{500} \sum_{m=1}^{500} (\hat{\gamma}_{m,j} - \gamma_j)^2$. We report the average of estimated standard errors of the estimates across 500 datasets (ie, $\frac{1}{500} \sum_{m=1}^{500} \sqrt{\hat{V}(\gamma_{m,j})}$) for each method: the average of the OLS estimated standard error for the direct regression estimates, the average of the asymptotic standard error for the partial regression estimates, the average of the bootstrap standard error for the constrained ML estimates, and the average posterior standard deviation for each of the 3 Bayes estimates. For the constrained ML solution, we also provide the bootstrap bias-corrected estimate.

For comparing prediction power across different methods, we calculate the average out-of-bag (OOB) R^2 for prediction error in a validation dataset of size 100: $\text{OOB } R^2 = 1 - \sum_{i=1}^{100} (Y_i - \sum_{j=0}^p \hat{\gamma}_j X_{ij} - \hat{\gamma}_{p+1} B_i)^2 / \sum_{i=1}^{100} (Y_i - \bar{Y})^2$.

Table 2 summarizes the simulation results for 3-covariate simulation scenario. The bootstrap corrected constrained ML, partial regression, and informative full Bayes give estimates of the regression coefficients with low bias. The OOB R^2 of \mathbf{Y} regressed on $\mathbf{X}_1, \mathbf{X}_2$ is low. By looking at relative efficiency of regression parameters γ_1 and γ_2 , we find the constrained methods greatly improve the estimation efficiency of coefficients of \mathbf{X} . For γ_1 and γ_2 , the partial regression, informative full Bayes and Bayesian transformation approach reduce the MSE by more than 50%. As expected, for γ_3 , these constrained estimates do not improve systematically from the unconstrained estimates in terms of MSE, as the historical data provide no information on the additional predicting variable \mathbf{B} . This finding agrees with the conclusion in Qin et al¹⁰ that there is large improvement in the coefficients of \mathbf{X} but not in the coefficient of \mathbf{B} . Among all the methods, the 2 constrained Bayesian methods, informative full Bayes and Bayesian transformation approach have highest prediction power. They lead to an increase of 41% and 36% in terms of OOB R^2 , respectively, compared to direct regression.

Table 3 summarizes the simulation results for the 5-covariate simulation scenario. The OOB R^2 of \mathbf{Y} regressed on $\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3, \mathbf{X}_4$ is moderate. By looking at relative efficiency of regression parameters $\gamma_1, \gamma_2, \gamma_3, \gamma_4$, this table again tells us that the constrained methods could substantially improve the estimation efficiency of coefficients of \mathbf{X} . In fact, the efficiency of coefficient γ_3 and that of γ_4 triple for the partial regression, informative full Bayes and the Bayesian transformation approach compared to direct regression. The informative full Bayes and the Bayesian transformation approach have the highest prediction power, as measured by OOB R^2 . They increase by 25% and 23% the OOB R^2 , respectively,

TABLE 2 Simulation results of 3-covariate scenario: comparison of different methods

Method	γ_1	γ_2	γ_3	OOBR ²
True value	3	3	2	
Direct regression	3.25(1.00)	3.07(1.00)	1.96(1.00)	0.270
MSE	5.21	5.90	1.92	
Avg.SE	2.20	2.23	1.31	
Constrained ML	2.82(1.59)	2.79(2.20)	2.27(0.80)	0.334
MSE	3.27	2.74	2.46	
Avg.Boot.SE	1.93	2.31	3.06	
Constrained ML _{bc}	3.01(1.40)	3.01(2.59)	2.00(0.88)	0.346
Partial regression	3.03(2.26)	3.01(2.62)	1.96(1.00)	0.346
MSE	2.29	2.25	1.92	
Avg.Asy.SE	1.58	1.56	1.34	
Standard Bayes	3.24(1.01)	3.06(1.00)	1.97(1.00)	0.270
MSE	5.24	5.93	1.93	
Avg.PSD	2.43	2.46	1.44	
Informative full Bayes	3.06(2.63)	2.99(2.97)	1.98(1.11)	0.382
MSE	1.97	1.98	1.74	
Avg.PSD	1.45	1.48	1.30	
Transformation	3.16(2.14)	3.09(2.31)	1.84(0.88)	0.366
MSE	2.40	2.56	2.22	
Avg.PSD	1.74	1.78	1.65	

OOB R^2 denotes average out-of-bag prediction ability. For each method, the first row includes mean (relative efficiency w.r.t. direct regression) of each regression coefficient and OOB R^2 of this method. The second row shows the MSE of each coefficient and the third row is the average of the standard error across 500 datasets. For constrained ML, we also report a bootstrap bias-corrected constrained ML estimate. A linear regression on Y on X_1, X_2 has an OOB R^2 of 0.212.

compared to direct regression. The constrained ML solution also increases by 23% the OOB R^2 when the bootstrap bias-correction procedure is applied.

We also conduct additional simulation studies with different sample sizes for the 3-covariate scenario and for the 5-covariate scenario. The simulation results are shown in Appendix S6. As expected, when the sample size increases, the gain in estimating efficiency and predictive power by incorporating the external information is not as significant as it is in small sample sizes settings.

4 | APPLICATION TO THE NORMATIVE AGING STUDY

We illustrate our methodology by enhancing a published prediction model for bone lead levels in terms of blood lead and other covariates,²⁰ with a new biomarker defined through a continuous genetic risk score (\mathbf{B} , in terms of the notation used in previous sections). It is known that up to 95% of the total body burden of lead is accumulated in the skeleton.²¹ While blood lead is widely used as a biomarker of recent lead exposure due to the convenience of collecting blood samples, its short half-life (~ 30 days) limits its utility in chronic disease epidemiology research. Therefore, bone lead reflects cumulative lead exposure and is considered a better biomarker when examining chronic diseases. Recent development of K X-ray fluorescence instruments makes it possible to take direct measurements of bone lead concentrations.²² However, K X-ray fluorescence measurements can only be taken at very few locations in the entire country and thus direct bone lead level measurements are not commonly available. Thus, prediction models of bone lead level were constructed in terms of blood lead levels and other covariates that are more readily available in other studies.

Our study uses data from the Normative Aging Study, a longitudinal study established by the Veterans Administration in 1961. The Normative Aging Study enrolled 2280 men, aged 21 to 80, living in the Greater Boston area. Participants were recruited to represent a range of socioeconomic characteristics in terms of education and occupation.²³ Every 3 to 5 years,

TABLE 3 Simulation results of 5-covariate scenario: comparison of different methods

Method	γ_1	γ_2	γ_3	γ_4	γ_5	OOB R^2
True value	3	3	2	2	2	
Direct regression	3.11(1.00)	3.02(1.00)	1.93(1.00)	2.04(1.00)	1.92(1.00)	0.421
MSE	3.54	4.28	1.32	1.33	1.22	
Avg.SE	1.84	1.85	1.08	1.08	1.07	
Constrained ML	2.77(1.91)	2.66(1.67)	2.03(2.70)	1.94(2.35)	2.33(0.76)	0.492
MSE	1.91	2.67	0.49	0.56	1.69	
Avg.Boot.SE	1.54	1.59	0.74	0.73	1.76	
Constrained ML_{bc}	3.07(2.50)	2.97(2.19)	2.04(3.67)	1.94(3.13)	1.95(0.98)	0.519
Partial regression	3.08(2.41)	3.01(2.28)	2.09(3.23)	1.91(3.13)	1.92(1.00)	0.500
MSE	1.48	1.89	0.42	0.43	1.22	
Avg.Asy.SE	1.28	1.29	0.62	0.61	1.11	
Standard Bayes	3.11(1.00)	3.02(1.00)	1.93(1.00)	2.04(1.00)	1.92(1.00)	0.421
MSE	3.55	4.28	1.32	1.33	1.22	
Avg.PSD	1.98	2.00	1.17	1.16	1.16	
Informative full Bayes	3.02(2.92)	2.93(2.54)	2.05(3.93)	1.92(3.33)	2.00(1.26)	0.526
MSE	1.24	1.72	0.34	0.40	0.96	
Avg.PSD	1.14	1.17	0.56	0.56	1.01	
Transformation	3.16(2.34)	3.08(2.07)	2.04(3.44)	1.95(3.23)	1.81(0.81)	0.516
MSE	1.52	2.09	0.38	0.41	1.51	
Avg.PSD	1.42	1.46	0.67	0.66	1.41	

OOB R^2 denotes average out-of-bag prediction ability. For each method, the first row includes mean (relative efficiency w.r.t. direct regression) of each regression coefficient and OOB R^2 of this method. The second row shows the MSE of each coefficient and the third row is the average of the standard error across 500 datasets. For constrained ML, we also report a bootstrap bias-corrected constrained ML estimate. A linear regression on Y on X_1, X_2, X_3, X_4 has an OOB R^2 of 0.350.

participants returned for follow-up visits and information about age, smoking, education level, disease status, medication use, physical activity, and dietary intake was recorded. Beginning in 1991, K X-ray fluorescence was used to measure bone lead levels of participants at 2 sites: tibia (representing cortical bone) and patella (representing trabecular bone).

Park et al²⁰ developed a prediction model for tibia lead level using blood lead levels, age, smoking status, pack-years of cigarette, education, and occupation based on 550 participants of the Normative Aging Study. These 6 predictors were selected because they could be routinely collected in epidemiological studies. Table 4 shows the estimated tibia lead prediction model presented in their paper.

Park et al²⁰ commented that bone lead levels differ by genetic makeup and including some relevant genetic polymorphisms to the existing model may provide improved prediction accuracy. We want to use the published tibia lead prediction model as external information and see if genes in the lead toxicokinetics and toxicodynamics pathway can enhance the prediction power. We use the data from the Normative Aging study that not only has bone lead levels and these 6 covariates but also has 19 single-nucleotide polymorphisms (SNPs) relevant to the lead toxicokinetics and toxicodynamics pathway. We would like to include in the model a composite genetic risk score based on the unweighted sum of the risk allele counts of the 19 SNPs.

We exclude those individuals that have missingness in more than 3 SNPs. The remaining missing values in each genotype are imputed by the average number of risk alleles. The composite genetic risk score is then constructed as the summation of the risk allele counts of these 19 relevant SNPs. The genetic risk score is noted to be roughly normally distributed. We also remove those individuals who have missing values in any of the 6 predictors or the response tibia lead level. Our dataset then consists of 156 observations, including first measurements of 100 participants used as a training dataset and follow-up measurements of another 56 participants used as a testing dataset. These 2 datasets are independent of each other and our training dataset is independent of the original training data of Park et al.²⁰

Table 3 in Appendix S4 shows the characteristics of the training dataset and the characteristics of the testing dataset. There are no significant differences in variables age, pack-years of cigarette, genetic score, blood lead in these 2 datasets. There are apparent differences in variables smoking status and education. We would like to build the expanded tibia lead prediction model using our training dataset and the testing dataset will be used for validating our model.

TABLE 4 Regression coefficients externally imported from the tibia lead prediction model ($n = 550$) in Park et al (2009) and regression coefficients of this tibia prediction model estimated based on our training dataset ($n = 100$)

Variable	Tibia lead prediction model ($n=550$)			Our training dataset ($n=100$)		
	β	SE	P	β	SE	P
Intercept	-20.27	5.34	.0002	-22.20	9.80	.03
Bone lead	1.03	0.13	< .0001	0.93	0.30	.002
Age	0.59	0.07	< .0001	0.55	0.13	< .0001
Education						
High school diploma	-3.65	1.75	.04	3.21	3.42	.35
≥ 4 yr of college	-7.05	2.09	.0008	0.02	4.05	1.00
White collar	-3.21	1.18	.01	-4.46	2.40	.07
Cumulative cigarette smoking (pack-yr)	0.04	0.03	.17	0.22	0.06	.0005
Smoking status						
Former smoker	1.80	1.34	.18	-1.96	2.52	.44
Current smoker	0.05	2.48	.98	-19.22	5.43	.0006
R^2	0.27			0.42		

The expanded tibia lead prediction model will be estimated by both the unconstrained methods and the constrained methods we described in Section 2.

For comparing coefficient estimation across different methods, we report the estimated coefficients and their standard errors. For comparing prediction power, we calculate R^2 in the training dataset and OOB R^2 in the testing dataset. We also estimate a 6-predictor model without SNP information as in Park et al²⁰ in our training dataset and find out that the R^2 and OOB R^2 of this tibia lead model without SNP information are 0.42 and 0.16, respectively. The estimated coefficients are shown in Table 4.

Table 5 presents the expanded tibia lead prediction model fitted to the training dataset. We find that while the R^2 of this model does not increase much comparing to that of the prediction model without SNP information, the OOB R^2 increases 88% if we incorporate external information from Table 4 into our model estimating procedure. If we compare the standard errors across different methods, it is easily seen that the constrained methods will reduce the standard errors of regression coefficients compared to direct regression. For the constrained ML, informative full Bayes and the Bayesian transformation approach, the standard errors of the parameters of the variables blood lead, age, education, white collar, pack-years of cigarette, and smoking status decrease at least 50% comparing to the standard errors in direct regression. Meanwhile, partial regression estimates of parameters of variables blood lead, age, education, white collar, pack-years of cigarette, and smoking status have more than 80% reduction in standard errors comparing to direct regression. Therefore, it could be easier to identify statistically significant predictors based on these constrained methods. The reason that the partial regression estimates have the smallest estimated standard errors among all constrained solutions is that the standard errors of the regression coefficients from the tibia lead prediction model in Park et al²⁰ have different scales. By using partial regression method, we assume that there is no variation in the estimated coefficients in this tibia lead model (we only plug in the point estimates and do not make use of the standard errors) while the other three constrained solutions take into account the precision in these estimates.

5 | DISCUSSION

In this study, we demonstrate how to incorporate external information on regression coefficients in linear regression model estimation and prediction. We formulate the problem in an inferential framework in which the historical information is translated into nonlinear inequality constraints for coefficients and propose 4 constrained solutions: constrained ML, partial regression, informative full Bayes, and Bayesian transformation approach. We use simulation studies to assess the performance of these proposed methods and show that incorporating external information can improve the efficiency of model estimation and increase the prediction accuracy. The application to the Normative Aging Study shows that the estimation accuracy of regression coefficients and the predictive power of a tibia lead prediction model that includes a

TABLE 5 Regression coefficients of the expanded tibia lead prediction model with the genetic score (n = 100)

Variable	Direct regression	Standard bayes	Constrained ML	Partial regression	Informative full bayes	Bayesian transformation approach	Constrained ML_{bc}
Intercept	-23.85(11.27) ^a	-23.40(11.53)	-27.47(7.27)	-21.77(5.83)	-25.54(5.23)	-20.56(8.82)	-30.98(-46.71,-16.80) ^b
Blood lead	0.95(0.30)	0.95(0.31)	0.93(0.13)	1.05(0.06)	1.06(0.13)	1.00(0.14)	0.85(0.71, 1.17)
Age	0.55(0.13)	0.55(0.14)	0.64(0.06)	0.59(0.01)	0.60(0.05)	0.58(0.06)	0.68(0.52, 0.72)
Education							
High school diploma	3.19(3.44)	3.11(3.44)	-1.92(1.42)	-3.68(0.15)	-3.13(1.40)	-2.05(1.24)	-0.47(-3.47, 0.33)
≥ 4 yr of college	-0.06(4.08)	-0.11(4.12)	-5.04(1.60)	-7.13(0.33)	-6.62(1.64)	-5.22(1.56)	-3.66(-7.40,-2.62)
White collar	-4.48(2.42)	-4.48(2.45)	-2.84(1.02)	-3.24(0.11)	-3.07(1.03)	-3.61(0.98)	-2.77(-4.98,-1.77)
Cumulative cigarette smoking	0.22(0.06)	0.22(0.06)	0.06(0.02)	0.03(0.005)	0.04(0.02)	0.06(0.02)	0.08(0.02, 0.10)
Smoking status							
Former smoker	-1.92(2.54)	-1.95(2.56)	2.89(0.98)	1.84(0.19)	2.27(1.20)	0.70(1.00)	3.25(0.29, 3.74)
Current smoker	-19.07(5.48)	-19.16(5.60)	-2.28(1.90)	0.20(0.61)	-2.44(2.17)	0.19(2.54)	-3.51(-4.65, 1.90)
Genetic risk score	0.13(0.44)	0.12(0.45)	0.15(0.48)	0.13(0.45)	0.36(0.31)	-0.10(0.67)	0.13(-0.76, 1.10)
R ²	0.42	0.42	0.36	0.32	0.35	0.30	0.37
OOB R ²	0.17	0.17	0.29	0.29	0.30	0.28	0.28

^astandard error. ^b95% confidence interval.

composite genetic risk score as a new biomarker can be improved, when information from a previously published model with non-genetic data is incorporated.

Among the constrained solutions, our Bayesian transformation approach, motivated by Gunn and Dunson's transformation method, is a simple and effective computation method. The main underlying idea for the Bayesian transformation approach is to first obtain rapid draws from a simple sampling algorithm ignoring the constraints and then transform the draws by minimizing the squared normalized Euclidean distance between unconstrained draws and constrained draws, subject to these constraints.

One point of future consideration is the choice of tuning parameter. We have a quantity labeled d that controls the degree of trust in the historical information and we select it by drawing from a half normal distribution $|N(0, \sigma_d^2)|$ in Bayesian transformation approach and we fix $d = 1$ in the constrained ML approach. Although we fix $\sigma_d^2 = 1$ or $d = 1$, it can also be considered as a tuning parameter and adaptively selected for a particular dataset. Using simulation studies, we tried to use cross validation to conduct a grid search of d by choosing the optimal d that gives the lowest cross validation prediction error. However, this choice of d did not have optimal performance in terms of the accuracy and prediction metrics we evaluated (MSE, OOB R^2). Thus, how to select this tuning parameter in a principled optimal manner is not justified yet. In the Supporting Information, we show how the value of d will affect the constrained ML estimates. With bigger d , the standard errors of the regression coefficients will increase, and the predictive power in the validation dataset will decrease. When $d = 10$, these box constraints will be very weak and the estimated model based on the constrained ML is very similar to the estimated model based on direct regression, although the standard errors are smaller.

Another challenge is to demonstrate whether our Bayesian transformation approach provides a good approximation to the informative full Bayes solution. We claim to obtain approximate posterior draws from the Bayesian transformation approach without writing down the posterior distribution functions while informative full Bayes does give the exact posterior distribution. To validate that the Bayesian transformation approach is a good approximation to the true posterior, we compare the posterior distribution to that obtained from informative full Bayes in simulation studies and find good correspondence (results not shown). But this empirical observation needs further justification.

Among the constrained solutions, the constrained maximum likelihood estimate method and the informative full Bayes approach both depend on the likelihood of $\mathbf{Y}|\mathbf{X}, \mathbf{B}$ and $\mathbf{B}|\mathbf{X}$. In this manuscript, we primarily discuss the situation when $\mathbf{Y}|\mathbf{X}, \mathbf{B}$ and $\mathbf{B}|\mathbf{X}$ are both normal. However, if one or both $\mathbf{Y}|\mathbf{X}, \mathbf{B}$ or $\mathbf{B}|\mathbf{X}$ is not normal, but still continuous, the joint likelihood function can be modified and these 2 constrained solutions can be directly extended. To perform standard Bayesian inference to produce the initial raw draws, Bayesian transformation approach also depends on the likelihood function of $\mathbf{Y}|\mathbf{X}, \mathbf{B}$ and $\mathbf{B}|\mathbf{X}$ which can again be modified to a nonnormal likelihood if needed. The simulation study and the data analysis are based on a single continuous variable \mathbf{B} . However, these strategies to incorporating the external coefficient information can be extended if \mathbf{B} is multivariate. The exact relationship between parameters shown in Equation 6 can be extended to the case that \mathbf{B} is multivariate normal with L dimensions:

$$\beta_j = \gamma_j + \sum_{l=1}^L \gamma_{p+l} \theta_{lj}, j = 0, \dots, p. \quad (18)$$

As a consequence, when \mathbf{B} is multivariate normal, the constrained ML, partial regression, informative full Bayes, and Bayesian transformation approach are still applicable.

We consider prediction models to predict a continuous outcome. In future work, we will consider predictions of a binary outcome, which are also common, particularly in medical applications, where logistic regression models are frequently used for predicting the risk of a binary disease indicator.

6 | SOFTWARE

Software in the form of R code, together with a sample input data set and complete documentation, is available on request from the corresponding author (chengwt@umich.edu).

ACKNOWLEDGEMENTS

This research was partially supported by the NSF grant DMS 1406712 and NIH grants CA129102 and ES20811. The VA Normative Aging Study is supported by the Cooperative Studies Program/Epidemiology Research and Information Center of the US Department of Veterans Affairs and is a component of the Massachusetts Veterans Epidemiology Research

and Information Center, Boston, Massachusetts. The views expressed in this paper are those of the authors and do not necessarily represent the views of the US Department of Veterans Affairs.

CONFLICT OF INTEREST

None declared.

ORCID

Wenting Cheng  <http://orcid.org/0000-0003-2199-1467>

Sung Kyun Park  <http://orcid.org/0000-0001-9981-6250>

REFERENCES

1. D'Agostino RB, Grundy S, Sullivan LM, Wilson P, the CHD Risk Prediction Group. Validation of the framingham coronary heart disease prediction scores: results of a multiple ethnic groups investigation. *The J Am Med Assoc*. 2001;286(2):180-187.
2. Thompson IM, Ankerst DP, Chi C, et al. Assessing prostate cancer risk: results from the prostate cancer prevention trial. *J Nat Cancer Inst*. 2006;98(8):529-534.
3. Gail MH, Brinton LA, Byar DP, et al. Projecting individualized probabilities of developing breast cancer for white females who are being examined annually. *J Nat Cancer Inst*. 1989;81(24):1879-1886.
4. Geweke J. Exact inference in the inequality constrained normal linear regression model. *J Appl Econometrics*. 1986;1(2):127-141.
5. Dunson DB, Neelon B. Bayesian inference on order-constrained parameters in generalized linear models. *Biometrics*. 2003;59(2):286-295.
6. Gunn LH, Dunson DB. A transformation approach for incorporating monotone or unimodal constraints. *Biostatistics*. 2005;6(3):434-449.
7. Chen A, Owen AB, Shi M. Data enriched linear regression. *Electron J Stat*. 2015;9(1):1078-1112.
8. Imbens GW, Lancaster T. Combining micro and macro data in microeconomic models. *The Rev Econ Stud*. 1994;61(4):655-680.
9. Qin J. Combining parametric and empirical likelihoods. *Biometrika*. 2000;87(2):484-490.
10. Qin J, Zhang H, Li P, Albanes D, Yu K. Using covariate-specific disease prevalence information to increase the power of case-control studies. *Biometrika*. 2015;102(1):169-180.
11. Chatterjee N, Chen Y-H, Maas P, Carroll RJ. Constrained maximum likelihood estimation for model calibration using summary-level information from external big data sources. *J Am Stat Assoc*. 2016;111(513):107-117.
12. Carlin BP, Louis TA. *Bayesian Methods for Data Analysis*. Boca Raton: CRC Press; 2009.
13. Lesaffre E, Lawson AB. Choosing the prior distribution. In: Lesaffre E, Lawson AB, eds. *Bayesian Biostatistics*. West Sussex: John Wiley and Sons, Ltd; 2012:104-138.
14. Efron B, Tibshirani R. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Stat Sci*. 1986;1(1):54-75.
15. Efron B. Nonparametric standard errors and confidence intervals. *The Can J Stat / La Revue Canadienne de Statistique*. 1981;9(2):139-158.
16. Carpenter J, Bithell J. Bootstrap confidence intervals: when, which, what? a practical guide for medical statisticians. *Stat Med*. 2000;19(9):1141-1164.
17. Abdi H. Partial regression coefficients. In: Lewis-Beck MS, Bryman A, Liao TF, eds. *The Encyclopedia of Social Science Research Methods*. CA: Sage Publications, Inc.; 2004:796-798.
18. Roberts GO. Markov chain concepts related to sampling algorithms. In: Gilks WR, Richardson S, Spiegelhalter D, eds. *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall; 1996:45-54.
19. Gilks WR, Roberts GO. Strategies for improving mcmc. In: Gilks WR, Richardson S, Spiegelhalter D, eds. *Markov chain monte carlo in practice*. London: Chapman and Hall; 1996:89-110.
20. Park SK, Mukherjee B, Xia X, et al. Bone lead level prediction models and their application to examining the relationship of lead exposure and hypertension in the third national health and nutrition examination survey (nhanes-iii). *J Occup Environ Med / Am College Occup Environ Med*. 2009;51(12):1422-1436.
21. Barry PSI, Mossman DB. Lead concentrations in human tissues. *Br J Ind Med*. 1970;27(4):339-351.
22. Hu H, Shih R, Rothenberg S, Schwartz BS. The epidemiology of lead toxicity in adults: measuring dose and consideration of other methodologic issues. *Environ Health Perspectives*. 2007;115(3):455-462.
23. Bell B, Rose CL, Damon A. The normative aging study: an interdisciplinary and longitudinal study of health and aging. *The Int J Aging Human Dev*. 1972;3(1):5-17.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Cheng W, Taylor JMG, Vokonas PS, Park SK, Mukherjee B. Improving estimation and prediction in linear regression incorporating external information from an established reduced model. *Statistics in Medicine*. 2018;37:1515–1530. <https://doi.org/10.1002/sim.7600>