

BAYESIAN SHRINKAGE METHODS FOR PARTIALLY OBSERVED DATA WITH MANY PREDICTORS¹

BY PHILIP S. BOONSTRA, BHRAMAR MUKHERJEE
AND JEREMY M. G. TAYLOR

University of Michigan

Motivated by the increasing use of and rapid changes in array technologies, we consider the prediction problem of fitting a linear regression relating a continuous outcome Y to a large number of covariates \mathbf{X} , for example, measurements from current, state-of-the-art technology. For most of the samples, only the outcome Y and surrogate covariates, \mathbf{W} , are available. These surrogates may be data from prior studies using older technologies. Owing to the dimension of the problem and the large fraction of missing information, a critical issue is appropriate shrinkage of model parameters for an optimal bias-variance trade-off. We discuss a variety of fully Bayesian and Empirical Bayes algorithms which account for uncertainty in the missing data and adaptively shrink parameter estimates for superior prediction. These methods are evaluated via a comprehensive simulation study. In addition, we apply our methods to a lung cancer data set, predicting survival time (Y) using qRT-PCR (\mathbf{X}) and microarray (\mathbf{W}) measurements.

1. Introduction. The ongoing development of array technologies for assaying genomic information has resulted in an abundance of data sets with many predictors and presents both statistical opportunities and challenges. As an example, Chen et al. (2011) analyzed a gene-expression microarray data set of 439 lung adenocarcinomas from four cancer centers in the United States, with the goal of using gene expression to improve predictions of survival time relative to using clinical covariates alone. Expression was measured using Affymetrix oligonucleotide microarray technology. After pre-screening the probes for consistency between centers, the authors initially evaluated 13,306 probes for construction of their predictor.

A clinical challenge to a candidate model which uses Affymetrix data is its application for predictions in new patients. The underlying complexity of Affymetrix data, including necessary preprocessing, requires specialized laboratory facilities, which will be locally unavailable at smaller hospitals. On the other hand, quantitative real-time polymerase chain reaction (qRT-PCR) offers a faster and more

Received September 2012; revised June 2013.

¹Supported by NSF Grant DMS-10-07494 and the National Institutes of Health [CA129102, CA156608, ES020811].

Key words and phrases. High-dimensional data, Markov chain Monte Carlo, missing data, measurement error, shrinkage.

efficient assay of the same underlying genomic information, making a qRT-PCR-based prediction model clinically applicable. The trade-off comes from the limited number of genes which may be assayed on a single qRT-PCR card. Thus, from the Affymetrix data, 91 promising genes were first identified. These 91 genes were then re-assayed with qRT-PCR. Because of tissue availability issues owing to the multi-center-nature of the study, only 47 out of 439 tumors were re-assayed by qRT-PCR, creating a significant missing data problem.

Motivated by this problem, in this paper we consider the analysis of a data set with many predictors in which a large block of covariates are missing, a situation for which there is limited previous literature. To maintain relevance to the application which drives our methodology, we assume the data have two distinctive features. First, the number of covariates, that is, genes, is of moderate size, approximately the same order as the number of observations. This precludes both a more traditional regression situation as well as an “ultra-high-dimensional” regression and reflects that an initial screening has identified a subset of potentially informative genes. Second, there are two versions of the genomic data: measurements from a prior technology, which are complete for all observations, and measurements from a newer, more efficient technology, which are observed only on a small subset of the observations. Owing to the inherent variability in parameter estimates induced by both the missing data and the dimensionality of the problem, we consider Bayesian approaches, which allow for the application of shrinkage methods, in turn offering better prediction.

Translating this into statistical terminology, we consider predicting an outcome Y given length- p covariates \mathbf{X} . Assuming Y is continuous and fully observed, we use the linear model

$$(1) \quad Y = \beta_0 + \mathbf{X}^\top \boldsymbol{\beta} + \sigma \varepsilon.$$

All observations contain Y and \mathbf{W} , which is an error-prone length- p surrogate for the true covariate \mathbf{X} . On a small number of observations of size n_A , subsample A, we also observe \mathbf{X} , which is missing for the remaining subjects, constituting subsample B, of size n_B . Complete observations, then, contain an outcome Y , covariates \mathbf{X} and surrogates \mathbf{W} . Subsample A is written as $\{y_A, \mathbf{x}_A, \mathbf{w}_A\}$ and subsample B as $\{y_B, \mathbf{w}_B\}$. The true covariates from subsample B, \mathbf{x}_B , are unmeasured. The data are schematically presented in Figure S1 of the supplemental article [Boonstra, Mukherjee and Taylor (2013)].

Our goal is a predictive model for $Y|\mathbf{X}$ as in equation (1), but because \mathbf{W} is correlated with \mathbf{X} , subsample B contains information about $\boldsymbol{\beta}$. Moreover, shrinkage of regression coefficients may alleviate problems associated with multicollinearity of covariates. Boonstra, Taylor and Mukherjee (2013) proposed a class of targeted ridge (TR) estimators of $\boldsymbol{\beta}$, shrinking estimates toward a target constructed using subsample B, making a bias-variance trade-off. The amount of shrinkage can be data-adaptive with a tuning parameter, say, λ . In a simulation study of data sets

with many predictors, they showed that two biased methods, a modified regression calibration algorithm and a “hybrid” estimator, which is a linear combination of multiple TR estimators with data-adaptive weights, uniformly out-perform standard regression calibration, an unbiased method, in terms of mean-squared prediction error (MSPE):

$$\begin{aligned}
 & \text{MSPE}(\hat{\beta}_0, \hat{\beta}) \\
 (2) \quad & = E[(Y_{\text{new}} - \hat{\beta}_0 - \mathbf{X}_{\text{new}}^\top \hat{\beta})^2] \\
 & = \sigma^2 + (E[\beta_0 - \hat{\beta}_0 + \mathbf{X}_{\text{new}}^\top \beta - \mathbf{X}_{\text{new}}^\top \hat{\beta}])^2 + \text{Var}[\hat{\beta}_0 + \mathbf{X}_{\text{new}}^\top \hat{\beta}],
 \end{aligned}$$

where the expectation is over $Y_{\text{new}}, \mathbf{X}_{\text{new}}, \mathbf{y}_A, \mathbf{y}_B | \mathbf{x}_A, \mathbf{w}_A, \mathbf{w}_B$.

However, there are reasons to consider alternative strategies. The authors showed the TR estimator can be viewed as a missing data technique: make an imputation $\tilde{\mathbf{x}}_B$ of the missing \mathbf{x}_B and calculate $\hat{\beta}$ treating the data as complete. When the shrinkage is data-adaptive through the tuning parameter λ , there is an intermediate stage: choose λ given $\tilde{\mathbf{x}}_B$. Uncertainty in $\tilde{\mathbf{x}}_B$ or λ is not propagated in the TR estimators, thus, it can be viewed as improper imputation [Little and Rubin (2002)]. Moreover, to choose λ , a generalized cross-validation (GCV) criterion was applied to subsample A. Although GCV asymptotically chooses the optimal value of λ [Craven and Wahba (1979)], it can overfit in finite sample sizes, and an approach for estimating λ which also uses information in subsample B is preferred. Finally, constructing prediction intervals corresponding to the point-wise predictions generated by the class of TR estimators requires use of the bootstrap. This resampling process is computationally intensive and provides coverage that may not be nominal.

These reasons, that is, characterizing prediction uncertainty and unifying shrinkage, imputation of missing data and an adaptive choice of λ , motivate a fully Bayesian approach toward the same goal of improving predictions using auxiliary data. Consider the generic hierarchical model presented in Figure 1. Known (unknown, resp.) quantities are bounded by square (circular) nodes. Instead of splitting the data into subsamples (cf. Figure S1), we classify it more broadly into observed (\mathbf{U}^{obs}) and missing (\mathbf{U}^{mis}) components. Let ϕ denote parameters of interest and nuisance parameters in the underlying joint likelihood of $\{\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}}\}$.




FIG. 1. A hierarchical model with missing data \mathbf{U}^{mis} and observed data \mathbf{U}^{obs} . The shrinkage penalty parameters η are the hyperparameters of ϕ , the quantity(ies) of primary interest.

Regularization of ϕ is achieved through the shrinkage parameter η , equivalently interpreted in Figure 1 as the hyperparameters which index a prior distribution on ϕ . One can impose another level of hierarchy through a hyperprior distribution on η . Using $[\cdot]$ and $[\cdot|\cdot]$ to denote marginal and conditional distributions, draws from $[\mathbf{U}^{\text{mis}}, \phi, \eta | \mathbf{U}^{\text{obs}}]$, the distribution of unknown random quantities conditional on the observed data, constitute proper imputation and incorporate all of the information in the data. Summary values, like posterior means, as well as measurements of uncertainty, like highest posterior density credible intervals and prediction intervals, can easily be calculated based on posterior draws.

Placing the shrinkage parameter η in a hierarchical framework allows the flexibility to determine both which components of ϕ to shrink and to what extent. As an example of the former, [Boonstra, Taylor and Mukherjee \(2013\)](#) shrink estimates of the regression coefficients β , tuned by the parameter λ . However, for improved prediction of the outcome Y , it may be beneficial to shrink the parameters generating the missing data \mathbf{x}_B . For example, in a nonmissing-data context, the SCOUT method [[Witten and Tibshirani \(2009\)](#)] shrinks the estimate of $\text{Var}(\mathbf{X})$ for better prediction. As for the extent of shrinkage, the hyperparameter-equivalence of the tuning parameters allows for the use of Empirical Bayes algorithms to estimate η . This has been used in the Bayesian Lasso [[Park and Casella \(2008\)](#), [Yi and Xu \(2008\)](#)].

This paper makes two primary contributions. First, in Section 2 we discuss variants of the Gibbs sampler [[Geman and Geman \(1984\)](#)], a key algorithm for fitting hierarchical models with missing data. Here, we keep the context broad, assuming a generic hierarchical model indexed by ϕ with missing data \mathbf{U}^{mis} and unspecified hyperparameters η , as in Figure 1. One variant, Data Augmentation [[Tanner and Wong \(1987\)](#)], is a standard Bayesian approach to missing data, and all unknown quantities have prior distributions. Two others are Empirical Bayes methods: the Monte Carlo expectation–maximization algorithm [[Wei and Tanner \(1990\)](#)] and the Empirical Bayes Gibbs sampler [[Casella \(2001\)](#)]. Although proposed for seemingly different problems, we argue that the sampling strategies in each are special cases of that in Figure 1: variants of the same general algorithm, which we call EM-within-Gibbs. This previously-unrecognized link is important, given the increasing role Empirical Bayes methods play in modern applications. The second primary contribution builds on this proposed framework (Section 3), namely, a comparison of several fully Bayesian and Empirical Bayes options and their application to our motivating genomic analysis. Of note in the data are two crucial features: first, ϕ , comprised of β_0 , β , σ^2 plus parameters for modeling the distribution of \mathbf{X} , is of a significant dimension, so that fitting a model with *no* missing data would still be somewhat challenging, and, second, the number of partial observations where \mathbf{X} is missing is larger than the number of complete observations. Meaningful analysis then requires the regularization, or shrinkage, of ϕ via an appropriate specification of the hierarchy and choice of η . We propose

to shrink several different components of ϕ , making use of the simultaneous interpretation of η as a shrinkage penalty and a hyperparameter on ϕ . We evaluate these methods via a comprehensive simulation study (Section 4), also considering robustness of these methods under model misspecification. Finally, we analyze the Chen et al. data set (Section 5). We include ridge regression [Hoerl and Kennard (1970)] as a reference, because the additional modeling assumptions of the other likelihood-based methods offer efficiency gains only when they are satisfied.

2. Gibbs sampler variants. In this section we discuss four existing variants of the Gibbs sampler relevant to our analysis. As we will argue, two of these are special cases of a more general variant, which we call ‘‘Empirical Bayes Within Gibbs’’ (EWiG), an equivalence that has not been established previously, leaving three distinct variants. We define a ‘‘variant’’ here as the characterization of a posterior distribution plus an algorithm for fitting the model. All variants are summarized in Table 1.

Data augmentation (DA+, DA) [Tanner and Wong (1987)].

$$\text{Posterior: } [\phi, \mathbf{U}^{\text{mis}}, \eta | \mathbf{U}^{\text{obs}}] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} | \phi] \times [\phi | \eta] \times [\eta]$$

Algorithm: at iteration i ,

$$\mathbf{U}^{\text{mis}(i)} \sim [\mathbf{U}^{\text{mis}} | \mathbf{U}^{\text{obs}}, \phi^{(i-1)}]$$

$$\phi^{(i)} \sim [\phi | \mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}(i)}, \eta^{(i-1)}]$$

$$\eta^{(i)} \sim [\eta | \phi^{(i)}].$$

These two variants are natural Bayesian treatments of missing data: \mathbf{U}^{mis} and ϕ are both unobserved random variables. In DA+, which is given above, the hyperparameters η are also unknown [Gelfand and Smith (1990)]. In DA, a value for η is chosen. In either case, draws of ϕ and \mathbf{U}^{mis} are sequentially made from their

TABLE 1

A comparison of the general form of the Gibbs sampler variants from Section 2 as they were originally proposed. Differences between posteriors depend on the presence of missing data \mathbf{U}^{mis} and whether the hyperparameters η are fully known. Differences in algorithms depend on how the lowest level of the hierarchy, which is unknown, is treated. In particular, MCEM differs from DA because it returns only an estimate of the posterior mode

Variant	Posterior	Prior on η
DA [Tanner and Wong (1987)]	$[\phi, \mathbf{U}^{\text{mis}} \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} \phi] \times [\phi \eta]$	No
DA+ [Gelfand and Smith (1990)]	$[\phi, \mathbf{U}^{\text{mis}}, \eta \mathbf{U}^{\text{obs}}] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} \phi] \times [\phi \eta] \times [\eta]$	Yes
MCEM [Wei and Tanner (1990)]	$[\phi, \mathbf{U}^{\text{mis}} \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} \phi] \times [\phi \eta]$	No
EBGS [Casella (2001)]	$[\phi \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}} \phi] \times [\phi \eta]$	No
EWiG	$[\phi, \mathbf{U}^{\text{mis}} \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} \phi] \times [\phi \eta]$	No

conditional posteriors. In **DA+** only, η is also sampled from its conditional posterior. Then, in either **DA** or **DA+**, the whole process is iterated. Tanner and Wong prove that iterations will eventually yield a draw from the true posterior distribution of interest, $[\phi, \mathbf{U}^{\text{mis}}, \eta | \mathbf{U}^{\text{obs}}]$ for **DA+** or $[\phi, \mathbf{U}^{\text{mis}} | \mathbf{U}^{\text{obs}}, \eta]$ for **DA**. The full conditional distribution $[\phi | \mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}}, \eta]$ may be difficult to specify. Suppose instead a set of partial conditional distributions is available, $[\phi_J | \phi_{(J)}, \mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}}, \eta]$, where the set of J 's forms a partition of the vector ϕ . Then under mild conditions, repeated iterative sampling from these partial conditional distributions will also yield draws from the true posterior [Geman and Geman (1984)].

Monte Carlo expectation–maximization (MCEM) [Wei and Tanner (1990)].

$$\text{Posterior: } [\phi, \mathbf{U}^{\text{mis}} | \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}} | \phi] \times [\phi | \eta]$$

Algorithm: at iteration i ,

for $k = 1, \dots, K$,

$$\mathbf{U}^{\text{mis}(i,k)} \sim [\mathbf{U}^{\text{mis}} | \mathbf{U}^{\text{obs}}, \phi^{(i-1)}]$$

$$\phi^{(i)} = \arg \max_{\phi} \frac{1}{K} \sum_{k=1}^K \ln[\phi | \mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}(i,k)}, \eta].$$

MCEM provides a point estimate of ϕ rather than an estimate of the posterior distribution, as with **DA/DA+**. It is a modification of the original EM algorithm [Dempster, Laird and Rubin (1977)], replacing an intractable expectation with a Monte Carlo average of multiple imputations. K draws of \mathbf{U}^{mis} are sampled conditional on the current value of $\phi : \phi^{(i-1)}$. The expected posterior is updated with a Monte Carlo average and maximized with respect to ϕ . When ϕ has a flat prior, as in the originally proposed **MCEM**, $\{\phi^{(i)}\}$ will converge to the maximum likelihood estimate (MLE) of ϕ . If an informative prior is specified through a particular choice of η , the sequence will converge to a *penalized* MLE [Green (1990)].

Empirical Bayes Gibbs sampling (EBGS) [Casella (2001)].

$$\text{Posterior: } [\phi | \mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}} | \phi] \times [\phi | \eta]$$

Algorithm: at iteration i ,

for $k = 1, \dots, K$,

$$\phi^{(i,k)} \sim [\phi | \mathbf{U}^{\text{obs}}, \eta^{(i-1)}]$$

$$\eta^{(i)} = \arg \max_{\eta} \frac{1}{K} \sum_{k=1}^K \ln[\phi^{(i,k)} | \eta].$$

EBGS allows the data to determine a value for the hyperparameter η . In the context of Casella, there are no missing data \mathbf{U}^{mis} . However, ϕ is considered missing for purposes of determining η : choose η which maximizes its *marginal* log-likelihood, $\ln[\mathbf{U}^{\text{obs}} | \eta]$. As in **MCEM**, an EM-type algorithm can maximize this intractable

log-likelihood. K draws of ϕ are made from the current estimate of its posterior, and η is updated by maximizing a Monte Carlo estimate of $E[\ln[\phi|\eta]]$, where the expectation is over the distribution $[\phi|\mathbf{U}^{\text{obs}}, \eta^{(i)}]$. This expected complete-data log-likelihood relates to the desired marginal log-likelihood as follows. First observe that

$$\begin{aligned} [\mathbf{U}^{\text{obs}}|\eta][\phi|\mathbf{U}^{\text{obs}}, \eta] &= [\phi|\eta][\mathbf{U}^{\text{obs}}|\phi, \eta] \\ &= [\phi|\eta][\mathbf{U}^{\text{obs}}|\phi]. \end{aligned}$$

Let $C = E[\ln[\mathbf{U}^{\text{obs}}|\phi]]$, which is constant with respect to η . Then,

$$\ln[\mathbf{U}^{\text{obs}}|\eta] = E[\ln[\phi|\eta]] - E[\ln[\phi|\mathbf{U}^{\text{obs}}, \eta]] + C.$$

Because $E[\ln[\phi|\mathbf{U}^{\text{obs}}, \eta]] \leq E[\ln[\phi|\mathbf{U}^{\text{obs}}, \eta^{(i)}]]$ for any η , we have the result that maximizing $E[\ln[\phi|\eta]]$ (or a Monte Carlo approximation thereof) over η will increase $\ln[\mathbf{U}^{\text{obs}}|\eta]$ and converge to a local maximum.

EM-within-Gibbs (EWiG).

Posterior: $[\phi, \mathbf{U}^{\text{mis}}|\mathbf{U}^{\text{obs}}, \eta] \propto [\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}}|\phi] \times [\phi|\eta]$

Algorithm: at iteration i ,

for $k = 1, \dots, K$,

$$\mathbf{U}^{\text{mis}(i,k)} \sim [\mathbf{U}^{\text{mis}}|\mathbf{U}^{\text{obs}}, \phi^{(i,k-1)}]$$

$$\phi^{(i,k)} \sim [\phi|\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}(i,k)}, \eta^{(i-1)}]$$

$$\eta^{(i)} = \arg \max_{\phi} \frac{1}{K} \sum_{k=1}^K \ln[\phi^{(i,k)}|\eta].$$

Importantly, both **MCEM** and **EBGS** allow the lowest level of the hierarchy to be adaptively determined by the data rather than chosen a priori. In **MCEM**, this lowest level is ϕ , and in **EBGS**, it is η . However, **MCEM** can be expanded in the presence of an unknown η by putting both \mathbf{U}^{mis} and ϕ into the imputation step, so ϕ is sampled rather than optimized. The maximization step determines η . This returns to the original goal of **DA+/DA**, which is determining the posterior distribution of ϕ . Equivalently, we can take the perspective of expanding **EBGS**: add an imputation step for sampling \mathbf{U}^{mis} , keeping the maximization step the same. As a result of this equivalence, expanding either **MCEM** or **EBGS** yields the same result, what we call **EWiG**, given above. Because η is unknown, the hierarchical model here is the same as that given in Figure 1.

In summary, we have asserted that **MCEM** and **EBGS** are special cases of **EWiG**, so there are three distinct variants which we apply to our problem in the following section: **DA**, **DA+**, and **EWiG**.

3. Specification of likelihood and priors. The discussion so far has been deliberately generic. We now specify a likelihood for our problem of interest, which in turn gives ϕ , and apply these Gibbs variants to several combinations of (i) choices of priors $[\phi|\eta]$ and (ii) values of the hyperparameter η . Translating the quantities in Figure 1 to our problem, we have $\mathbf{U}^{\text{obs}} = \{\mathbf{y}_A, \mathbf{y}_B, \mathbf{x}_A, \mathbf{w}_A, \mathbf{w}_B\}$ and $\mathbf{U}^{\text{mis}} = \mathbf{x}_B$. A commonly used factorization of the joint likelihood is $[Y, \mathbf{X}, \mathbf{W}] = [Y|\mathbf{X}][\mathbf{W}|\mathbf{X}][\mathbf{X}]$, which makes a conditional independence assumption $[Y|\mathbf{X}, \mathbf{W}] = [Y|\mathbf{X}]$. An alternative factorization is $[Y|\mathbf{X}][\mathbf{X}|\mathbf{W}]$, which we do not consider, as it is inconsistent with the application-driven measurement error structure of \mathbf{W} and \mathbf{X} . We make the following assumptions:

$$(3) \quad \begin{aligned} Y|\mathbf{X} &= N\{\beta_0 + \mathbf{X}^\top \boldsymbol{\beta}, \sigma^2\}, & \mathbf{W}|\mathbf{X} &= N_p\{\psi \mathbf{1}_p + \nu \mathbf{X}, \tau^2 \mathbf{I}_p\}, \\ \mathbf{X} &= N_p\{\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X\}. \end{aligned}$$

The likelihood has an outcome model relating Y to \mathbf{X} , a measurement error model relating the error-prone \mathbf{W} to \mathbf{X} , and a multivariate distribution for \mathbf{X} . Thus, $\phi = \{\beta_0, \boldsymbol{\beta}, \sigma, \psi, \nu, \tau, \boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X\}$, and η is described below. Of interest is prediction of a new value Y_{new} given \mathbf{X}_{new} , for example, $\hat{Y}_{\text{new}} = \beta_0^* + \mathbf{X}_{\text{new}}^\top \boldsymbol{\beta}^*$, where β_0^* and $\boldsymbol{\beta}^*$ are posterior summaries of β_0 and $\boldsymbol{\beta}$. Uncertainty is quantified using the empirical distribution of $\hat{Y}_{\text{new}}^{(t)} = \beta_0^{(t)} + \mathbf{X}_{\text{new}}^\top \boldsymbol{\beta}^{(t)} + \sigma^{2(t)} \varepsilon^{(t)}$, where $\{\beta_0^{(t)}, \boldsymbol{\beta}^{(t)}, \sigma^{2(t)}\}$ is the set of posterior draws and $\varepsilon^{(t)} \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$. If $\mathbf{x}_B = \mathbf{U}^{\text{mis}}$ were observed, the *complete* log-likelihood would be

$$(4) \quad \begin{aligned} \ell_C &= \ln[\mathbf{U}^{\text{obs}}, \mathbf{U}^{\text{mis}}|\phi] \\ &= \ln[\mathbf{y}_A|\mathbf{x}_A, \beta_0, \boldsymbol{\beta}, \sigma^2] + \ln[\mathbf{w}_A|\mathbf{x}_A, \psi, \nu, \tau^2] + \ln[\mathbf{x}_A|\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X] \\ &\quad + \ln[\mathbf{y}_B|\mathbf{x}_B, \beta_0, \boldsymbol{\beta}, \sigma^2] + \ln[\mathbf{w}_B|\mathbf{x}_B, \psi, \nu, \tau^2] + \ln[\mathbf{x}_B|\boldsymbol{\mu}_X, \boldsymbol{\Sigma}_X]. \end{aligned}$$

The log-likelihood gives the imputation step:

$$(5) \quad \mathbf{x}_B|\mathbf{U}^{\text{obs}}, \phi = N_{n_B \times p}\{\tilde{\mathbf{x}}_B(\mathbf{U}^{\text{obs}}, \phi), \boldsymbol{\Gamma}(\phi)\},$$

where $\boldsymbol{\Gamma}(\phi) = [\boldsymbol{\beta}\boldsymbol{\beta}^\top/\sigma^2 + (\nu^2/\tau^2)\mathbf{I}_p + \boldsymbol{\Sigma}_X^{-1}]^{-1}$ and $\tilde{\mathbf{x}}_B(\mathbf{U}^{\text{obs}}, \phi) = [(\mathbf{y}_B - \beta_0 \mathbf{1}_{n_B})\boldsymbol{\beta}^\top/\sigma^2 + (\nu/\tau^2)(\mathbf{w}_B - \psi \mathbf{1}_{n_B} \mathbf{1}_p^\top) + (\mathbf{1}_{n_B} \boldsymbol{\mu}_X^\top) \boldsymbol{\Sigma}_X^{-1}] \boldsymbol{\Gamma}(\phi)$. Note that the mean is an $n_B \times p$ matrix, each row representing the mean vector corresponding to a length- p observation, but the covariance is shared. The imputation is defined only by the likelihood and is common to all methods we consider; the differences lie in the choice of prior $[\phi|\eta]$ and the hyperparameter η . These crucially determine the nature and extent of shrinkage induced on ϕ . In what follows, we propose several options, summarized in Table 2.

VANILLA. As a baseline approach, we apply **DA** to the problem. The choice of prior is

$$(6) \quad [\phi|\eta] \propto (\sigma^2 \tau^2)^{-1} |\boldsymbol{\Sigma}_X^{-1}|^{(2p-1)/2} \exp\left\{-\frac{2p-1}{2} \text{Tr}(\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A]) \boldsymbol{\Sigma}_X^{-1})\right\},$$

TABLE 2

A summary of all Gibbs samplers and choices of priors we considered. Λ is constrained to the class of diagonal matrices. VANILLA and EBSIGMAX require that $p \leq n_A + n_B$

Method	$[\beta \eta] \propto$	$[\Sigma_X^{-1} \eta] \propto \Sigma_X^{-1} ^{(2p-1)/2} \times$	η	Variant
VANILLA	1	$\exp\{-\frac{2p-1}{2} \text{Tr}(\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])\Sigma_X^{-1})\}$	{}	DA
HIERBETAS	$(\frac{\lambda}{\sigma^2})^{p/2} \exp\{-\frac{1}{2}\frac{\lambda}{\sigma^2}\beta^\top\beta\}$	$\exp\{-\frac{2p-1}{2} \text{Tr}(\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])\Sigma_X^{-1})\}$	$\{\lambda\}$	DA+
EBBETAS	$(\frac{\lambda}{\sigma^2})^{p/2} \exp\{-\frac{1}{2}\frac{\lambda}{\sigma^2}\beta^\top\beta\}$	$\exp\{-\frac{2p-1}{2} \text{Tr}(\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])\Sigma_X^{-1})\}$	$\{\lambda\}$	EWiG
EBSIGMAX	1	$ \Lambda ^{3p/2} \exp\{-(1/2) \text{Tr}(\Lambda\Sigma_X^{-1})\}$	$\{\Lambda\}$	EWiG
EBBOTH	$(\frac{\lambda}{\sigma^2})^{p/2} \exp\{-\frac{1}{2}\frac{\lambda}{\sigma^2}\beta^\top\beta\}$	$ \Lambda ^{3p/2} \exp\{-(1/2) \text{Tr}(\Lambda\Sigma_X^{-1})\}$	$\{\lambda, \Lambda\}$	EWiG

where $\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])$ is the diagonal part of the empirical covariance of \mathbf{x}_A . This is a Jeffreys prior on each component of ϕ except Σ_X^{-1} (see Remark 1 below), and η is known. The product of expressions (4) and (6) yields the full conditional distributions of each component of ϕ . For brevity, we present only the Gibbs steps for β and Σ_X^{-1} ; the complete set of full conditional distributions are given in the supplemental article [Boonstra, Mukherjee and Taylor (2013)]:

$$\beta \sim N_p\{(\mathbf{x}_A^\top\mathbf{x}_A + \mathbf{x}_B^\top\mathbf{x}_B)^{-1}(\mathbf{x}_A^\top[\mathbf{y}_A - \beta_0\mathbf{1}_{n_A}] + \mathbf{x}_B^\top[\mathbf{y}_B - \beta_0\mathbf{1}_{n_B}]), \sigma^2(\mathbf{x}_A^\top\mathbf{x}_A + \mathbf{x}_B^\top\mathbf{x}_B)^{-1}\},$$

$$(7) \quad \Sigma_X^{-1} \sim W\{3p + n_A + n_B,$$

$$\begin{aligned} & ((2p - 1) \text{diag}(\widehat{\text{Var}}[\mathbf{x}_A]) + (\mathbf{x}_A - \mathbf{1}_{n_A}\mu_X^\top)^\top (\mathbf{x}_A - \mathbf{1}_{n_A}\mu_X^\top) \\ & + (\mathbf{x}_B - \mathbf{1}_{n_B}\mu_X^\top)^\top (\mathbf{x}_B - \mathbf{1}_{n_B}\mu_X^\top))^{-1}\}. \end{aligned}$$

The Wishart distribution with d degrees of freedom, $W\{d, \mathbf{S}\}$, has mean $d\mathbf{S}$. As made clear in the matrix inversion in (7), VANILLA may only be implemented when $p \leq n_A + n_B$.

REMARK 1. A Jeffreys prior on Σ_X^{-1} , $\Sigma_X^{-1} \sim W\{0, \mathbf{0I}_p\}$, may result in an improper joint posterior if $n_B \gg n_A$ and p is large, that is, when the fraction of missing data is large. From our numerical studies and monitoring of trace plots, even a minimally proper prior on Σ_X^{-1} , that is, using $p + 1$ degrees of freedom, does not ensure a proper posterior. We assume a priori $\Sigma_X^{-1} \sim W\{3p, (2p - 1)^{-1}[\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])]^{-1}\}$, a data-driven choice, the density of which is given in (6). The prior mean of Σ_X^{-1} is $\frac{3p}{2p-1}[\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])]^{-1}$, and the prior mean of Σ_X is $\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])$. Heuristic numeric evidence shows that $3p$ degrees of freedom works well, but we have not demonstrated a theoretical optimality for this. Other values that ensure convergence are equally defensible.

We call the Gibbs sampler which uses this mildly informative prior specification VANILLA. All the other methods we propose will have modified Gibbs steps for two components of $\phi : \beta$ and Σ_X^{-1} . Shrinking β is a clear choice: from (3), β is closely tied to prediction of $Y|\mathbf{X}$. As for Σ_X^{-1} , this determines in part the posterior variance of \mathbf{x}_B (5); as this variance increases, the posterior variance of β decreases (7), thereby shrinking draws β . Other factors in the variance of \mathbf{x}_B , like τ^2 , are additional candidates for shrinkage, but we do not pursue this here.

3.1. *Adaptive prior on β* . Since we are interested in regularizing predictions of the outcome Y , a natural candidate for shrinkage via an informative prior is the parameter vector β , which yields the conditional mean of $Y|\mathbf{X}$. Ridge regression offers favorable predictive capabilities [Frank and Friedman (1993)], and the ℓ_2 penalty on the norm of β is conjugate to the Normal log-likelihood. For these reasons, we replace the Jeffreys prior on β in (6) with

$$(8) \quad [\beta|\sigma^2, \lambda] \propto \left(\frac{\lambda}{\sigma^2}\right)^{p/2} \exp\left\{-\frac{1}{2}\frac{\lambda}{\sigma^2}\beta^\top\beta\right\}.$$

This normal prior on β is analogous to Bayesian ridge regression. λ is a hyperparameter, that is, $\eta = \{\lambda\}$. Conditional upon λ , the Gibbs step for β is

$$\beta \sim N_p\left\{(\mathbf{x}_A^\top\mathbf{x}_A + \mathbf{x}_B^\top\mathbf{x}_B + \lambda\mathbf{I}_p)^{-1}(\mathbf{x}_A^\top\mathbf{y}_A + \mathbf{x}_B^\top\mathbf{y}_B), \sigma^2(\mathbf{x}_A^\top\mathbf{x}_A + \mathbf{x}_B^\top\mathbf{x}_B + \lambda\mathbf{I}_p)^{-1}\right\}.$$

Thus, the posterior mean of β is shrunk toward zero and with smaller posterior variance. As we have outlined in Section 2, there are several options for the treatment of λ .

HIERBETAS. Following Gelfand and Smith (1990), we can treat the hyperparameter λ as random (DA+) with prior distribution $[\lambda] \propto \lambda^{-1}$. Then, we have the following additional posterior step: $\lambda \sim G\{p/2, \beta^\top\beta/(2\sigma^2)\}$. This Bayesian ridge regression with posterior sampling of λ is denoted by HIERBETAS.

EBBETAS. Alternatively, we may apply EWIG to estimate λ . That is, integrate $\log[\beta|\sigma^2, \lambda]$ with respect to the density $[\phi|\mathbf{U}^{\text{obs}}, \lambda]$, differentiate with respect to λ , and solve for λ . The resulting EWIG update is $\lambda \leftarrow p[(1/K)\sum_{k=1}^K \beta^{(k)\top}\beta^{(k)}/\sigma^{2(k)}]^{-1}$. This is a Monte Carlo estimate of $p\{E[(\beta^\top\beta)/(\sigma^2)]\}^{-1}$, the maximum of the marginal likelihood of λ . The update occurs at every K th iteration of the algorithm using the previous K draws of β and σ^2 ; larger values of K yield a more precise estimate. This Bayesian ridge with an Empirical Bayes update of λ is denoted by EBBETAS.

3.2. *Adaptive prior on Σ_X^{-1}* (EBSIGMAX, EBBOTH). We noted previously that an informative prior on Σ_X^{-1} is necessary to ensure a proper joint posterior: $\Sigma_X^{-1} \sim W\{3p, (2p - 1)^{-1}[\text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])]^{-1}\}$, which has inverse scale matrix

$(2p - 1) \text{diag}(\widehat{\text{Var}}[\mathbf{x}_A])$. As we have noted, shrinkage of $\Sigma_{\mathbf{X}}^{-1}$ is closely related to that of β . This was exploited by [Witten and Tibshirani \(2009\)](#) in the SCOUT procedure, suggesting that prediction can be improved through adaptive regularization of $\Sigma_{\mathbf{X}}^{-1}$. Leaving the inverse scale matrix unspecified, the prior is

$$(9) \quad [\Sigma_{\mathbf{X}}^{-1} | \Lambda] \propto |\Lambda|^{3p/2} |\Sigma_{\mathbf{X}}^{-1}|^{(2p-1)/2} \exp\{-(1/2) \text{Tr}(\Lambda \Sigma_{\mathbf{X}}^{-1})\}.$$

Λ is the unknown positive-definite matrix of hyperparameters. The full conditional distribution of $\Sigma_{\mathbf{X}}^{-1}$ becomes

$$(10) \quad \Sigma_{\mathbf{X}}^{-1} \sim W\{3p + n_A + n_B, \\ (\Lambda + (\mathbf{x}_A - \mathbf{1}_{n_A} \mu_{\mathbf{X}}^{\top})^{\top} (\mathbf{x}_A - \mathbf{1}_{n_A} \mu_{\mathbf{X}}^{\top}) \\ + (\mathbf{x}_B - \mathbf{1}_{n_B} \mu_{\mathbf{X}}^{\top})^{\top} (\mathbf{x}_B - \mathbf{1}_{n_B} \mu_{\mathbf{X}}^{\top}))^{-1}\}.$$

Λ may be random or it can be updated with an **EWiG** step. Given the potential difficulty in precisely estimating an unconstrained matrix which maximizes the marginal likelihood, we constrain Λ to be diagonal. Under this constraint, the **EWiG** update for the i th diagonal of Λ is $\Lambda_{ii} \leftarrow 3p((1/K) \sum_{k=1}^K \Sigma_{\mathbf{X}}^{-1(k)}(ii))^{-1}$, where $\Sigma_{\mathbf{X}}^{-1}(ii)$ indicates the i th diagonal element of $\Sigma_{\mathbf{X}}^{-1}$. Then, $\Lambda = \text{diag}\{\Lambda_{11}, \dots, \Lambda_{pp}\}$. This is a Monte Carlo approximation of $3p \text{diag}\{E[\Sigma_{\mathbf{X}}^{-1}]_{11}, \dots, E[\Sigma_{\mathbf{X}}^{-1}]_{pp}\}^{-1}$, the minimizer of $E[\log[\Sigma_{\mathbf{X}}^{-1} | \Lambda]]$ with respect to Λ , subject to the diagonal constraint, with $[\Sigma_{\mathbf{X}}^{-1} | \Lambda]$ as in (9). This approach is denoted as **EBSIGMAX**. Like **VANILLA**, **EBSIGMAX** may only be implemented when $p \leq n_A + n_B$. Finally, let **EBBOTH** be the approach which uses both priors in (8) and (9) with **EWiG** updates for λ and Λ . These alternatives are all summarized in Table 2.

REMARK 2. Adaptively estimating the diagonal inverse scale matrix parameter Λ modifies the variance components of \mathbf{X} . Alternatively, one might apply an **EWiG** update to the degrees of freedom parameter, say, d , which modifies the partial correlations of \mathbf{X} . For example, when $d = p + 1$, the induced prior on each partial correlation is uniform on $[-1, 1]$ [[Gelman and Hill \(2006\)](#)]. Larger values of d place more prior mass closer to zero. Allowing the data to specify d is a reasonable alternative; however, we encountered numerical difficulties in implementing this approach. The **EWiG** update cannot be expressed in closed form and must be estimated numerically. Additionally, the ‘‘complete-data log-likelihood’’ in the M -step is often flat, and a wide range of values for d will return nearly equivalent log-likelihoods.

3.3. *Estimation under predictive loss.* A fitted model may be summarized by measures of uncertainty, for example, a posterior predictive interval $(\hat{Y}_{\text{new}}^{PL}, \hat{Y}_{\text{new}}^{PH})$, as well as point predictions, $\hat{Y}_{\text{new}} = \beta_0^* + \mathbf{X}_{\text{new}}^{\top} \beta^*$ using summary values β_0^*

and β^* . These are calculated with draws from the posterior distribution, $\{\phi^{(t)}\}$. Predictive intervals are given by empirical quantiles of $\{Y_{\text{new}}^{(t)}\}$, where $\hat{Y}_{\text{new}}^{(t)} = \beta_0^{(t)} + \mathbf{X}_{\text{new}}^\top \beta^{(t)} + \sigma^{2(t)} \varepsilon^{(t)}$ and $\varepsilon^{(t)} \stackrel{\text{i.i.d.}}{\sim} N\{0, 1\}$. For point predictions, a summary value of β_0 is given by $\hat{\beta}_0 = (1/T) \sum_t \beta_0^{(t)}$. For β , we minimize posterior predictive loss of \hat{Y}_{new} . Specifically, we define the posterior predictive mean by $\beta^{\text{ppm}} = \arg \min_{\mathbf{b}} E_{\phi, \mathbf{X}_{\text{new}} | \mathbf{U}_{\text{obs}}} (\beta - \mathbf{b})^\top \mathbf{X}_{\text{new}} \mathbf{X}_{\text{new}}^\top (\beta - \mathbf{b})$. This is in contrast to the posterior mean: $\beta^{\text{pm}} = \arg \min_{\mathbf{b}} E_{\beta | \mathbf{U}_{\text{obs}}} (\beta - \mathbf{b})^\top (\beta - \mathbf{b})$. Estimates of these quantities are given by

$$(11) \quad \hat{\beta}^{\text{ppm}} = \sum_t (\Sigma_{\mathbf{X}}^{(t)} + \mu_{\mathbf{X}}^{(t)} \mu_{\mathbf{X}}^{(t)\top})^{-1} \sum_t (\Sigma_{\mathbf{X}}^{(t)} + \mu_{\mathbf{X}}^{(t)} \mu_{\mathbf{X}}^{(t)\top}) \beta^{(t)},$$

$$(12) \quad \hat{\beta}^{\text{pm}} = (1/T) \sum_t \beta^{(t)}.$$

To summarize, different posterior summaries of β come from minimizing different loss functions; we have two estimates of β for each method and, as a consequence, two choices of point predictions for Y_{new} . In contrast, we have only one posterior predictive interval, that derived from the empirical quantiles of $\{Y_{\text{new}}^{(t)}\}$.

4. Simulation study. We conducted a simulation study based upon the motivating data to evaluate these methods. The assumed model of the data satisfied the generating model, as given in (3); violations to these modeling assumptions are considered later. We fixed $n_A = 50$ and $n_B = 400$. The diagonal and off-diagonal elements of $\Sigma_{\mathbf{X}}$ were 1 and 0.15, respectively. The regression coefficients were $\beta = \{\frac{j}{100}\}_{j=-49}^{j=49}$ (a diffuse signal) or $\beta = \{0.1\}_{k=1}^{k=8}, 1\}_{j=1}^{j=11}$ (a signal concentrated in a limited number of coefficients). Values of R^2 were either 0.1 or 0.4. Given β , $\Sigma_{\mathbf{X}}$ and R^2 , σ^2 was determined by solving $\beta^\top \Sigma_{\mathbf{X}} \beta / (\beta^\top \Sigma_{\mathbf{X}} \beta + \sigma^2) = R^2$. β_0 was set to zero. This yielded four unique simulation settings: two choices each for β and R^2 . The covariates \mathbf{x}_A and \mathbf{x}_B were sampled from $N\{\mathbf{0}_p, \Sigma_{\mathbf{X}}\}$, and $\mathbf{y}_A | \mathbf{x}_A$ and $\mathbf{y}_B | \mathbf{x}_B$ were drawn for each combination of β and σ^2 . We set $\psi = 0$ and $\nu = 1$ and repeated each of the four settings for $\tau \in (0, 2)$, drawing $\mathbf{w}_A | \mathbf{x}_A$ and $\mathbf{w}_B | \mathbf{x}_B$, the auxiliary data, based on the measurement error model in (3).

After a burn-in period of 2500, we stored 1000 posterior draws. We calculated $\hat{\beta}_0$, $\hat{\beta}^{\text{ppm}}$ (11) and $\hat{\beta}^{\text{pm}}$ (12). For VANILLA, HIERBETAS, EBBETAS, EBSIGMAX, EBBOTH, we estimated the MSPE using $\hat{\beta}^{\text{ppm}}$ on 1000 new observations: $\widehat{\text{MSPE}}(\hat{\beta}_0, \hat{\beta}^{\text{ppm}}) = (1/1000) \sum_{j=1}^{1000} (Y_{\text{new},j} - \hat{\beta}_0 - \mathbf{X}_{\text{new},j}^\top \hat{\beta}^{\text{ppm}})^2$. $\{Y_{\text{new},j}, \mathbf{X}_{\text{new},j}\}$ are resampled from the same generating distribution for each simulation. As a comparison, we fit a ridge regression (RIDG) on subsample A only, choosing the tuning parameter with the GCV function. Figure 2 plots $\widehat{\text{MSPE}}$, averaged over 250 simulated data sets, over τ . Smaller values are better, and the smallest theoretical value is σ^2 , which is also plotted for reference. We also estimated MSPE using $\hat{\beta}^{\text{pm}}$.

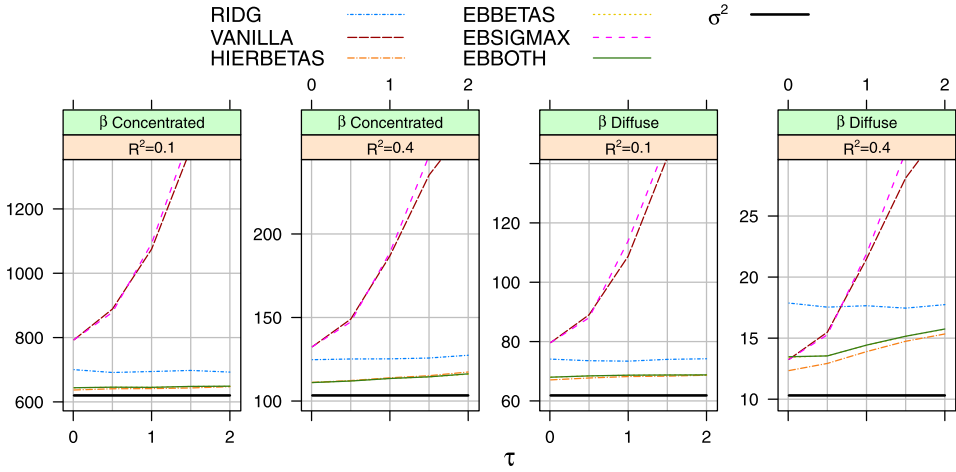


FIG. 2. $MSPE(\hat{\beta}^{ppm})$ plotted against τ , the standard deviation of the ME model, for four simulation settings. For each method, β was estimated from 250 independent training data sets, and MSPE was estimated from 250 validation data sets of size 1000. The thick, solid bar (σ^2) corresponds to predictions made using the true generating parameters. The three best-performing methods, HIERBETAS, EBBETAS and EBBOTH, are virtually indistinguishable.

Numerical values are given in Tables S1 and S2, which also contain results from additional parameter configurations. Finally, we computed prediction intervals for the new observations (Section 3.3). Although frequentist in nature, it is still desirable for a Bayesian prediction interval to achieve nominal coverage; the average coverage rates of $Y_{new,j}$, nominally 95%, are given in Figure 3.

From Figure 2, HIERBETAS, EBBETAS and EBBOTH give about equally good predictions and are consistently the best overall scenarios. EBSIGMAX, which corresponds to shrinkage on Σ_X^{-1} alone, predicts poorly, and VANILLA does only slightly better. RIDG does not beat the better-performing Bayesian methods. Even though the quality of the imputations for x_B depends on the signal in the ME model, the resulting prediction error of HIERBETAS, EBBETAS and EBBOTH varies little over the values of τ we evaluated.

Coverage properties. HIERBETAS, EBBETAS and EBBOTH maintain close-to-nominal prediction coverage (Figure 3). In contrast, larger values of τ drastically decrease the coverage of VANILLA and EBSIGMAX. Prediction intervals for RIDG are not automatic but may be calculated using the bootstrap. This is included in our primary data analysis.

Mean squared error. The results discussed above and reported in Figure 2 use $\hat{\beta}^{ppm}$, which minimizes predictive loss, and are evaluated by MSPE. If instead we use $MSE(\hat{\beta}^{ppm})$ or $MSE(\hat{\beta}^{pm})$, HIERBETAS, EBBETAS and EBBOTH remain the preferred methods (results not given).

Computation time. All Bayesian methods had approximately equal run-times, each requiring about 110 seconds per data set under these simulation settings; run-

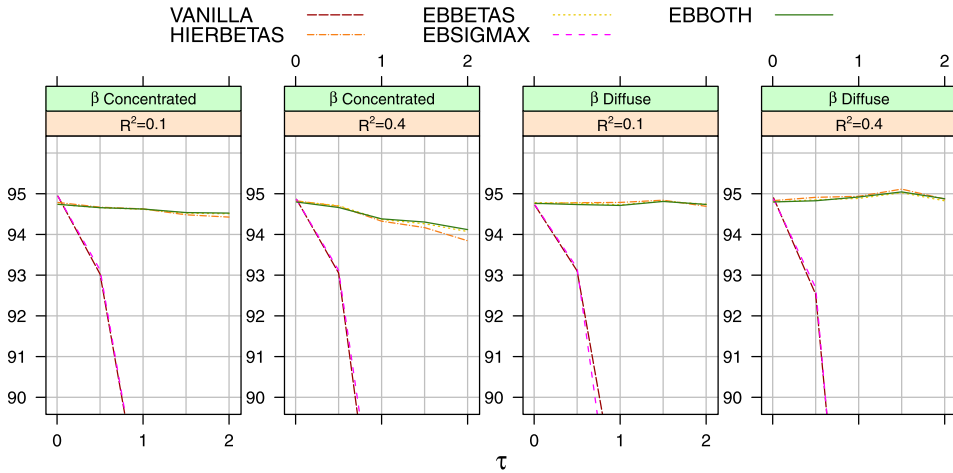


FIG. 3. Average coverage of prediction intervals plotted against τ , the standard deviation of the ME model, for four simulation settings. For each method, prediction intervals were created using draws of β from the converged Gibbs sampler, and coverage was averaged over 250 validation data sets of size 1000. Nominal coverage is 95. The lines for the three methods that are closest to maintaining nominal coverage, HIERBETAS, EBBETAS and EBBOTH, are virtually indistinguishable.

times would increase with p , the dimension of β . While RIDG required only 1–2 seconds for each data set, it does not give automatic prediction intervals, so a direct comparison of run-times here would be improper. In the data analysis (Section 5), we implement a bootstrap algorithm to construct prediction intervals, allowing for a fair comparison of computational time. Full computational details are in the supplemental article [Boonstra, Mukherjee and Taylor (2013)].

Violations to modeling assumptions. As we have noted, these likelihood-based approaches depend on the assumed model approximately matching the true generating model. We evaluated robustness by considering the following violations of the model assumptions: (i) the distribution of ε is skewed, shifted to maintain a zero mean: $\varepsilon + 1 \sim G\{1, 1\}$, (ii) the measurement error model is misspecified $\mathbf{W}|\mathbf{X} \sim N_p\{\psi \mathbf{1}_p + \nu \mathbf{X}^2, \tau^2 \mathbf{I}_p\}$, where we use \mathbf{X}^2 to denote the element-wise square, or (iii) \mathbf{X} comes from a mixture of distributions: $\mathbf{X}|Z \sim N_p\{1_{[Z=2]}(3 \times \mathbf{1}_p) - 1_{[Z=3]}(3 \times \mathbf{1}_p), \Sigma_{\mathbf{X}}\}$, where $1_{[\cdot]}$ is the indicator function and $Z \stackrel{i.i.d.}{\sim} \text{Unif}\{1, 2, 3\}$.

The results of these modeling violations are given in Tables S3–S8. When ε is skewed (S3, S4), the rankings change little; the Bayesian ridge methods are equally preferred. The case is similar for the misspecified measurement error model (S5, S6). When \mathbf{X} comes from a mixture of distributions, the results change depending on whether the signal in β is concentrated (S7) or diffuse (S8). In the former, EBBOTH is best by a large margin for larger values of τ , even over the other Bayesian ridge methods, HIERBETAS and EBBETAS. In this case, then, what is

required is the joint, adaptive shrinkage of $\Sigma_{\mathbf{x}}^{-1}$ and β . This difference in performance is not observed when the signal is diffuse (S8), and the Bayesian ridge methods are all equally good.

A general conclusion of this study is that the shrinkage induced by a Bayesian ridge regression is adaptable to many scenarios and robust to modeling violations. The Gibbs sampler allows for the use of the additional information in subsample B despite \mathbf{x}_B being missing, and the ridge prior on β is effective at controlling variability, thereby increasing precision in predictions. Most important is that this holds even when the signals in the outcome model and the ME model are both very weak, a challenge commonly encountered in the analysis of genomic data.

5. Data analysis. We now consider the motivating problem of efficiently using the auxiliary information in the data from [Chen et al. \(2011\)](#), containing 91 genes representing a broad spectrum of relevant biological functions, to build a predictive model for survival. Expression using Affymetrix is measured on 439 tumors, and qRT-PCR measurements are collected on a subset of 47 of these. Clinical covariates, age, gender and stage of cancer [I–III], are also available. Because qRT-PCR is the clinically applicable measurement for future observations, the goal is a qRT-PCR + clinical covariate model for predicting survival time after surgery. An independent cohort of 101 tumors with qRT-PCR measurements and clinical covariates is available for validation. After some necessary preprocessing of the data, as described in the supplemental article [[Boonstra, Mukherjee and Taylor \(2013\)](#)], the available data had $n_A = 47$, $n_B = 389$, and the validation sample is size 100.

Because our methodology was developed for continuous outcomes, censoring necessitated some adjustments to the data in order to fit our models. We first imputed each censored log-survival time from a linear model of the clinical covariates, conditional upon the censoring time. This model was fit to the training data, but censored survival times in both the training and validation data were imputed. Given completed log-survival times, we refit this same model and calculated residuals from both the training and validation data. These residuals were considered as outcomes, and the question is whether any additional variation in the residuals is explained by gene expression. While there are other ways of dealing with coarsened data and additional covariates in the likelihood-based framework, processing the data this way allows for RIDG to serve as a reference. To more realistically model the data, we allow for a gene-specific ME model: $w_{ij} = \psi_j + v_j x_{ij} + \tau \xi_{ij}$. To incorporate this modification into our model, we put independent flat priors on ψ_j and v_j , $j = 1, \dots, p$. The modified Gibbs steps are included in the supplemental article [[Boonstra, Mukherjee and Taylor \(2013\)](#)].

We applied each Bayesian approach, running each chain of the Gibbs sampler for 4000 iterations and storing posterior draws from the subsequent 4000 iterations. [Table 3](#) presents numerical results: the estimated MSPE from predicting the uncensored residuals in the validation data and the average prediction coverage of

TABLE 3

Results from lung adenocarcinoma analysis. \widehat{MSPE} is the empirical prediction error in the validation data, SIBS is the Scaled Integrated Brier Score, Avg. Coverage is average coverage of the prediction intervals, $\text{Avg}(\hat{Y}_{\text{new}}^{97.5} - \hat{Y}_{\text{new}}^{2.5})$ gives the average prediction interval length for the validation sample, and Computation gives the time, in seconds, to calculate coefficient estimates and prediction intervals

	RIDG	VANILLA	HIERBETAS	EBBETAS	EBSIGMAX	EBBOTH
$\widehat{MSPE}(\hat{\beta}_0, \hat{\beta}^{\text{ppm}})$	0.620	1.251	0.555	0.555	1.230	0.561
$\widehat{MSPE}(\hat{\beta}_0, \hat{\beta}^{\text{pm}})$	–	1.768	0.559	0.558	1.932	0.560
SIBS($\hat{\beta}^{\text{ppm}}$)	0.544	0.629	0.394	0.393	0.632	0.396
SIBS($\hat{\beta}^{\text{pm}}$)	–	0.796	0.395	0.395	0.848	0.395
Avg. Coverage	0.92	0.88	0.96	0.97	0.87	0.96
$\text{Avg}(\hat{Y}_{\text{new}}^{97.5} - \hat{Y}_{\text{new}}^{2.5})$	3.37	3.98	3.11	3.11	3.93	3.09
Computation (sec)	298	268	269	268	269	269

these residuals. Additionally, Table 3 presents the Scaled Integrated Brier Score [SIBS, Graf et al. (1999)], which is a scoring method for right-censored data, on the original, unadjusted validation data.

To calculate the SIBS, which is a function of predicted survival probabilities, we used the survival function from the Normal distribution, estimating the mean log-survival time by adding the linear predictor of the genomic data to the linear predictor of the clinical covariates. At each unique time of last follow-up (either time of death or censoring), the squared difference in predicted survival probability for each individual minus current dead/alive status was calculated and averaged over all individuals and integrated over all time points, with censored individuals only contributing to the calculation of the score until their censoring time. This quantity was scaled by a reference score, that from plugging in 0.5 as a predicted survival probability everywhere, to get the SIBS. Thus, any model that does better than random guessing has a SIBS in the interval (0,1), and a smaller SIBS is better.

Based upon MSPE, HIERBETAS, EBBETAS and EBBOTH were about equally good, with MSPEs of 0.555, 0.555 and 0.561, respectively, using $\hat{\beta}^{\text{ppm}}$. These MSPEs are smaller than those from RIDG (0.620) as well as VANILLA (1.251), and EBSIGMAX (1.230). Using $\hat{\beta}^{\text{pm}}$, the estimated posterior mean of β , the three best methods gave almost identical results, while VANILLA and EBSIGMAX had worse prediction error. Similarly, HIERBETAS, EBBETAS and EBBOTH had the smallest SIBS (resp., 0.394, 0.393 and 0.396), and the remaining methods had larger SIBS.

Considering coverage of the prediction intervals, HIERBETAS (0.96), EBBETAS (0.97) and EBBOTH (0.96) all had rates close to their nominal values, and their prediction intervals widths are smallest. This contrasts with VANILLA and EBSIGMAX, whose coverage rates are less than nominal (0.88, 0.87). We created prediction intervals for RIDG using a bootstrap algorithm; the resulting coverage is 0.92.

The required computational time is 298 seconds for RIDG, including the bootstrap algorithm to calculate prediction intervals, and about 268–269 seconds for each Bayesian method. Although p , n_A and n_B were about the same as in the simulation study, fitting the methods took longer (268 vs. 110 seconds) because the number of total MCMC iterations increased (8000 vs. 3500).

To summarize the analysis thus far, a Bayesian ridge regression, which uses all observations in the data, offers better overall predictive performance in our validation data and, compared to a ridge regression on the complete observations alone, narrower prediction intervals that still achieve nominal coverage. This is a reflection of the extra information that is available in the incomplete observations. Beyond the question of *how* to use the auxiliary genomic information in a prediction model, which has been already been covered in detail, more fundamental to the application is whether one of the Bayesian ridge regressions, for example, EBBETAS, can do better than an analysis using clinical covariates alone, of which complete information is available on all observations. The natural comparison would be an accelerated failure time (AFT) regression, modeling censored log-survival time as a linear function of the clinical covariates and gaussian noise. Predictions from this AFT model could be directly compared to the outcome model in (1).

The SIBS from fitting the AFT model is 0.394, nearly equal to that of EBBETAS. Exploring this comparison further, Figure 4 gives risk-indexed Kaplan–Meier plots of the validation data, comparing predictions using EBBETAS (calculated by adding together the genomic linear predictors to the clinical covariate linear predictors described at the beginning of this section) to that of the AFT model. For each model, patients in the validation sample were indexed based on their predicted survival time: less than 30 months, between 30 and 60 months, or longer than 60 months. From the figure, the clearest distinction is in the low-risk group,

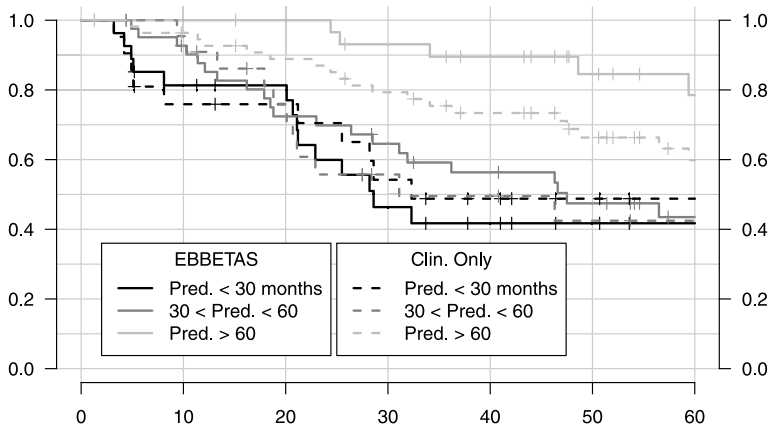


FIG. 4. Comparison of risk-indexed Kaplan–Meier plots. For both EBBETAS and an accelerated failure time model using only the clinical covariates, the validation data was grouped based on predicted survival time (less than 30 months, between 30 and 60 months, and longer than 60 months).

those predicted to live longer than 60 months. In the low-risk, “>60 month” group as defined by EBBETAS, 25 out of 31 patients, or about 80%, were alive at 60 months’ time. This contrasts with the AFT model: 56 patients were predicted to live beyond 60 months, and 36, or about 64%, were alive at 60 months’ time. Also distinctive is that the survival curves for the medium- and high-risk groups of the AFT model cross several times and generally show less separation compared to EBBETAS. The estimated median survival times for these two groups are 28.6 (high) and 47.5 (med.) months under the EBBETAS-based grouping versus 32.3 (high) and 31.1 (med.) under the AFT grouping. Thus, despite nearly equal values of the SIBS, which are aggregate measures of predictive performance, EBBETAS appears to have better individual predictions and discrimination between the three groups.

6. Discussion. Driven by a need to incorporate genomic information into prediction models, we have considered the problem of shrinkage in a model with many covariates when a large proportion of the data are missing. Predictions for future observations are of primary interest. We discuss the primary contributions of this paper in two parts as follows.

6.1. *Shrinkage via the Gibbs sampler.* A likelihood-based approach confers a number of advantages, these being the inclusion of shrinkage into the likelihood and the proper accounting of uncertainty in predictions coming from the unobserved data. A number of existing Bayesian approaches for the treatment of missing data and/or implementation of shrinkage methods are easily adapted here. We have shown how two such approaches, the Monte Carlo EM [Wei and Tanner (1990)], a Gibbs sampler which multiply imputes missing data, and the Empirical Bayes Gibbs Sampler [Casella (2001)], a Gibbs sampler which adaptively shrinks parameter estimates, generalize to the same algorithm, which we call EM-within-Gibbs.

We proposed specific choices of prior specification aimed at improving prediction with shrinkage methods. The various flavors of the Bayesian ridge, denoted as HIERBETAS, EBBETAS and EBBOTH, stand out as the methods of choice, indicating that shrinkage of β , which is the vector of regression coefficients in the outcome model, is most crucial, over and above no shrinkage at all (VANILLA) or shrinkage of Σ_X^{-1} alone (EBSIGMAX). Our simulation study and data analysis showed the Bayesian ridge to be best under a number of scenarios using several criteria, including MSPE and prediction coverage, and robust to several modeling violations. In addition, the Bayesian ridge does not require $p \leq n_A + n_B$, in contrast to VANILLA or EBSIGMAX. As for the specific choice of which Bayesian ridge regression is best, we found little evidence to recommend any one variant.

That shrinkage of Σ_X^{-1} alone, as we have implemented it, does not improve predictions (and sometimes actually worsens predictions) may be due to the specific nature of the shrinkage we implemented. The mean of the conditional distribution

of $\Sigma_{\mathbf{X}}^{-1}$ given in (10) is a convex combination of $\Lambda/(3p)$, which is the inverse of its prior mean, and the sample variance of \mathbf{x}_A and \mathbf{x}_B . In contrast, ridge regression may be viewed as simply adding $\lambda \mathbf{I}_p$ to the sample variance of the covariates. The Wishart prior cannot mimic this effect, and the construction of a different, nonconjugate prior for $\Sigma_{\mathbf{X}}^{-1}$ may be required to induce ridge-type shrinkage.

6.2. *Using genomic information in prediction models.* Figure 5 plots coefficient estimates and 95% credible intervals for the 91 genes according to EBBETAS. They are ordered by the ratio of their posterior mean to posterior standard deviation, an estimate of statistical significance. The ten most significant genes are annotated, according to the R package `annotate` [Gentleman (2012)]. Even the most significant gene, *ERBB3*, is not significant at the 0.05 level. Although these are preselected genes that were deliberately chosen to represent a wide spectrum




FIG. 5. Coefficient estimates (X) and 95% credible intervals ($- -$) of the 91 genes according to EBBETAS, ordered from top to bottom by the magnitude of the ratio of posterior predictive mean to posterior standard deviation. The top ten genes are highlighted and annotated.

of biological functions, many of which have already been implicated in different cancers, this lack of significance for individual genes is not unexpected. The genomic effect is likely to be at the pathway-level rather than individual expressions, which a plot like Figure 5 is too coarse to detect. Despite this lack of individual significance, the small genomic effects collectively yield an overall improvement, albeit small, in predictive ability when the information is properly incorporated, and the Bayesian ridge regression appears best-equipped to do so.

SUPPLEMENTARY MATERIAL

Supplemental article (DOI: [10.1214/13-AOAS668SUPP](https://doi.org/10.1214/13-AOAS668SUPP); .pdf). Here we give the full derivation of the Gibbs steps, computational details and the results from the simulation study. The data from Section 5 and the code for its analysis are available at <http://www-personal.umich.edu/~philb>.

REFERENCES

- BOONSTRA, P. S., MUKHERJEE, B. and TAYLOR, J. M. G. (2013). Supplement to “Bayesian shrinkage methods for partially observed data with many predictors.” DOI:[10.1214/13-AOAS668SUPP](https://doi.org/10.1214/13-AOAS668SUPP).
- BOONSTRA, P. S., TAYLOR, J. M. G. and MUKHERJEE, B. (2013). Incorporating auxiliary information for improved prediction in high-dimensional datasets: An ensemble of shrinkage approaches. *Biostatistics* **14** 259–272.
- CASELLA, G. (2001). Empirical Bayes Gibbs sampling. *Biostatistics* **2** 485–500.
- CHEN, G., KIM, S., TAYLOR, J. M. G., WANG, Z., LEE, O., RAMNATH, N., REDDY, R. M., LIN, J., CHANG, A. C., ORRINGER, M. B. and BEER, D. G. (2011). Development and validation of a qRT-PCR-classifier for lung cancer prognosis. *Journal of Thoracic Oncology* **6** 1481–1487.
- CRAVEN, P. and WAHBA, G. (1979). Smoothing noisy data with spline functions. *Numer. Math.* **31** 377–403.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **39** 1–38. [MR0501537](https://doi.org/10.2307/2346178)
- FRANK, I. E. and FRIEDMAN, J. H. (1993). A statistical view of some chemometrics regression tools. *Technometrics* **35** 109–135.
- GELFAND, A. E. and SMITH, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409. [MR1141740](https://doi.org/10.2307/2289140)
- GELMAN, A. and HILL, J. (2006). *Data Analysis Using Regression and Multilevel Hierarchical Models*. Cambridge Univ. Press, New York.
- GEMAN, S. and GEMAN, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-6** 721–741.
- GENTLEMAN, R. (2012). Annotate: Annotation for microarrays. R package version 1.36.0.
- GRAF, E., SCHMOOR, C., SAUERBREI, W. and SCHUMACHER, M. (1999). Assessment and comparison of prognostic classification schemes for survival data. *Stat. Med.* **18** 2529–2545.
- GREEN, P. J. (1990). On use of the EM algorithm for penalized likelihood estimation. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **52** 443–452. [MR1086796](https://doi.org/10.2307/2346178)
- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* **12** 55–67.

- LITTLE, R. J. A. and RUBIN, D. B. (2002). *Statistical Analysis with Missing Data*, 2nd ed. Wiley, Hoboken, NJ. [MR1925014](#)
- PARK, T. and CASELLA, G. (2008). The Bayesian lasso. *J. Amer. Statist. Assoc.* **103** 681–686. [MR2524001](#)
- TANNER, M. A. and WONG, W. H. (1987). The calculation of posterior distributions by data augmentation. *J. Amer. Statist. Assoc.* **82** 528–550. [MR0898357](#)
- WEI, G. C. G. and TANNER, M. A. (1990). A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *J. Amer. Statist. Assoc.* **85** 699–704.
- WITTEN, D. M. and TIBSHIRANI, R. (2009). Covariance-regularized regression and classification for high dimensional problems. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **71** 615–636. [MR2749910](#)
- YI, N. and XU, S. (2008). Bayesian Lasso for quantitative trait loci mapping. *Genetics* **179** 1045–1055.

DEPARTMENT OF BIOSTATISTICS
UNIVERSITY OF MICHIGAN
1415 WASHINGTON HEIGHTS
ANN ARBOR, MICHIGAN 48109
USA
E-MAIL: philb@umich.edu
bhramar@umich.edu
jmgt@umich.edu